



**Universitat**  
de les Illes Balears

DOCTORAL THESIS

2023

**NOVEL STRATEGIES FOR CONTINUOUS  
GRAVITATIONAL WAVE SEARCHES IN THE ERA  
OF THE ADVANCED DETECTORS**

Rodrigo Tenorio Márquez





**Universitat**  
de les Illes Balears

DOCTORAL THESIS

2023

Doctoral Programme in Physics

**NOVEL STRATEGIES FOR CONTINUOUS  
GRAVITATIONAL WAVE SEARCHES IN THE ERA  
OF THE ADVANCED DETECTORS**

Rodrigo Tenorio Márquez

Thesis Supervisor: Alicia Magdalena Sintés Olives

Thesis Supervisor: David Benjamin Keitel

Thesis Tutor: Carlos Palenzuela Luque

Doctor by the Universitat de les Illes Balears



# Resum

Les ones gravitacionals contínues (CWs, per les sigles en anglès) són formes de radiació gravitacional de llarga durada encara no detectades. S'espera que aquests senyals siguin emesos per estels de neutrons (NSs, per les sigles en anglès) no axisimètrics en rotació a la nostra galàxia, si bé s'han proposat fonts més exòtiques, com ara l'evaporació de núvols de bosons al voltant de forats negres en rotació.

Aquesta tesi presenta tres nous mètodes per post-processar i seguir els resultats de cerques de CW a cegues, així com els resultats de dues cerques de CW a cegues utilitzant dades dels detectors Advanced LIGO. Cadascun d'aquests capítols és una adaptació de dos articles que, respectivament, han estat publicats a la revista *Physical Review D*. El capítol 2, que aporta una descripció general del panorama actual dels mètodes de cerca de CW, és una adaptació d'un article de revisió publicat a la revista *Universe*. Durant el desenvolupament d'aquesta tesi he contribuït directament a un total de 15 articles, els quals es troben publicats a les revistes *Physical Review D*, *Monthly Notices of the Royal Astronomical Society*, *Astrophysical Journal*, i *Astrophysical Journal Letters*. A més, estic involucrat en tres paquets de programari d'anàlisi de dades de codi obert: **LALSuite**, com a col·laborador; **PyFstat**, com a desenvolupador principal tal com es reflecteix en la publicació corresponent al *Journal of Open Source Software*; i **distromax** com a desenvolupador principal i mantenidor. L'impacte d'aquests nous tres mètodes presentats en aquesta tesi es fa palès a en l'ús extensiu d'aquests mètodes en una fracció significativa de les cerques de CW produïdes sobre les dades d'O3 dels detectors Advanced LIGO.

La detecció directa d'una CW ampliaria el nostre coneixement sobre la població de NS galàctica. En primer lloc, les NS són alguns dels objectes observats més extrems de l'Univers perquè contenen una massa entre una i dues vegades la del Sol dins d'un radi d'entre deu i quinze quilòmetres. En conseqüència, són un laboratori ideal per entendre el comportament de la matèria en condicions extremes. L'emissió de CW pot ser conseqüència de diferents processos físics, com per exemple una deformació a les capes exteriors de l'estel o l'oscil·lació de modes inestables a la seva regió interna. La mesura d'aquest senyal ens permetria entendre quin tipus de física nuclear es duu a terme dins aquests estels. En segon lloc, la població d'NS coneguda comprèn aproximadament 3500 estels al moment d'escriure aquesta tesi; la taxa de supernoves de la nostra galàxia, però, suggereix una població total d'entre  $10^8$  i  $10^9$  NS. Aquesta enclotxa demogràfica es podria explicar mitjançant diversos arguments, un dels quals advoca per l'existència d'una població d'NS electromagnèticament inactiva, possiblement amb propietats molt diferents respecte a la població de púlsars observada. Les CW, per tant, podrien ser crucials per a comprendre la totalitat de la població de NS a la nostra galàxia.

Les estratègies de cerca de CW generalment es classifiquen d'acord amb la informació prèvia disponible sobre el tipus de font que cal analitzar. Les cerques dirigides, per una banda, intenten detectar CW de púlsars coneguts, per als quals se suposa que l'emissió de CW es troba en fase amb les observacions electromagnètiques. Això produeix una cerca computacionalment barata, capaç d'extreure informació astrofísicament rellevant d'objectes específics. La física d'aquests objectes, no obstant, és encara incerta. En conseqüència l'assumpció de sincronia entre els senyals electromagnètic i gravitacional podria ser incorrecte amb escenaris realistes i, de forma última, podria impedir la detecció d'una CW.

Les cerques a cegues, per altra banda, no se centren en una font específica sinó que més aviat cerquen un senyal gairebé monocromàtic modulad d'acord amb el corrent Doppler induït pel moviment del detector. Conseqüentment, l'espai de paràmetres a cobrir s'incrementa enormement, tant és així que les cerques han d'utilitzar mètodes subòptims, anomenats *semicoherents*, a fi de ser computacionalment assequibles. Aquesta compensació s'analitza breument al capítol 1. Al capítol 2 es discuteix la implementació de diverses cerques semicoherents utilitzades en l'anàlisi de dades reals. La conclusió general és que les cerques a cegues són les més cares en termes computacionals en el panorama actual a causa de l'extensió de l'espai de paràmetres sota consideració. A causa d'això, aquestes cerques a cegues tendeixen a produir un nombre elevat de candidats que han de ser degudament examinats amb mètodes més sensibles per a descartar candidats d'origen no astrofísic i identificar adequadament candidats consistents en un senyal CW astrofísic.

El primer dels nous mètodes proposats, introduït al capítol 3, se centra en la selecció de candidats interessants produïts per l'etapa principal d'una cerca CW. Concretament, introdueix una nova noció de distància entre senyals CW que és independent de la parametrització del model de senyal. Aquesta distància és utilitzada per agrupar candidats produïts en una cerca CW amb un origen comú, i comprimeix una gran quantitat de resultats en un nombre més petit de grups significatius als quals se'ls pot fer un seguiment més fàcilment. En descriure adequadament l'estructura de l'espai de paràmetres

subjacent, l'ús daquesta distància augmenta la sensibilitat d'una cerca pilot sobre dades d'O2 per a CW procedents de NS desconeguts a sistemes binaris entre un 5 i un 15% respecte a l'ús d'una distància euclidiana.

El capítol 4 analitza com avaluar el resultat d'un seguiment de diverses etapes d'un candidat CW interessant produït per una cerca CW genèrica. En concret, proposa un nou factor de Bayes per establir si el comportament d'un candidat de CW al llarg d'un seguiment per etapes és consistent amb el d'un senyal astrofísic. Per això, es proposen distribucions de dades sota les hipòtesis de senyal i soroll: La distribució sota la hipòtesi de senyal avalua la consistència de les amplituds del senyal al llarg de diferents etapes. La distribució sota la hipòtesi de soroll, per altra banda, empra la teoria de valors extrems per estimar la significança del candidat més significatiu produït per un procés de soroll. L'efectivitat d'aquest nou factor de Bayes s'exemplifica amb l'anàlisi trenta candiats atípics produïts per diverses cerques a les dades obertes del segon període d'observació dels detectors Advanced LIGO (O2). Cap d'aquests valors atípics no resulta consistent amb un senyal astrofísic.

El capítol 5 presenta **distromax**, un nou mètode per estimar la significança del candidat més significatiu produït pel soroll en una cerca d'ones gravitacionals. El mètode aplica resultats de teoria de valors extrems presentats al Capítol 4, els quals es basen en el comportament de la cua de les distribucions de probabilitat involucrades. Conseqüentment, aquest mètode permet la construcció de llindars de detecció fins i tot si es desconeix la distribució subjacent de l'estadística de detecció en ús. Aquesta propietat s'aprofita per avaluar el postprocessament d'una cerca en dades obertes d'O2 amb l'utilització d'una estadística de detecció més robusta contra els artefactes instrumentals. Aquesta estadística no es va utilitzar a la cerca original perquè la seva distribució sota la hipòtesi del soroll és desconeguda; **distromax**, per altra banda, pot ser aplicat sense més problemes.

Aquests mètodes són aplicables a altres tipus de cerques, tal com es discuteix en els seus respectius capítols, i s'han utilitzat de forma rutinària en cerques de CW realitzades per la col·laboració LIGO–Virgo–KAGRA durant el seu tercer període d'observació.

Els capítols 6 i 7 presenten cerques CW a cegues sobre dades del tercer període d'observació dels detectors LIGO–Virgo–KAGRA. El primer s'enfoca a NS en sistemes binaris, mentre que el segon cerca NS aïllats. Aquestes cerques emplearen els mètodes d'anàlisi de dades discutides als capítols anteriors, i obtingueren en conseqüència els resultats més sensibles fins ara als espais de paràmetres analitzats, així com una precisió sense precedents a l'hora de recuperar els paràmetres de senyals generats artificialment.

Tot i la manca de detecció d'un senyal CW, els mètodes desenvolupats en aquesta tesi representen un pas endavant cap a l'anàlisi eficient d'àmplies regions de l'espai de paràmetres. El postprocessament i les estratègies de seguiment presentades serviran com a base per a definir les futures cerques a cegues sobre dades del quart període d'observació de la xarxa avançada de detectors interferomètrics. Aquests mètodes també seran rellevants a mesura que futurs detectors comencin a sondejar freqüències més baixes de l'espectre d'ones gravitacionals, atès que llavors un major nombre de sistemes, com per exemple NS en sistemes binaris o forats negres binaris, produiran senyals d'ones gravitacionals compatibles amb el model CW.

# Resumen

Las ondas gravitacionales continuas (CWs, por sus siglas en inglés) son formas de radiación gravitacional de larga duración aun no detectadas. Se espera que tales señales sean emitidas por estrellas de neutrones (NSs, por sus siglas en inglés) no axisimétricas en nuestra galaxia girando a altas velocidades, si bien fuentes más exóticas se han propuesto en la literatura, como por ejemplo la evaporación de nubes de bosones alrededor de agujeros negros en rotación.

Esta tesis presenta tres nuevos métodos para post-procesar y dar seguimiento a los resultados de búsquedas de CW a ciegas, así como los resultados de dos búsquedas de CW a ciegas utilizando datos de los detectores Advanced LIGO. Cada uno de estos capítulos es una adaptación de un artículo correspondiente publicado en la revista *Physical Review D*. El capítulo 2, que brinda una descripción general del panorama actual de los métodos de búsqueda de CW, es una adaptación de un artículo de revisión publicado en la revista *Universe*. Durante el desarrollo de esta tesis he contribuido directamente a un total de 15 artículos, publicados en *Physical Review D*, *Monthly Notices of the Royal Astronomical Society*, *Astrophysical Journal*, y *Astrophysical Journal Letters*. Además, estoy involucrado en tres paquetes de software de análisis de datos de código abierto: `LALSuite`, como colaborador; `PyFstat`, como desarrollador principal tal y como se refleja en la publicación correspondiente en el *Journal of Open Source Software*; y `distromax` como principal desarrollador y mantenedor. El impacto de los nuevos desarrollos presentados en esta tesis se hace patente en el uso extensivo de estos métodos en una fracción significativa de las búsquedas de CW producidas en los datos de O3 de los detectores Advanced LIGO.

La detección directa de una CW ampliaría nuestro conocimiento sobre la población de NS galáctica. En primer lugar, las NS son algunos de los objetos observados más extremos del Universo, pues contienen una masa entre una y dos veces la del Sol dentro de un radio de entre diez y quince kilómetros. En consecuencia, son un laboratorio ideal para entender el comportamiento de la materia en condiciones extremas. La emisión de CWs sería consecuencia de diferentes procesos físicos, como por ejemplo una deformación en las capas exteriores de la estrella o la oscilación de modos inestables en su región interna. Medir tal señal nos permitiría entender qué tipo de física nuclear está ocurriendo dentro de estas estrellas. En segundo lugar, la población de NS conocida comprende alrededor de 3500 estrellas en el momento de escribir esta tesis; la tasa de supernovas de nuestra galaxia, sin embargo, sugiere una población total de entre  $10^8$  y  $10^9$  NS. Esta brecha demográfica podría explicarse por varios argumentos, uno de los cuales sostiene la existencia de una población de NS electromagnéticamente inactiva, posiblemente con propiedades muy diferentes respecto a la población de púlsares observada. Las CW, por tanto, podrían desempeñar un papel crucial para comprender la totalidad de la población de NS en nuestra galaxia.

Las estrategias de búsqueda de CW generalmente se clasifican de acuerdo a la información previa disponible sobre el tipo de fuente a analizar. Las búsquedas dirigidas, por un lado, intentan detectar CW de púlsares conocidos, para los cuales se supone que la emisión de CW está en fase con las observaciones electromagnéticas. Esto produce una búsqueda computacionalmente barata que es capaz de extraer información astrofísicamente relevante de objetos específicos. No obstante, la física de estos objetos es todavía incierta; en consecuencia, es posible que la asunción de sincronía entre las señales electromagnética y gravitacional sea incorrecta en una situación realista e impeda, en última instancia, la detección de una CW.

Las búsquedas a ciegas, por otro lado, no se centran en una fuente específica; más bien, buscan una señal casi monocromática modulada de acuerdo con el corrimiento Doppler inducido por el movimiento del detector. En consecuencia, el espacio de parámetros a cubrir se incrementa enormemente, tanto así que las búsquedas deben utilizar métodos subóptimos, denominados *semicoherentes*, a fin de ser computacionalmente asequibles. Esta compensación se analiza brevemente en el capítulo 1. En el capítulo 2 se discute la implementación de diversas búsquedas semicoherentes utilizadas en análisis de datos reales. La conclusión general es que las búsquedas a ciegas son las más caras en términos computacionales en el panorama actual debido a la extensión del espacio de parámetros en consideración; en consecuencia, estas búsquedas tienden a producir un número elevado de candidatos que deben ser debidamente examinados utilizando métodos más sensibles a fin de descartar aquellos de origen no astrofísico e identificar adecuadamente aquellos consistentes con una señal CW astrofísica.

El primero de los nuevos métodos propuestos, introducido en el capítulo 3, se centra en la selección de candidatos interesantes producidos por la etapa principal de una búsqueda CW. En concreto, se introduce una nueva noción de distancia entre señales CW que es independiente de la parametrización del modelo de señal. Esta distancia es utilizada para agrupar candidatos producidos en una búsqueda

CW con un origen común, y comprime una gran cantidad de resultados en un número más pequeño de grupos significativos a los que se les puede hacer un seguimiento más fácilmente. Al describir adecuadamente la estructura del espacio de parámetros subyacente, el uso de esta distancia aumenta la sensibilidad de una búsqueda piloto en datos de O2 para CW de NS desconocidas en sistemas binarios entre un 5 y un 15 % con respecto al uso de una distancia euclidiana.

El capítulo 4 analiza cómo evaluar el resultado de un seguimiento de varias etapas de un candidato CW interesante producido por una búsqueda CW genérica. En concreto, propone un nuevo factor de Bayes para establecer si el comportamiento de un candidato de CW a lo largo de un seguimiento por etapas es consistente con el de una señal astrofísica. Para ello, se proponen distribuciones de datos bajo las hipótesis de señal y ruido: la distribución bajo la hipótesis de señal evalúa la consistencia de las amplitudes de la señal a lo largo de diferentes etapas; la distribución bajo la hipótesis del ruido, por otro lado, hace uso de la teoría de valores extremos para estimar la significancia del candidato más significativo producido por un proceso de ruido. La efectividad de este nuevo factor de Bayes se ejemplifica analizando treinta candiadtos atípicos producidos por varias búsquedas en los datos abiertos del segundo periodo de observación de los detectores Advanced LIGO (O2). Ninguno de estos valores atípicos resultó consistente con una señal astrofísica.

El capítulo 5 presenta *distromax*, un nuevo método para estimar la significancia del candidato más significativo producido por el ruido en una búsqueda de ondas gravitacionales. El método aplica resultados de teoría de valores extremos presentados en el Capítulo 4, los cuales se basan en el comportamiento de la cola de las distribuciones de probabilidad involucradas. Consecuentemente, este método permite la construcción de umbrales de detección incluso si se desconoce la distribución subyacente de la estadística de detección en uso. Esta propiedad se aprovecha para evaluar el posprocesamiento de una búsqueda en datos abiertos de O2 utilizando una estadística de detección más robusta contra los artefactos instrumentales. Esta estadística no se utilizó en la búsqueda original ya que su distribución bajo la hipótesis del ruido es desconocida; *distromax*, por otro lado, puede ser aplicado sin mayores problemas.

Estos métodos son aplicables a otros tipos de búsquedas, tal y como se discute en sus respectivos capítulos, y se han utilizado de forma rutinaria en búsquedas de CW realizadas por la colaboración LIGO–Virgo–KAGRA en el tercer periodo de observación.

Los capítulos 6 y 7 presentan búsquedas CW a ciegas utilizando datos del tercer periodo de observación de los detectores LIGO–Virgo–KAGRA. El primero se enfoca en NS en sistemas binarios, mientras que el segundo busca NS aisladas. Estas búsquedas hicieron uso de los métodos de análisis de datos discutidos en los capítulos anteriores, y obtuvieron en consecuencia los resultados más sensibles hasta el momento en los espacios de parámetros analizados, así como una precisión sin precedentes a la hora de recuperar los parámetros de señales generadas artificialmente.

Pese a la falta de detección de una señal CW, los métodos desarrollados en esta tesis representan un paso adelante hacia el análisis eficiente de amplias regiones del espacio de parámetros. El postprocesamiento y las estrategias de seguimiento aquí presentadas servirán como base para definir las futuras búsquedas a ciegas en datos del cuarto periodo de observación de la red avanzada de detectores interferométricos. Estos métodos también serán relevantes a medida que futuros detectores comiencen a sondear frecuencias más bajas del espectro de ondas gravitacionales, ya que entonces un mayor número de sistemas, como NS binarias o agujeros negros binarios, producirán señales de ondas gravitacionales compatibles con el modelo CW.



# Summary

Continuous gravitational waves (CWs) are long-lasting forms of gravitational radiation whose detection is yet to be achieved. The expected sources of such signals are rapidly-spinning non-axisymmetric neutron stars (NSs) within our galaxy, even though more exotic sources, such the evaporation of boson clouds around spinning black holes, have been also considered in the literature.

This thesis presents three new methods to post-process and follow-up the results of all-sky CW searches. Two searches for CWs from isolated and binary unknown NSs conducted on Advanced LIGO data are also presented. Each of these chapters is an adaptation of a corresponding paper published in the journal *Physical Review D*. Chapter 2, which gives an overview of the current landscape of CW search methods, is an adaptation of a review paper published in the journal *Universe*. During the development of this thesis, I have directly contributed to a total of 15 papers, which include publications in *Monthly Notices of the Royal Astronomical Society*, the *Astrophysical Journal*, and the *Astrophysical Journal Letters*. Additionally, I am involved in three open source data-analysis software packages: `LALSuite`, as a contributor; `PyFstat`, as a main developer as reflected in the corresponding publication in the *Journal of Open Source Software*; and `distromax` as main developer and maintainer. The impact of the new developments in this thesis is manifested by the extensive use of these methods in a significant fraction of the CW searches produced in O3 Advanced LIGO data.

A direct detection of a CW would expand our knowledge of the galactic NS population. First, NSs are some of the most extreme observed objects in the Universe, packing a mass one to two times that of the Sun inside a radius of ten to fifteen kilometers. As such, they are an ideal laboratory to understand the behaviour of matter in extreme conditions. The emission of CWs would be a consequence of different physical processes, such as a sustained deformation in the outer layers of the star or the oscillation of unstable modes in its inner region. Measuring such a signal would allow us to understand what sort of nuclear physics are happening inside these stars. Second, the known NS population comprises about 3500 stars at the time of writing; the supernova rate of our galaxy, however, suggests a total population of about  $10^8$  to  $10^9$  NSs. This demographic gap could be explained by several arguments one of which is the existence of an electromagnetically quiet NS population, possibly with vastly different properties with respect to the observed pulsar population. CW signals, thus, could play a crucial role in understanding the entire population of NSs in our galaxy.

Search strategies for CWs are usually classified according to the available prior information on the expected sources. Targeted searches, at one end, aim to detect CWs from known pulsars, for which CW emission is assumed to be phase locked to electromagnetic observations. This yields a computationally cheap search that is able to extract astrophysically relevant information from specific objects. Assuming an exact phase lock between the CW and electromagnetic signals from a pulsar, however, may prevent us from achieving a CW detection, as the physics of such objects are still uncertain and may well be in conflict with such assumption.

Blind searches, more commonly referred to as *all-sky* searches, on the other hand, do not target a specific source; rather, they look for a quasi-monochromatic signal modulated according to the Doppler-shift induced by the detector's movement. Thus, the parameter space to cover is vastly increased, so much so that searches using optimal methods are computationally unaffordable. This is counterbalanced by using non-optimal methods, called *semicoherent*, which reduce the computing cost of a search by using less restrictive signal models. This tradeoff is briefly discussed in Chapter 1. The various implementations of semicoherent searches used in real-data analyses are then discussed in Chapter 2. The overarching conclusion is that all-sky searches are the most expensive kind of search in the current gravitational-wave-search landscape due to the prior volume under consideration; this fact, in turn, tends to produce an elevated number of resulting CW candidates that must be further scrutinized using more sensitive methods in order to rule out non-astrophysical outliers and properly identify those consistent with an astrophysical CW signal.

The first of the newly proposed methods, introduced in Chapter 3, is concerned with the selection of interesting candidates resulting from the main stage of a CW search. Specifically, it introduces a new notion of distance among CW signals that is independent of the chosen parameterisation of the signal model. The distance is then used to cluster nearby candidates produced by a CW search. This strategy compresses a large amount of candidates into a smaller number of meaningful groups that can be more easily followed up. By appropriately capturing the structure of the underlying parameter space,

it increases the sensitivity of a pilot search on O2 data for CWs from unknown NSs in binary systems by about 5 to 15 % with respect to using an ad-hoc Euclidean distance.

Chapter 4 discusses how to evaluate the result of a multi-stage follow-up of an interesting CW candidate produced by a generic CW search. Concretely, it proposes a new Bayes factor to establish whether the behaviour of a CW candidate throughout a series of follow-up stages is consistent with that of an astrophysical signal. To do so, we propose specific data distributions under the signal and noise hypotheses: The distribution under the signal hypothesis follows from first principles, and evaluates the consistency of signal amplitudes across different follow-up stages with different sensitivities; the distribution under the noise hypothesis, on the other hand, makes use of extreme-value-theory results to estimate the expected loudest candidate produced by a background. The effectiveness of this new Bayes factor is demonstrated by analyzing thirty outliers produced by several open-data searches using the second observing run of the Advanced LIGO detectors (O2). None of these outliers was deemed consistent with an astrophysical source.

Chapter 5 presents **distromax**, a new method to estimate the loudest candidate produced by the background in a gravitational-wave search. The method builds on the extreme-value-theory results introduced in Chapter 4, which are based on the tail behaviour of probability distributions. Thus, it allows for the construction of meaningful detection thresholds even if the underlying distribution of the detection statistic at hand is unknown. This property is then exploited to re-evaluate the post-processing of a search on O2 open data to use a detection statistic more robust to instrumental artifacts. This statistic was not used in the original search as its distribution under the noise hypothesis is unknown; **distromax**, on the other hand, can be applied without any major trouble.

These newly developed methods are applicable to other kinds of searches, as discussed in their respective chapters, and have been routinely used in searches for CWs conducted by the LIGO–Virgo–KAGRA collaboration in the third observing run.

Chapters 6 and 7 present blind CW searches for unknown NSs using data from the third observing run of the LIGO–Virgo–KAGRA detectors. The former focuses on NSs in binary systems, while the latter looks for isolated NSs. These searches made use of data-analysis strategies derived from previous chapters to deliver the most sensitive results in the analyzed parameter spaces and an unprecedented precision in the recovery of parameters from artificially-generated signals.

Despite the lack of CW signal detection, the methods developed in this thesis represent a step forward towards the effective analysis of broad parameter-space regions. The post-processing and follow-up strategies here presented will serve as a basis for searching wide parameter-space regions in the forthcoming runs of the advanced network of interferometric detectors. These approaches may also become relevant as future detectors start to probe lower frequencies of the gravitational-wave spectrum, as then a higher number of systems, such as binary NSs or binary black holes, produce gravitational-wave signals compatible with the CW model.

## Publication list

Throughout the development of this thesis I have been involved in the production of the scientific publications listed below. The publications in bold are part of this thesis and have a chapter devoted to them. Amongst them I include two publications I led for the LVK. I mark with a dagger (<sup>†</sup>) those LVK search publications for which I contributed using the data-analysis tools developed in this thesis.

1. *Convolutional neural network search for long-duration transient gravitational waves from glitching pulsars*  
Luana M. Modafferi, Rodrigo Tenorio, David Keitel  
[arXiv:2303.16720 \[astro-ph.HE\]](#) (accepted in Phys. Rev. D)
2. *Search for Gravitational Waves from Scorpius X-1 in LIGO O3 Data With Corrected Orbital Ephemeris*  
John T. Whelan, Rodrigo Tenorio, Jared K. Wofford, James A. Clark, Edward J. Daw, David Keitel, Alicia M. Sintes, Katelyn J. Wagner, Graham Woan, Thomas L. Killestein, Danny Steeghs  
*Astrophys. J.*, **949**, 117 (2023) – [arXiv:2302.10338 \[astro-ph.HE\]](#)
3. *Prospects for detecting transient quasi-monochromatic gravitational waves from glitching pulsars with current and future detectors*  
Joan Moragues, Luana M. Modafferi, Rodrigo Tenorio, David Keitel  
*Mon. Not. Roy. Astron. Soc.* **519** (2023) 4, 5161-5176 – [arXiv:2210.09907 \[astro-ph.HE\]](#)  
DOI: 10.1093/mnras/stac3665
4. **Empirically estimating the distribution of the loudest candidate from a gravitational-wave search**  
Rodrigo Tenorio, Luana M. Modafferi, David Keitel, Alicia M. Sintes  
*Phys. Rev. D* **105**, 044029 (2022) – [arXiv:2111.12032 \[gr-qc\]](#)  
DOI: 10.1103/PhysRevD.105.044029
5. **All-sky search for continuous gravitational waves from isolated neutron stars using Advanced LIGO and Advanced Virgo O3 data**  
R. Abbott et al. (LIGO–Virgo–KAGRA Collaboration)  
*Phys. Rev. D* **106**, 102008 (2022) – [arXiv:2201.00697 \[gr-qc\]](#)  
DOI: 10.1103/PhysRevD.106.102008
6. <sup>†</sup>*Model-based cross-correlation search for gravitational waves from the low-mass X-ray binary Scorpius X-1 in LIGO O3 data*  
R. Abbott et al. (LIGO–Virgo–KAGRA Collaboration)  
*Astrophys. J. Lett.*, **941**, L30 (2022) – [arXiv:2209.02863 \[astro-ph.HE\]](#)  
DOI: 10.3847/2041-8213/aca1b0
7. <sup>†</sup>*Search for continuous gravitational wave emission from the Milky Way center in O3 LIGO–Virgo data*  
R. Abbott et al. (LIGO–Virgo–KAGRA Collaboration)  
*Phys. Rev. D* **106**, 042003 (2022) – [arXiv:2204.04523 \[astro-ph.HE\]](#)  
DOI: 10.1103/PhysRevD.106.042003
8. <sup>†</sup>*Narrowband searches for continuous and long-duration transient gravitational waves from known pulsars in the LIGO–Virgo third observing run*  
R. Abbott et al. (LIGO–Virgo–KAGRA Collaboration)  
*Astrophys. J.*, **932**, 133 (2022) – [arXiv:2112.10990 \[gr-qc\]](#)  
DOI: 10.3847/1538-4357/ac6ad0
9. <sup>†</sup>*Search for continuous gravitational waves from 20 accreting millisecond x-ray pulsars in O3 LIGO data*  
R. Abbott et al. (LIGO–Virgo–KAGRA collaboration)  
*Phys. Rev. D* **105**, 022002 (2022) – [arXiv:2109.09255 \[astro-ph.HE\]](#)  
DOI: 10.1103/PhysRevD.105.022002

10. **Search methods for continuous gravitational-wave signals from unknown sources in the advanced-detector era**  
Rodrigo Tenorio, David Keitel, Alicia M. Sintes  
[Universe 2021, 7\(12\), 474](#) – [arXiv:2111.12575 \[gr-qc\]](#)  
DOI: 10.3390/universe7120474
11. † *All-sky search for continuous gravitational waves from isolated neutron stars in the early O3 LIGO data*  
R. Abbott et al. (LIGO–Virgo–KAGRA collaboration)  
[Phys. Rev. D 104, 082004 \(2021\)](#) – [arXiv:2107.00600 \[gr-qc\]](#)  
DOI: 10.1103/PhysRevD.104.082004
12. **Application of a hierarchical MCMC follow-up to Advanced LIGO continuous gravitational-wave candidates**  
Rodrigo Tenorio, David Keitel, Alicia M. Sintes  
[Phys. Rev. D 104, 084012 \(2021\)](#) – [arXiv:2105.13860 \[gr-qc\]](#)  
DOI: 10.1103/PhysRevD.104.084012
13. *PyFstat: a Python package for continuous gravitational-wave data analysis*  
David Keitel, Rodrigo Tenorio, Gregory Ashton, Reinhard Prix  
[J. open source softw. 6 \(2021\) 60, 3000](#) – [arXiv:2101.10915 \[gr-qc\]](#)  
DOI: 10.21105/joss.03000
14. **All-sky search in early O3 LIGO data for continuous gravitational-wave signals from unknown neutron stars in binary systems**  
R. Abbott et al. (LIGO–Virgo Collaboration)  
[Phys. Rev. D 103, 064017 \(2021\)](#) – [arXiv:2012.12128 \[gr-qc\]](#)  
DOI: 10.1103/PhysRevD.103.064017
15. **Time-frequency track distance for comparing continuous gravitational wave signals**  
Rodrigo Tenorio, David Keitel, Alicia M. Sintes  
[Phys. Rev. D 103, 064053 \(2021\)](#) – [arXiv:2012.05752 \[gr-qc\]](#)  
DOI: 10.1103/PhysRevD.103.064053

# Acknowledgements

I wish to thank my advisors, Prof. Alicia M. Sintes and Dr. David Keitel, for their support during the development of this thesis. I would also like to thank Prof. Sascha Husa for his support on the proper use of computing resources.

I thank Prof. Graham Woan and Dr. Chris Messenger for their support during my visit at the University of Glasgow, and Dr. Andrew Spencer and Dr. Jennifer Wright for their hospitality. I am gratefully to Prof. Ian Jones, Dr. Andrew Lundgren, and Dr. Laura Nuttall for their support during my brief visit at the University of Southampton and the University of Portsmouth. Also, I would like to thank the CW working group of the LIGO-Virgo-KAGRA collaboration for many useful discussions throughout the development of this thesis.

I thank Rafel, Pep, and Dr. Miquel Oliver for their spiritual guidance before, during, and beyond this thesis. I am grateful to Michael, Joe, and the significant population of PhD students and early career scientists gravitating around the common room of the Kelvin building, for their comradery and scientific insight. I thank Amanda and Catalina for comments on some language issues.

I acknowledge support from the Spanish Ministerio de Ciencia, Educación y Universidades (refs. FPU 18/00694 and EST21/00643), the Universitat de les Illes Balears (UIB), the Spanish Ministry of Science and Innovation (MCIN) and the Spanish Agencia Estatal de Investigación (AEI) grants PID2019-10 6416GB-I00/MCIN/AEI/10.13039/501100011033, the MCIN with funding from the European Union NextGenerationEU (PRTR-C17.I1), the FEDER Operational Program 2021-2027 of the Balearic Islands, the Comunitat Autònoma de les Illes Balears through the Direcció General de Política Universitaria i Recerca with funds from the Tourist Stay Tax Law ITS 2017-006 (PRD2018/24, PDR2020/11), the Conselleria de Fons Europeus, Universitat i Cultura del Govern de les Illes Balears, and EU COST Actions CA18108 and CA17137. This thesis has been assigned document number LIGO-T2300155.





# List of Abbreviations

<b>ATNF</b>	Australia Telescope National Facility
<b>BOINC</b>	Berkeley Open Infrastructure for Network Computing
<b>CE</b>	Cosmic Explorer
<b>CW</b>	Continuous wave
<b>ET</b>	Einstein Telescope
<b>HESS</b>	High Energy Stereoscopic System
<b>GC</b>	Galactic center
<b>GCT</b>	Global correlation transform
<b>GR</b>	General relativity
<b>GW</b>	Gravitational wave
<b>JKS</b>	Jaranoswki, Królak, and Schutz [ <a href="#">Phys. Rev. D 58 063001 (1998)</a> ]
<b>KAGRA</b>	Kamioka Gravitational Wave Detector
<b>LIGO</b>	Laser Interferometry Gravitational-waves Observatory
<b>LMXB</b>	Low-mass X-ray binary
<b>LUT</b>	Look-up table
<b>LVC</b>	LIGO–Virgo Collaboration
<b>LVK</b>	LIGO–Virgo–KAGRA Collaboration
<b>MCMC</b>	Markov-chain Monte-Carlo
<b>NS</b>	Neutron star
<b>O2</b>	Second observing run of the Advanced LIGO, and Advanced Virgo detectors
<b>O3</b>	Third observing run of the Advanced LIGO, Advanced Virgo, and KAGRA detectors
<b>O3a</b>	First half of the third observing run of the Advanced LIGO and Advanced Virgo detectors
<b>PSD</b>	Power spectral density
<b>SFT</b>	Short Fourier Transform
<b>SNR</b>	Signal to noise ratio, Supernova remnant
<b>SOAP</b>	Snakes on a plane
<b>SSB</b>	Solar system barycenter
<b>tCW</b>	Transient continuous wave
<b>TT</b>	Transverse-traceless
<b>TTD</b>	Time-frequency track distance





*Que hi hagi pau damunt la Terra*



# Contents

<b>Resum</b>	<b>iii</b>
<b>Resumen</b>	<b>v</b>
<b>Summary</b>	<b>vii</b>
<b>Publication list</b>	<b>ix</b>
<b>Acknowledgements</b>	<b>xi</b>
<b>List of abbreviations</b>	<b>xiii</b>
<b>Contents</b>	<b>xvii</b>
<b>1 A brief introduction to continuous gravitational-wave searches</b>	<b>1</b>
1.1 Neutron stars as continuous-wave sources . . . . .	1
1.2 Continuous waves from spinning NSs . . . . .	2
1.3 Basics of detecting continuous waves . . . . .	3
1.3.1 Continuous-wave searches as Bayesian hypothesis testing . . . . .	3
1.3.2 How to claim a detection . . . . .	6
1.3.3 Multi-hypothesis testing . . . . .	7
1.3.4 Semicoherent searches . . . . .	9
1.4 Toward realistic continuous-wave searches . . . . .	11
<b>2 Search methods for continuous gravitational-waves from unknown sources in the advanced-detector era</b>	<b>17</b>
2.1 Introduction . . . . .	17
2.2 Wide parameter-space search pipelines . . . . .	19
2.2.1 $\mathcal{F}$ -statistic searches . . . . .	20
GCT hierarchical search . . . . .	20
Weave . . . . .	21
Time-domain $\mathcal{F}$ -statistic . . . . .	21
2.2.2 Fourier-transform-based searches . . . . .	22
PowerFlux & Falcon . . . . .	22
Cross-Correlation . . . . .	23
2.2.3 Hough-transform semicoherent searches . . . . .	23
SkyHough . . . . .	23
FrequencyHough . . . . .	24
2.2.4 Viterbi searches . . . . .	24
2.2.5 Machine learning . . . . .	25
2.3 Post-processing strategies . . . . .	26
2.3.1 Coincidences . . . . .	26
2.3.2 Parameter-space clustering . . . . .	26
2.3.3 Detector-consistency vetoes . . . . .	27
2.3.4 $\chi^2$ vetoes . . . . .	28
2.3.5 Vetoing narrow spectral features . . . . .	28
2.3.6 Null-hypothesis vetoes . . . . .	29
2.4 Follow-up . . . . .	29

2.4.1	Single-stage follow-up . . . . .	30
2.4.2	Multi-stage follow-up . . . . .	30
2.5	Upper bounds on $h_0$ . . . . .	31
2.6	Summary . . . . .	32
<b>3</b>	<b>Time-frequency track distance for comparing continuous gravitational wave signals</b>	<b>47</b>
3.1	Introduction . . . . .	47
3.2	Parameter-space distances and clustering algorithms . . . . .	48
3.2.1	Grid-based clustering methods . . . . .	48
3.2.2	Machine Learning-based clustering . . . . .	49
3.3	A new parameter space distance . . . . .	50
3.3.1	Deriving a distance from the $\mathcal{F}$ -statistic . . . . .	50
3.3.2	Distance implementation . . . . .	51
3.3.3	Discussion . . . . .	52
3.4	Application to clustering for an all-sky binary CW search . . . . .	55
3.4.1	Algorithm overview . . . . .	55
3.4.2	Reachable sets . . . . .	56
3.4.3	Coincidental sets . . . . .	56
3.5	Improving the sensitivity of all-sky searches . . . . .	56
3.5.1	The O2 data software injection campaign . . . . .	56
3.5.2	TTD: Clustering performance . . . . .	58
3.5.3	TTD: Sensitivity improvement . . . . .	59
3.6	Conclusion . . . . .	61
<b>4</b>	<b>Application of a hierarchical MCMC follow-up to Advanced LIGO continuous gravitational-wave candidates</b>	<b>67</b>
4.1	Introduction . . . . .	67
4.2	Continuous-wave data analysis: Search and follow up . . . . .	68
4.2.1	Coherent and semicoherent searches . . . . .	69
4.2.2	MCMC-based follow-ups . . . . .	69
4.2.3	A coherence-time ladder . . . . .	70
4.3	Evaluating the hierarchical follow-up with a Bayes factor . . . . .	72
4.3.1	Noise hypothesis . . . . .	72
4.3.2	Signal hypothesis . . . . .	74
4.3.3	Bayes factor . . . . .	75
4.4	Follow-up of outliers in LIGO O2 data . . . . .	77
4.4.1	Continuous-wave search outliers from O2 data . . . . .	77
	All-sky <b>Falcon</b> search . . . . .	77
	Directed <b>Einstein@Home</b> search . . . . .	79
	Fomalhaut b <b>Viterbi</b> search . . . . .	79
	H.E.S.S. <b>Viterbi</b> search . . . . .	79
4.4.2	Follow-up setup . . . . .	79
4.5	Results . . . . .	82
4.5.1	Injections in Gaussian noise . . . . .	82
4.5.2	Follow-up of CW outliers from Advanced LIGO O2 data . . . . .	85
4.6	Conclusion . . . . .	88
	Appendices	
4.A	On the distribution of the maximum $\mathcal{F}$ -statistic and the effective number of templates . . . . .	91
<b>5</b>	<b>Empirically estimating the distribution of the loudest candidate from a gravitational-wave search</b>	<b>103</b>
5.1	Introduction . . . . .	103
5.2	Continuous wave searches . . . . .	104
5.2.1	$\mathcal{F}$ -statistic under the noise hypothesis . . . . .	105
5.2.2	Template-bank correlations . . . . .	106
5.2.3	$\mathcal{F}$ -statistic under the signal hypothesis . . . . .	106

5.3	Loudest candidates, the “effective number of templates”, and extreme value theory . . . .	107
5.4	How to estimate the distribution of the loudest outlier: an empirical approach . . . . .	109
5.4.1	Basic formulation . . . . .	110
5.4.2	Introducing <b>distromax</b> . . . . .	111
5.4.3	The MaxProp operator . . . . .	112
5.4.4	Characterizing the batchmax distribution . . . . .	112
5.4.5	Parameter estimation accuracy and comparison to previous approaches . . . . .	114
5.4.6	Discussion . . . . .	117
5.5	Application to O2 data . . . . .	118
5.5.1	CW detection statistics . . . . .	118
5.5.2	tCW detection statistics . . . . .	119
5.6	Conclusion . . . . .	120
Appendices		
5.A	Basic results of extreme value theory . . . . .	121
5.B	Addressing disturbed data . . . . .	125
5.C	Robustness against injections . . . . .	128
<b>6</b>	<b>All-sky search in early O3 LIGO data for continuous gravitational-wave signals from unknown neutron stars in binary systems</b>	<b>139</b>
6.1	Introduction . . . . .	139
6.2	Signal Model . . . . .	140
6.3	Data used . . . . .	141
6.4	The Search Pipeline . . . . .	142
6.4.1	Ranking statistics . . . . .	143
6.4.2	Toplist construction . . . . .	144
6.5	Post Processing . . . . .	145
6.5.1	Clustering . . . . .	145
6.5.2	Line Veto . . . . .	145
6.6	Sensitivity . . . . .	146
6.7	Follow Up . . . . .	150
6.7.1	MCMC follow-up configuration . . . . .	150
	Sampler configuration . . . . .	150
	Prior choice . . . . .	151
6.7.2	Setting up a threshold . . . . .	151
6.7.3	Surviving Outliers . . . . .	152
	Line-crossing outliers . . . . .	152
	Detector consistency veto . . . . .	153
	Powerline sidebands . . . . .	154
6.8	Conclusion . . . . .	154
Appendices		
6.A	Frequency bands containing outliers . . . . .	155
<b>7</b>	<b>All-sky search for continuous gravitational waves from isolated neutron stars using Advanced LIGO and Advanced Virgo O3 data</b>	<b>161</b>
7.1	Introduction . . . . .	161
7.2	Data sets used . . . . .	162
7.3	Common Aspects of Search Pipelines . . . . .	162
7.3.1	Signal model . . . . .	162
7.3.2	Parameter space analyzed . . . . .	163
7.3.3	Detection statistics . . . . .	163
7.3.4	Outlier follow-up . . . . .	163
7.3.5	Upper limits . . . . .	165
7.4	The SkyHough search . . . . .	165
7.4.1	Parameter space . . . . .	165
7.4.2	Description of the search . . . . .	165

7.4.3	Candidate follow-up . . . . .	167
7.4.4	Sensitivity estimation . . . . .	171
7.5	Conclusions . . . . .	175
<b>8</b>	<b>Conclusion and future developments</b>	<b>181</b>
<b>A</b>	<b>Coauthors' authorization</b>	<b>187</b>

## Chapter 1

# A brief introduction to continuous gravitational-wave searches

Gravitational waves (GWs) are traveling perturbations of the spacetime metric predicted by General Relativity (GR) [1, 2] that were finally detected in 2015 by the LIGO Scientific Collaboration (LSC) [3]. The implications of such an initial discovery and the following ones during the several observing runs of the LIGO–Virgo–KAGRA Collaboration (LVK) are still lively discussed in the current scientific literature.

To date, the international network of interferometric GW detectors is composed of the two Advanced LIGO detectors [4], located in the USA; the Advanced Virgo detector [5], located in Italy; the GEO600 detector [6], located in Germany; and the KAGRA detector [7], in Japan. Another interferometric detector is currently under construction in India [8]. In conjunction, the detectors in this network have been operative for a total of three observing runs, with an increasing sensitivity due to progressive technological upgrades. After the first detection, and at the time of this writing, these detectors have achieved about 100 confident detections of GW signals produced in the coalescence of compact binary systems, such as those formed by black holes or neutron stars [9]. Among these events we can highlight the first multi-messenger observation of the coalescence of a binary neutron star system [10], and the first observation of a neutron-star-black-hole coalescence [11].

At the time of writing, the LIGO–Virgo–KAGRA detector network is about to start the fourth observing run, with a planned duration of about two years. This is projected to be the longest and most-sensitive run ever performed by the current generation of detectors. Upon completion, the advanced detectors will undergo an upgrade process for two years in order to get ready for the fifth observing run, which is currently projected for the end of this decade. In the new decade, the third generation of ground-based detectors, the Einstein Telescope (ET) [12] and Cosmic Explorer (CE) [13], is expected to enter in operation. The improvements in sensitivity will allow for significant increase in the number of GW events, and will likely bring the detection of new GW signals, such as those produced by rapidly-spinning neutrons stars.

This chapter introduces the search for continuous waves (CWs). These undetected signals are produced by long-duration quadrupolar variations, such as those produced by a non-axisymmetric spinning body. As discussed in Sec. 1.1, the expected sources for such signals are non-axisymmetric rapidly-spinning neutron stars (NSs), and their amplitude is expected to be several orders of magnitude weaker than that of a binary-black-hole coalescence. Their duration, however, is such that they will be present throughout the entire duration of an observing run (which tend to have a duration on the order of years). These two properties complement each other, and allow for these signals to be detectable if enough data is integrated.

## 1.1 Neutron stars as continuous-wave sources

Neutron stars (NSs) are one of the possible results of the core-collapse supernovae of stars with masses between 8 to 30 times that of the Sun. They were theoretically predicted in the early twentieth century by Walter Baade and Fritz Zwicky [14] (and, in a sense, by Lev Landau around the same period [15]) and finally discovered in the decade of the 60s, both as pulsating radio sources [16] and accreting from a binary companion [17]. To date, about 3500 NSs have been discovered using electromagnetic means (see [18, 19] for an updated catalogue).

A significant fraction of the known NS population has been entirely observed by electromagnetic means as pulsating stars denominated *pulsars* [20]. They typically have a mass of about 1.4 solar masses and a radius of 10 to 15 kilometers. The origin of the observed pulsation is related to the strong dipolar magnetic field sustained by these objects, which leads to the emission of radiation beams from their magnetic poles that can be observed from the Earth if they are properly oriented. The observed spin frequencies range from fractions of a hertz to about 700 Hz, being some of the most stable emitters ever observed.

The structure of NSs is still under discussion [21]. It is generally accepted that NSs are composed of an outer crust of ordinary atomic nuclei, and a core of neutron-rich matter. The physics of the core are currently poorly understood: plausible models include several phases of nuclear matter (typically referred to as “nuclear pasta”), strange matter, or quark-degenerate matter among others.

NSs can emit GW signals through a variety of mechanisms [22]. In this thesis we shall focus on CW emission due to the presence of “mountains”, which are parameterised by the equatorial ellipticity of the star

$$\varepsilon = \frac{I_1 - I_2}{I_3}. \quad (1.1)$$

Here,  $I_3$  is the moment of inertia with respect to the spinning axis of the star and  $I_{1,2}$  are the moments of inertia with respect to two orthogonal directions. This quantity is proportional to the quadrupolar moment of the star, which, as we shall see later in this section, is the ultimate cause of gravitational radiation. Two types of mountains are usually considered, namely magnetic mountains, sustained due to magnetic fields [23, 24, 25], and thermal mountains, produced as a consequence of accretion from a companion star [26, 27]. The former kind can take place in a general NS; the latter, however, is restricted to those in a binary system, as it requires accretion for the build-up of a temperature anisotropy. Specific values of  $\varepsilon$  are highly dependent on the specifics of the proposed mechanism. Latest results in the literature, however, argue that the maximum sustainable ellipticity by a conventional NS is about  $\varepsilon \approx 7.4 \times 10^{-6}$  [28, 29]. If exotic matter is included, this maximum ellipticity can reach values on the order of  $10^{-4}$  [30].

Detecting CW signals will lead to the understanding of different characteristics of the galactic NS population. Aside from obtaining direct measurements of  $\varepsilon$ , which will allow to further probe the structure of these objects, we will also be able to extend our knowledge in the demography of NSs. First, estimations of the rate of core-collapse supernovae in our galaxy suggest it should contain about  $10^8$  NSs,  $10^5$  of which should be detectable as pulsars [31, 32, 33]; as previously mentioned, only about 3500 have been observed to date [19]. The new generation of radiotelescopes (e.g. [34]) is expected to increase this number, but among the unknown NSs there is still the possibility to find a subpopulation whose properties diverge from those of pulsars, such as gravitars [35], for which CWs may be the most favorable channel to detect them. CW, thus, may be a *required* observational channel to observe the complete population of galactic NS. A second issue is related to the spinning frequency of the observed NS population, which is well below the Kepler limit beyond which NSs become structurally unstable [36]. The fastest-spinning observed NS are *millisecond* pulsars, which are believed to be the result of a slowly-rotating NSs spinning up due the accretion of material from a companion star. It is argued that such a spin-up process is counteracted by the emission of CW due to the formation of thermal mountains in the NS, which yields binary NSs as plausible sources of CWs [37, 38, 39, 40, 41].

## 1.2 Continuous waves from spinning NSs

CW signals emitted from a NS can be described using linearized gravity [42], which treats GWs as traveling linear perturbations of the metric in a flat spacetime. Despite its simplicity, this approximation has provided a convincing solution to a series of problems throughout twentieth-century physics, such as the precession of Mercury’s perihelion [43]. Later in the century, the quadrupolar approximation was used to obtain indirect evidence of GW emission by a binary system of compact objects [44].

The metric perturbation  $h$  due to a quadrupolar acceleration in a source located at a distance  $d$  from an observer in the linearized theory of gravity is given by

$$h = \frac{1}{d} \frac{2G}{c^4} \ddot{Q}, \quad (1.2)$$



where  $G$  is the gravitational constant,  $c$  is the speed of light in vacuum and  $\ddot{Q}$  is the second time derivative of the quadrupole moment of the source. The corresponding polarizations, for the case of CW emission due to mountains, are given by [45]

$$h_+ = h_0 \frac{1 + \cos^2 \iota}{2} \cos(2\pi f_{\text{gw}} t), \quad (1.3)$$

$$h_\times = h_0 \cos \iota \sin(2\pi f_{\text{gw}} t), \quad (1.4)$$

where  $\cos \iota$  is the angle between the star's spinning axis and the line-of-sight,

$$h_0 = \frac{4\pi^2 G}{c^4} I f_{\text{gw}}^2 \frac{\varepsilon}{d} \approx 1.1 \times 10^{-24} \left( \frac{\varepsilon}{10^{-6}} \right) \left( \frac{I_3}{10^{38} \text{kg m}^2} \right) \left( \frac{f_{\text{gw}}}{1 \text{kHz}} \right)^2 \left( \frac{d}{1 \text{kpc}} \right)^{-1}, \quad (1.5)$$

and the GW frequency is twice the star's rotational frequency  $f_{\text{gw}} = 2f_{\text{rot}}$ . Other mechanisms, such as mass-current quadrupoles or free precession, yield emission of CWs at other multiples of the rotational frequency [46]. Hence, the frequency of a CW signal alone could already provide valuable information about the physics of the source if the rotational frequency of the source is known [47].

Due to the emission of electromagnetic or gravitational radiation, a NS will tend to diminish its rotational frequency. The specific rate at which this happens is characterized by the so-called *braking index*  $n$ , defined such that  $\dot{f}_{\text{gw}} \propto f_{\text{gw}}^n$ . The value of  $n$  characterizes the dominant channel of emission of a star (e.g.  $n = 3$  for magnetic dipolar radiation,  $n = 5$  for mass quadrupolar radiation) [48]. For sufficiently old NSs in the LIGO detector band, however, it suffices to consider a Taylor-expansion model for the frequency evolution

$$f_{\text{gw}} = f_0 + f_1(t - t_0) + \dots \quad (1.6)$$

where the specific number of spin-down terms is dependent upon the age and spinning frequency of the NS population being searched; usually, one or two spin-down terms are enough for most wide parameter-space searches [49, 50].

## 1.3 Basics of detecting continuous waves

Searching for CW signals involves an intricate interplay amongst probability theory, the properties of the signal itself upon its arrival to the detector, and the response of the detector itself to the signal. This section intends to give a brief introduction to the search for CW from unknown sources from the practical point of view. To do so, we shall simplify the general problem to that of detecting a monochromatic signal in Gaussian noise. The problem itself is unchanged, as we will still be able to set up a template bank, evaluate detection statistics and select interesting candidates in the same fashion as an actual CW search<sup>1</sup>. The use of purely monochromatic signals allows us to postpone the discussion of the signal's properties upon its arrival to the detector (and the response of the detector itself) to Sec. 1.4.

### 1.3.1 Continuous-wave searches as Bayesian hypothesis testing

The analysis of a monochromatic signal in Gaussian noise using Bayesian probability is thoroughly discussed in [51]. We shall follow some of the basic steps therein presented in order to draw a clear picture of the sort of analyses performed in a CW search. Let us start describing a monochromatic signal with a certain frequency  $f$  (in Hertz) as

$$h(t; f, \mathcal{A}) = \mathcal{A}_0 \cos(2\pi ft + \varphi_0) = \mathcal{A}_c \cos(2\pi ft) + \mathcal{A}_s \sin(2\pi ft). \quad (1.7)$$

$\mathcal{A}_0$  refers to the signal's amplitude and  $\varphi_0$  is the initial phase of the signal. It will be convenient to parameterise this signal in terms of zero-offset sinusoidal functions, for which we introduce the corresponding cartesian variables  $\mathcal{A}_c = \mathcal{A}_0 \cos \varphi_0$  and  $\mathcal{A}_s = -\mathcal{A}_0 \sin \varphi_0$ . Note that  $\mathcal{A}_0 = \sqrt{\mathcal{A}_c^2 + \mathcal{A}_s^2}$ . Throughout this discussion we will collectively refer to all amplitude-related quantities as  $\mathcal{A}$  regardless of whether we use  $(\mathcal{A}_0, \varphi_0)$  or  $(\mathcal{A}_c, \mathcal{A}_s)$ .

<sup>1</sup>In fact, the analysis here presented is equivalent to the first step of other, more complicated methods such as [49].

The standard situation in CW searches describes a data stream  $x$  consisting of a signal such as the one in Eq. (1.7) buried in zero-mean Gaussian noise

$$x(t) = n(t) + h(t; f, \mathcal{A}). \quad (1.8)$$

We assume data is taken without interruption for a time  $T$  at discrete timestamps  $t_j = j\Delta t$ , each labeled by  $j = 0, \dots, M-1$  (note that  $t_0 = 0$  without loss of generality.). Discretely-sampled quantities will be labeled using subindices such that  $x_j = x(t_j)$ . The collection of all quantities will be represented by an index-less variable  $x = \{x_j, j = 0, \dots, M-1\}$ .

The noise stream will be described as an uncorrelated and homoscedastic Gaussian process with standard deviation  $\sigma$ :

$$p(n_j|\sigma) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{1}{2}\left(\frac{n_j}{\sigma}\right)^2\right), \quad (1.9)$$

$$p(n|\sigma) = \prod_{j=0}^{M-1} p(n_j|\sigma) = \left(\frac{1}{\sqrt{2\pi}\sigma}\right)^M \exp\left(\frac{1}{2\sigma^2} \sum_{j=0}^{M-1} n_j^2\right). \quad (1.10)$$

In a realistic situation, the detector's noise would be non-stationary and non-Gaussian; within relatively short timescales, however, it would be characterized by its Power Spectral Density (PSD)  $S_n$  and overall results would be computed by weighing different time periods by their corresponding PSD. This quantity is related to  $\sigma$  by the Wiener-Khinchin theorem [49, 52]

$$\sigma^2 = \frac{1}{2\Delta t} S_n, \quad (1.11)$$

which allows for the treatment of  $\sigma$  and  $S_n$  on equal grounds, so the derivation here presented can trivially be re-stated using either of these quantities. It remains for us to appropriately deal with  $\sigma$ . The Bayesian way of proceeding is to marginalize Eq. (1.10) with respect to  $\sigma$  using a certain prior distribution  $p(\sigma)$ . Ref. [53] approached the problem by splitting the overall datastream into segments with similar data quality, so that segment-wise homoscedasticity could be assumed, and marginalizing each segment separately using Jeffrey's prior  $p(\sigma) \sim \sigma^{-1}$ , obtaining as a result a Student's t distribution. Similar distributions were obtained through a different reasoning in [54, 55]. Finally, [49, 52] assume  $\sigma$  to be known (which corresponds to marginalizing with a delta prior), and rely on the use of ad-hoc weights to compensate for the lack of homoscedasticity. Since marginalizing approximately amounts to estimating unknown quantities using the available data, these two approaches should behave similar given enough data. We take this last stance, which amounts to considering  $\sigma$  as a known, fixed parameter; we keep  $\sigma$  implicit in our notation from this point onwards for the sake of simplicity.

Since  $x(t) - h(t; f, \mathcal{A}) = n(t)$ , the distribution of the data under a specific signal hypothesis is simply

$$p(x|f, \mathcal{A}) = \left(\frac{1}{\sqrt{2\pi}\sigma}\right)^N \exp\left(-\frac{1}{2\sigma^2} \sum_{j=0}^{M-1} (x_j - s_j(\mathcal{A}, f))^2\right). \quad (1.12)$$

It will come in handy to define a scalar product in a similar manner to [56]<sup>2</sup>,

$$\langle a, b \rangle = \sum_{j=1}^N a_j b_j, \quad (1.13)$$

so that

$$\sum_{j=0}^{M-1} (x_j - s_j(\mathcal{A}, f))^2 = -2\langle x, s(\mathcal{A}, f) \rangle + \langle x, x \rangle + \langle s(\mathcal{A}, f), s(\mathcal{A}, f) \rangle. \quad (1.14)$$

---

<sup>2</sup>One could include  $\sigma$  or other quantities in this definition, as originally done in [56]; in this case, however,  $\sigma$  is simply a constant overall factor and can be kept out of the definition.

The signal model at hand [Eq. (1.7)] will lead to the manipulation of summations involving sines and cosines. Let us open up that discussion by quoting the following result, which follows from the summation of a convergent series:

$$\sum_{j=0}^{M-1} e^{iAj} = \frac{\sin(\frac{1}{2}AM)}{\sin(\frac{1}{2}A)} e^{i\frac{1}{2}A(M-1)}, \quad (1.15)$$

In the case at hand,  $A = 2\pi\Delta tf$ . Now, if we introduce  $T = M\Delta t$  and apply some basic trigonometric manipulation, we arrive at

$$\sum_{j=0}^{M-1} \cos(2\pi f\Delta tj) = \Re \sum_{j=0}^{M-1} e^{i2\pi\Delta tfj} = \frac{1}{2} \sin(2\pi Tf) \cot(\pi\Delta tf) + \frac{1}{2} [1 - \cos(2\pi Tf)] \leq \frac{1}{2} \cot(\pi\Delta tf) \quad (1.16)$$

and

$$\sum_{j=0}^{M-1} \sin(2\pi f\Delta tj) = \Im \sum_{j=0}^{M-1} e^{i2\pi\Delta tfj} = \frac{1}{2} [1 - \cos(2\pi Tf)] \cot(\pi\Delta tf) + \frac{1}{2} \sin(2\pi Tf) \leq \frac{1}{2} \cot(\pi\Delta tf), \quad (1.17)$$

where  $\Re$  and  $\Im$  denote the real and imaginary part of a complex number, respectively. The typical regime of a CW search will be such that  $M \gg 1$ ; as a result, we will remain at  $\mathcal{O}(M)$  and use Eq. (1.16) and Eq. (1.17) to neglect lower-order terms. More specifically, we will be concerned with frequencies  $1 \text{ Hz} \lesssim f \lesssim 1 \text{ kHz}$  and sampling frequencies  $\Delta t^{-1} \simeq 1 \times 10^{-4} \text{ Hz}$  (e. g.  $\Delta t^{-1} \approx 16 \text{ kHz}$  for Advanced LIGO [4]), which implies  $f\Delta t \lesssim 1$ . Also, we will assume  $T$  to be long enough so all the frequencies of interest are properly resolved,  $Tf \gg 1$  or, equivalently,  $M \gg (\Delta tf)^{-1}$ . Under these assumptions, cotangents can be Laurent-expanded as

$$\cot(x) = \frac{1}{x} - \mathcal{O}(x), x \lesssim 1, \quad (1.18)$$

and the leading order in both Eq. (1.16) and Eq. (1.17) is  $(\Delta tf)^{-1} \ll \mathcal{O}(M)$ ; hence, sums over sines or cosines alone can be neglected againsts  $\mathcal{O}(M)$  terms. This results hold as well if we duplicated the frequency, as both conditions  $2f\Delta t \lesssim 1$  and  $2fT \gg 1$  would still be fulfilled.

With these results in hand, the scalar product of sine and cosine functions can be easily expanded using the half-angle formula:

$$\langle \cos(2\pi ft), \cos(2\pi ft) \rangle = \frac{M}{2} + \frac{1}{2} \sum_{j=0}^{M-1} \cos(2\pi 2f\Delta tj) \simeq \frac{M}{2}, \quad (1.19)$$

$$\langle \sin(2\pi ft), \sin(2\pi ft) \rangle = \frac{M}{2} - \frac{1}{2} \sum_{j=0}^{M-1} \cos(2\pi 2f\Delta tj) \simeq \frac{M}{2}, \quad (1.20)$$

$$\langle \sin(2\pi ft), \cos(2\pi ft) \rangle = \frac{1}{2} \sum_{j=1}^N \sin(2\pi 2f\Delta tj) \simeq 0. \quad (1.21)$$

Thus,

$$\langle s(\mathcal{A}, f), s(\mathcal{A}, f) \rangle = \frac{M}{2} \mathcal{A}_0^2 \quad (1.22)$$

$$\langle x, s(\mathcal{A}, f) \rangle = \mathcal{A}_c x_c + \mathcal{A}_s x_s \quad (1.23)$$

where we defined

$$x_c = \langle x, \cos(2\pi ft) \rangle, \quad (1.24)$$

$$x_s = \langle x, \sin(2\pi ft) \rangle. \quad (1.25)$$

for later convenience.

The presence of a signal, in the context of Bayesian probability, is measured by a Bayes factor.

This quantity weights the support of a given dataset for either of two hypotheses. For the case of pure Gaussian noise versus a signal with amplitude  $\mathcal{A}$  and frequency  $f$  being present in addition to the noise, the Bayes factor is defined as

$$\mathcal{B}(x; \mathcal{A}, f) = \frac{p(x|\mathcal{A}, f)}{p(x|\mathcal{A}_0 = 0)} = \exp \left[ -\frac{1}{2} \frac{M}{2\sigma^2} (\mathcal{A}_c^2 + \mathcal{A}_s^2) + (\mathcal{A}_c x_c + \mathcal{A}_s x_s) \right], \quad (1.26)$$

where we use the fact that the noise hypothesis [Eq. (1.10)] is equivalent to a signal with amplitude  $\mathcal{A}_0 = 0$ . Note that Eq. (1.26) refers explicitly to a specific signal with amplitude parameters  $\mathcal{A}$  and frequency  $f$ . In a realistic situation, we rarely have access to the specific amplitude of the signal being searched; rather, we are interested in a *plausible range* of amplitude values. In a Bayesian context, this sort of information can be taken into account by marginalizing the signal hypothesis with respect to a prior on  $\mathcal{A}$ <sup>3</sup>

$$\mathcal{B}(x; f) = \int d\mathcal{A} p(\mathcal{A}) \mathcal{B}(x; \mathcal{A}, f). \quad (1.27)$$

This integral can be analytically approximated by a Gaussian integral if a uniform prior on  $\mathcal{A}$  is used, following a similar motivation to that of CW searches [58, 59].

Choosing a uniform prior on  $\mathcal{A}$  in this case, however, may be problematic, as the amplitude has no obvious upper bound and thus the probability distribution cannot be properly normalized.<sup>4</sup> This problem can be circumvented at the expense of assuming an arbitrary upper bound  $\mathcal{A}_*$

$$p(\mathcal{A}_0, \varphi_0) = \frac{1}{2\pi} \frac{1}{\pi \mathcal{A}_*^2}, \quad \mathcal{A}_0 \leq \mathcal{A}_*, \quad (1.28)$$

so that Eq. (1.29) is well approximated by a Gaussian integral

$$\mathcal{B}(x; f) = \frac{1}{\pi \mathcal{A}_*^2} \int_0^{\mathcal{A}_*} d\mathcal{A}_0 \mathcal{B}(x; \mathcal{A}, f) \approx \frac{\sigma}{\mathcal{A}_*^2} \sqrt{\frac{8}{M}} \exp \left[ \frac{1}{2} \frac{2\sigma^2}{M} (x_c^2 + x_s^2) \right] \quad (1.29)$$

The constant  $\mathcal{A}_*$  can be interpreted as an Occam factor (also known as a “trials factor”): Integrating a higher prior volume considers a higher number of hypotheses (“trials”) which diminishes the “significance” of a specific result with respect to using a smaller prior volume. This sort of constant, derived from the use of improper priors, is innocuous as long as only two classes of hypotheses are considered [60].

It is instructive to realize that the only quantity dependent on the data in Eq. (1.29) is the quadrature in the exponential term

$$x_c^2 + x_s^2. \quad (1.30)$$

Such a quantity corresponds to the squared modulus of the Fourier transform of the data [51] (for a suitable definition of Fourier transform). For the sake of simplicity, we will work in terms of a re-scaled quantity

$$\bar{s}(f) = \frac{2}{N\sigma^2} (x_c^2 + x_s^2). \quad (1.31)$$

### 1.3.2 How to claim a detection

The conclusion of Sec. 1.3.1 is that Eq. (1.31) is the appropriate quantity to use in order to assess the presence of a monochromatic signal with a certain frequency  $f$  in the data. We now tackle the problem of deciding a minimum value of  $\bar{s}$  to justify such a claim, which can be formalized in terms of decision theory [61]. Simply put, deciding on whether a signal is actually in the data or not based on the value of a Bayes factor (or a monotonic function of it) incurs on a certain risk of committing a mistake. We will mainly be concerned with two possible situations: either claiming the presence of a signal when the data contains only noise (i.e. a “false alarm”), or failing to claim the presence of a signal (i.e. a “false dismissal”). Both of these situations have an associated probability,  $p_{\text{fa}}$  and  $p_{\text{fd}}$  respectively, which allow us to tune the performance of the decision.

<sup>3</sup>Note that, since the noise hypothesis is independent of the signal parameters, the marginalization integral can be directly written as an integral over the Bayes factor. This is the reason the MCMC in [57] can be used in terms of Bayes factors rather than likelihoods.

<sup>4</sup>This would also occur if we used the alternative parameterisation  $\{\mathcal{A}_c, \mathcal{A}_s\}$ , as discussed in [51].

The standard procedure is to construct a threshold on the detection statistic  $\tilde{s}_t$  so that the false alarm probability is fixed at a specific value  $p_{\text{fa}}$ . Any data scoring above that threshold is deemed inconsistent with the noise hypothesis and subject to further analysis. This stems from the fact that a Bayes factor is the optimal detection statistic (i.e. it minimizes  $p_{\text{fd}}$  at a given  $p_{\text{fa}}$ ) provided that the amplitude prior  $p(\mathcal{A})$  is consistent with the population of signals being searched for [62, 63]. Different populations would be described by different priors and, in turn, would derive different optimal detection statistics.

A characterization of the statistical properties of  $\tilde{s}$  is explicitly given in [64, 65]. In short, under the noise hypothesis,  $x_c$  and  $x_s$  are the real and imaginary parts of a linear combination of Gaussian random variables with zero mean and unit variance, and as a result have a variance of  $\sigma^2 M/2$  each. Consequently,  $\tilde{s}$  is the sum of two squared Gaussian variables with zero mean and unitary variance, which yields a chi-squared distribution with two degrees of freedom

$$\tilde{s}|\text{Noise} \sim \chi_2^2. \quad (1.32)$$

Adding a signal with amplitude  $\mathcal{A}_0$  and following the same argument implies both  $x_c$  and  $x_s$  pick up a mean value of  $\mathcal{A}_c M/2$  and  $\mathcal{A}_s M/2$ , respectively. As a result,  $\tilde{s}$  is in this case a non-central chi-squared distribution with two degrees of freedom

$$\tilde{s}|\text{Signal} \sim \chi_2^2(\rho^2), \quad (1.33)$$

where the non-centrality parameter is given by

$$\rho^2 = \frac{M}{2} \left( \frac{\mathcal{A}_0}{\sigma} \right)^2. \quad (1.34)$$

This quantity is referred to as the (squared) signal-to-noise ratio (SNR) in the context of a CW search.

The performance of a detection statistic can be stated in terms of the minimum amplitude that can be distinguished at a certain false-alarm and false-dismissal level. Computing such a number involves solving the following implicit equation in  $\rho$

$$p_{\text{fd}} = p(\tilde{s} \leq \tilde{s}_t(p_{\text{fa}}) | \rho^2), \quad (1.35)$$

where

$$p_{\text{fa}} = p(\tilde{s} \geq \tilde{s}_t | \text{Noise}). \quad (1.36)$$

We postpone any discussion on the results of such an equation until Sec. 1.3.4, where an approximated analytical expression is given for  $\rho$ .

### 1.3.3 Multi-hypothesis testing

After reviewing the basics of signal detection, we are in a good position to discuss one of the characteristic difficulties of blind CW searches: to detect a monochromatic signal with an *unknown* frequency. The Bayesian tenet in this case is no different than for the amplitude parameters  $\mathcal{A}$ : Eq. (1.29) should be marginalized over  $f$  to construct a global Bayes factor testing the presence of any signal within the support of the priors. Such an integral cannot be usually expressed in closed form. Instead, the standard approach is to place a grid along the considered range of  $f$  and set an appropriate threshold to select those points scoring a high detection statistic value. We will use a simple criterion, based on estimating the loudest noise candidate of a search, for the sake of simplifying our exposition. In Chapter 2 we will review several of the approaches used in real-data searches.

Probing multiple frequencies is akin to generating several draws from the background. The higher the number of draws, the more likely we are to obtain an extreme event, hence the less “surprised” we should be by it<sup>5</sup>. This is related to the false-alarm probability of a search, which, for a given threshold value, tends to *increase* as the number of probed frequencies increases. In order to maintain a reasonable false-alarm probability, we ought to appropriately choose a threshold depending on the overall number of trials of our analysis.

The first question we should be answering, thus, is “How many different trials are we doing?”. The approach in the CW literature is to consider two frequencies as “different” if the loss in the detection

<sup>5</sup>Note that this is exactly the same interpretation we gave to the Occam factor  $\mathcal{A}_*$  in the previous subsection.

statistic  $\tilde{s}$  we incur in by mistakenly looking for one if the other one is present is sufficiently large. This allows for the construction of a characteristic frequency resolution to count the number trials in a certain frequency range.

Assume our data consists of a noiseless monochromatic signal with a certain frequency  $f_0$ :

$$x(t) = \mathcal{A}_c \cos(2\pi f_0 t) + \mathcal{A}_s \sin(2\pi f_0 t). \quad (1.37)$$

We are interested in the behaviour of  $\tilde{s}(f; f_0)$ , which corresponds to computing  $\tilde{s}$  for a frequency  $f$  on a datastream containing a signal with frequency  $f_0$ . First, let us note that

$$\tilde{s}(f_0; f_0) = \frac{2}{N\sigma^2} (\mathcal{A}_c^2 \langle \cos(2\pi f_0 t), \cos(2\pi f_0 t) \rangle^2 + \mathcal{A}_s^2 \langle \sin(2\pi f_0 t), \sin(2\pi f_0 t) \rangle^2) = \frac{M}{2} \left( \frac{\mathcal{A}_0}{\sigma} \right)^2 \quad (1.38)$$

by direct application of Eq. (1.21), which corresponds to the optimal SNR introduced in Eq. (1.34). A convenient choice to quantify the loss of a detection statistic due to mismatching a parameter is to use the relative loss, typically referred to as *mismatch* [66, 67, 68]

$$\tilde{\mu}(f; f_0) = 1 - \frac{\tilde{s}(f; f_0)}{\tilde{s}(f_0; f_0)}. \quad (1.39)$$

Let  $\delta f = f - f_0$ . The relevant quantities in this case will be

$$\langle \cos(2\pi f t), \cos(2\pi f_0 t) \rangle = \frac{1}{2} \sum_{j=0}^{M-1} [\cos(2\pi \delta f \Delta t j) + \cos(2\pi(f + f_0) \Delta t j)] \simeq \frac{1}{2} \sum_{j=0}^{M-1} \cos(2\pi \delta f \Delta t j) \quad (1.40)$$

and

$$\langle \sin(2\pi f t), \sin(2\pi f_0 t) \rangle = \frac{1}{2} \sum_{j=0}^{M-1} [\cos(2\pi \delta f \Delta t j) - \cos(2\pi(f + f_0) \Delta t j)] \simeq \frac{1}{2} \sum_{j=0}^{M-1} \cos(2\pi \delta f \Delta t j), \quad (1.41)$$

where  $\delta f$  is assumed to be small enough so that  $f + f_0$  behaves similarly to  $2f$  and thus can be neglected to order  $\mathcal{O}(M)$ . Specifically, we want to consider the case  $\delta f T \sim 1$  such that the frequency difference is properly resolved throughout the observing run. Equation (1.16) still applies, but we need to reconsider which terms can be neglected:

$$\sum_{j=0}^{M-1} \cos(2\pi \delta f \Delta t j) = \frac{1}{2} \sin(2\pi T \delta f) \cot(\pi \Delta t \delta f) + \sin^2(\pi T \delta f). \quad (1.42)$$

First, the squared sine is  $\mathcal{O}(1)$ , which can be neglected. The cotangent can still be Laurent-expanded since  $\Delta t \delta f \sim \mathcal{O}(M^{-1})$  and  $M \gg 1$ , but cannot be neglected to order  $\mathcal{O}(M)$ , yielding a cardinal sine function

$$\sum_{j=0}^{M-1} \cos(2\pi \delta f \Delta t j) \simeq M \operatorname{sinc}(2\pi T \delta f). \quad (1.43)$$

Therefore,

$$\tilde{s}(f; f_0) = \frac{M}{2} \left( \frac{\mathcal{A}_0}{\sigma} \right)^2 \operatorname{sinc}^2(2\pi T \delta f), \quad (1.44)$$

which is consistent with a similar calculation in [68]. The mismatch, thus, results in

$$\tilde{\mu}(f; f_0) = 1 - \operatorname{sinc}^2(2\pi T \delta f) = \frac{4\pi^2 T^2 \delta f^2}{3} + \mathcal{O}(T^4 \delta f^4) \quad (1.45)$$

where in the last Taylor expansion we assumed the probed frequency  $f$  was close to the signal's value  $f_0$ . The lack of a linear term in  $\delta f$  implies that  $f = f_0$  corresponds to a minimum mismatch value; i. e.  $\rho^2$

is the maximum value attainable by  $\tilde{s}$ . This last result also allows us to define a typical resolution

$$\delta f = \frac{\sqrt{3\tilde{\mu}}}{2\pi T}. \quad (1.46)$$

If more parameters were considered, Eq. (1.45) would be a quadratic form in the parameters' mismatch, and the corresponding matrix would be the parameter-space metric. Diagonalizing and inverting such a matrix would return the corresponding parameter resolutions. Note that Eq. (1.46) is independent of the parameter  $f$  itself. Counting the number of templates is thus a trivial operation in this case

$$\mathcal{N} = \left\lceil \frac{\Delta f}{\delta f} \right\rceil, \quad (1.47)$$

where  $\Delta f$  is the width of the frequency band to search over and  $\lceil \cdot \rceil$  represents the ceiling function.

Once the number of independent trials  $\mathcal{N}$  is known, we are interested in estimating the probability distribution of  $\tilde{s}^* = \max_{\mathcal{N}} \tilde{s}$  under the noise hypothesis. An in-depth discussion of such a topic is given in Chapter 5, and we limit ourselves here to the final result

$$\tilde{s}^* | \text{Noise} \sim \text{Gumbel}(\mu_{\mathcal{N}}, \sigma_{\mathcal{N}}), \quad (1.48)$$

$$\mu_{\mathcal{N}} = 2 \ln \mathcal{N}, \quad (1.49)$$

$$\sigma_{\mathcal{N}} = 2, \quad (1.50)$$

where the specific expressions for  $\mu_{\mathcal{N}}$  and  $\sigma_{\mathcal{N}}$  follow from the fact that  $\tilde{s}$  is a chi-squared distribution with two degrees of freedom under the noise hypothesis. The expected value of  $\tilde{s}^*$  is the expected value of the aforementioned Gumbel distribution, which is given by

$$\tilde{s}_t = 2(\ln \mathcal{N} + \gamma), \quad (1.51)$$

where  $\gamma \approx 0.577 \dots$  is the Euler-Mascheroni constant. We will use this quantity as a threshold in order to estimate the sensitivity of a broad parameter space search.

It is interesting to note that the concept of false-alarm probability does not explicitly appear in the construction of (1.51). We can use such a threshold to account for the large trials factor of a search, yet the specific false alarm at which a search operates is never explicitly stated. As explained in Chapter 2 and references therein, this is an ongoing trend whenever conducting a CW search; rarely, if ever at all, the global false alarm at which a pipeline operates is stated.

### 1.3.4 Semicohherent searches

Suppose now that we had access to a datastream with a duration  $TN_{\text{seg}}$  such that the number of templates as computed by Eq. (1.47) was too big to be feasible to evaluate  $\tilde{s}$  over the whole set. This is the standard situation in CW searches, as in such a case the number of templates scales with a large power of the run's duration time.

The standard approach is a so called *semicoherent* search [69], which consists in splitting the data into  $N_{\text{seg}}$  independent segments and computing  $\tilde{s}$  in each of them, to then combine these values into a new *semicoherent* statistic

$$\hat{s} = \sum_{n=0}^{N_{\text{seg}}-1} \tilde{s}^{(n)}. \quad (1.52)$$

From a Bayesian perspective, this corresponds to marginalizing with respect to  $\mathcal{A}$  parameters in each of the segments independently. As a result, the family of models considered by this detection statistic is less restrictive than a coherent search, as amplitude needs not to be consistent across segments. One should also note that this choice is not motivated by statistical optimality but, as we shall see, by computational efficiency.




FIGURE 1.1: Sensitivity estimates for different search setups. For all the cases in these panels,  $\Delta t = 10^{-4}$  s,  $\rho = 1$ , and  $p_{\text{fd}} = 0.05$ . The coherent search uses  $T = N_{\text{seg}} \cdot 1000$  s; the semicoherent search uses  $T = 1000$  s and the number of segments specified in the horizontal axis. Left panel: Single-template search using  $p_{\text{fa}} = 10^{-3}$  to derive a threshold. Scaling lines are only a visual aid. Right panel: Search over  $\Delta f = 0.1$  Hz using  $\tilde{\mu} = \hat{\mu} = 1$  to compute  $\mathcal{N}$ . The threshold corresponds to the expected loudest outlier over  $\mathcal{N}$ . Dashed lines, corresponding to the results in the left panel, are shown as a visual aid.

We proceed in the same manner as in Secs. 1.3.2 and 1.3.3. We start by studying the semicoherent mismatch. It is easy to show that

$$\hat{s}(f; f_0) = \sum_{n=0}^{N_{\text{seg}}-1} \tilde{s}^{(n)}(f; f_0) = N_{\text{seg}} \tilde{s}(f; f_0), \quad (1.53)$$

which implies, introducing a corresponding semicoherent mismatch,

$$\hat{\mu}(f; f_0) = 1 - \text{sinc}^2(2\pi T \delta f). \quad (1.54)$$

This means the number of templates required to run a semicoherent search is independent of the number of segments used to run the search<sup>6</sup>. One can split the data into  $N_{\text{seg}}$  segments in such a way that the resulting segment duration provides a manageable number of templates in order to tune the computing cost of a search. The downside of this approach is related to the optimal squared SNR

$$\hat{s}(f_0; f_0) = \sum_{n=0}^{N_{\text{seg}}-1} \tilde{s}^{(n)}(f_0; f_0) = N_{\text{seg}} \rho^2, \quad (1.55)$$

which increases *linearly* with the number of segments, rather than quadratically, as it would be the case in a coherent search. A semicoherent search is, thus, a tradeoff between constructing the most sensitive search and being actually able to run it in an affordable time.

The statistics of  $\hat{s}$  are similar to those of  $\tilde{s}$ . Both in the noise and signal case  $\hat{s}$  is distributed as a chi-squared distribution with  $2N_{\text{seg}}$  degrees of freedom, as now Gaussian variables are summed in quadrature for each data segment. Also, as shown in Eq. (1.55), the non-centrality parameter is linear rather than quadratic in  $N_{\text{seg}}$ . If we follow the simple approach outlined in [70] and assume  $N_{\text{seg}} \gg 1$ ,  $\hat{s}$  is distributed following a Gaussian distribution. As a result, we obtain the equivalent of Eq. (1.35)

$$p_{\text{fd}} = \frac{1}{2} \text{erfc} \left( \frac{N_{\text{seg}} \rho^2 - 2\hat{z}_t \sqrt{N_{\text{seg}}}}{2\sqrt{2N_{\text{seg}}(1 + \rho^2)}} \right) \quad (1.56)$$

<sup>6</sup>In a realistic CW search, the number of templates across other dimensions of the parameter space, such as sky position, would depend on  $N_{\text{seg}}$ . The resulting number of templates, however, would be reduced with respect to a coherent search, as such a dependency would scale with the length of the segments rather than the length of the datastream.



where

$$\hat{z}_t = \frac{\hat{s}_t - 2N_{\text{seg}}}{2\sqrt{N_{\text{seg}}}} \quad (1.57)$$

is a renormalization of the threshold  $\hat{s}_t$  and  $\text{erfc}$  is the complementary error function. Given a threshold  $\hat{s}_t$  and a false dismissal probability  $p_{\text{fd}}$ , Eq. (1.56) provides an estimate for the minimum detectable SNR of a search

$$\rho = \frac{2}{\sqrt{MN_{\text{seg}}^{1/4}}} \sqrt{\hat{z}_t + \frac{2q^2}{\sqrt{N_{\text{seg}}}} \left(1 + \sqrt{1 + \frac{N_{\text{seg}} + \hat{z}_t\sqrt{N_{\text{seg}}}}{q^2}}\right)} \quad (1.58)$$

where for convenience  $q = \sqrt{2}\text{erfc}^{-1}(2p_{\text{fd}})$ . All the terms inside the square root are  $\mathcal{O}(N_{\text{seg}}^{-1/2})$  at most. This is the so called ‘‘fourth root law’’ of semicoherent searches, usually stated as  $\rho \sim N_{\text{seg}}^{-1/4}$ ; it is clear from this derivation that the origin of such a scaling is in Eq. (1.52) where the the number of segments enters linearly rather than quadratically, and thus is a general consequence of combining data in terms of power rather than complex amplitude. As pointed out in [70, 71], such a description is inaccurate unless  $N_{\text{seg}}$  is at least between 100 to 1000.

Equation (1.58) is in fact valid for all values of  $N_{\text{seg}}$  provided that  $p_{\text{fa}}$  is low enough [70]; we can thus study the relative performance of the several situations discussed in these past subsections in a ‘‘strong signal’’ regime, i.e. assuming signals are strong enough not to be confused with background noise. Such a comparison is shown in Fig. 1.1.

The left panel displays the performance of a coherent search versus a semicoherent search assuming the frequency of the signal is known (that is, deriving the corresponding threshold from a chi-squared distribution). The sensitivity of a semicoherent search scales at a slower rate with data duration than a coherent one. This can be justified by considering the relative Occam factor between the amplitude priors used by each of the searches. Concretely, the family of signals considered by a semicoherent search is broader than that of a coherent search, as amplitude consistency is only imposed within each segment independently. The Occam factor in a semicoherent search is, thus, higher, and correspondingly the required amplitude to distinguish a signal from the noise at a given false alarm is higher.

The right panel shows the equivalent result assuming the frequency of the signal is unknown. In this case, we use the loudest expected candidate as a threshold in order to account for the increased trials factor due to searching across a frequency band. For the coherent case, increasing the trials factor has a negative impact on the sensitivity of the search compared to the single-frequency case, as the expected loudest maximum value is higher than the single-template threshold at the specified false alarm probability. This is untrue for the semicoherent search. The reason was already discussed in Sec. 1.3.3: selecting the loudest candidate as a threshold comes at the cost of setting an unknown false alarm probability. As a result, search results cannot be compared insofar that they would be assuming different levels of significance. In this case, the resulting threshold on the search over the template bank is lower than what was used in the single-template search (i.e. it corresponds to a higher false-alarm probability), yielding overall a lower detectable amplitude.

The interplay between raising a threshold due to a high number of templates and diminishing the sensitivity of a search due to an increase in the number of segments may produce a variety of different situations in this type of searches. The specific example we show suggests coherent searches are the most sensitive kind of search regardless of whether we evaluate one or many templates; in a general CW search, however, where the number of templates is several orders of magnitude higher, it could well be the case that beyond a certain parameter-space region a semicoherent search yields better results than a coherent one, as in such case the lower number of templates would tend to produce lower detection thresholds. Investigating such a statement is beyond the scope of this introduction.

## 1.4 Toward realistic continuous-wave searches

Throughout this chapter we have discussed the basic issues concerning a CW search. Sections 1.1 and 1.2 introduced the primary targets of these searches, namely rapidly-spinning NSs sustaining a quadrupolar deformation, and briefly discussed the information we could derive from a CW detection. In Sec. 1.3, we discussed the set up and sensitivity of a CW-like search for monochromatic signals in Gaussian noise; therein we derived some of the basic properties of CW searches, such as the sensitivity and computing

cost difference between coherent and semicoherent searches, or the derivation of a threshold based on the loudest background event. In this last section, we briefly discuss the remaining traits of a CW search not captured by the toy model in Sec. 1.3; these are related to both the CW signal model upon arrival at the detector and the response of the detector itself to such a signal.

Approximating a CW signal by a monochromatic model is only valid for short time-scales compared to the typical evolution time-scale of the source and the detector’s position in the Solar system. This is due to two main reasons. First, the detector is moving with respect to the Solar system barycenter (SSB), inducing a Doppler modulation on any CW signal emitted from a specific sky position. As a result, all-sky searches need to resolve the sky position of the source in order to apply the methods outlined in Sec. 1.3. Second, a CW may undergo intrinsic processes changing the frequency of the emitted signal; these include a diminishment in frequency due to the emission of energy, which can be modelled by the inclusion of spin-down terms [see Eq. (1.6)], or an extra Doppler modulation due to the presence of a binary companion, which requires the inclusion of orbital parameters in the signal model. These effects increase the number of parameters a CW search must consider to achieve a detection. As a result, the number of required templates to cover a specific parameter-space region grows with a large power of  $T$  [72, 73] [cf. the linear dependency on  $T$  for a single parameter in Eq. (1.39)], effectively bounding the sensitivity of an all-sky search by the available computing resources. The overall arguments to setup a search, however, remain unchanged: once a template bank is set, a detection statistic is evaluated in each template and candidates are selected following a suitable criterion.

Regarding interferometric detectors, we are concerned with two main issues. First, the detector’s noise is non-stationary and non-Gaussian due to instrumental and environmental causes [74, 75], which increases the difficulty of the detection and characterization of weak monochromatic signals. Several approaches exist to deal with specific kinds of noise. Mild non-stationarities can be handled by estimating the noise’s PSD and whitening the data [52]; more specific artifacts, such as narrow instrumental lines, require post-hoc vetoes [76] and extended noise hypotheses to be appropriately dealt with [77]. Second, by construction, the sensitivity of a detector is non-isotropic in the sky. Coupled with the movement of the detector around the SSB, this induces an amplitude modulation on a CW signal coming from a specific sky position. This modulation can be taken into account either by weighting the segment-wise statistics in a semicoherent search [78, 79, 80], or by including the amplitude modulations into the signal model, which leads to the definition of a more complicated detection statistic [58, 81].

In sum, all-sky CW searches consist of evaluating a detection statistic over a grid covering a specific parameter-space region. The detection statistics are similar to Fourier transforms, albeit including extra modulations due to the frequency and amplitude modulations affecting a CW signal. The required integration time to accumulate enough significance to claim a detection leads to a fine discretization of the parameter space, highly increasing the amount of templates to be analyzed in order not to miss a detection. The sensitivity of a CW search, thus, is limited by the available computing power. Chapter 2 will be devoted to describing several of the trade-offs employed to conduct searches in Advanced LIGO data.

# Bibliography

- [1] A. Einstein, “Approximative Integration of the Field Equations of Gravitation,” *Sitzungsber. Preuss. Akad. Wiss. Berlin (Math. Phys.)*, vol. 1916, pp. 688–696, 1916.
- [2] A. Einstein, “On Gravitational Waves,” *Sitzungsber. Preuss. Akad. Wiss. Berlin (Math. Phys.)*, vol. 1918, pp. 154–167, 1918.
- [3] B. P. Abbott *et al.*, “Observation of Gravitational Waves from a Binary Black Hole Merger,” *Phys. Rev. Lett.*, vol. 116, no. 6, p. 061102, 2016.
- [4] J. Aasi *et al.*, “Advanced LIGO,” *Class. Quant. Grav.*, vol. 32, p. 074001, 2015.
- [5] F. Acernese *et al.*, “Advanced Virgo: a second-generation interferometric gravitational wave detector,” *Class. Quant. Grav.*, vol. 32, no. 2, p. 024001, 2015.
- [6] K. L. Dooley *et al.*, “GEO 600 and the GEO-HF upgrade program: successes and challenges,” *Class. Quant. Grav.*, vol. 33, p. 075009, 2016.
- [7] T. Akutsu *et al.*, “KAGRA: 2.5 Generation Interferometric Gravitational Wave Detector,” *Nature Astron.*, vol. 3, no. 1, pp. 35–40, 2019.
- [8] M. Saleem *et al.*, “The science case for LIGO-India,” *Class. Quant. Grav.*, vol. 39, no. 2, p. 025004, 2022.
- [9] R. Abbott *et al.*, “GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo During the Second Part of the Third Observing Run.” arXiv:2111.03606 [gr-qc].
- [10] B. P. Abbott *et al.*, “GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral,” *Phys. Rev. Lett.*, vol. 119, no. 16, p. 161101, 2017.
- [11] R. Abbott *et al.*, “Observation of Gravitational Waves from Two Neutron Star–Black Hole Coalescences,” *Astrophys. J. Lett.*, vol. 915, no. 1, p. L5, 2021.
- [12] M. Maggiore *et al.*, “Science Case for the Einstein Telescope,” *JCAP*, vol. 03, p. 050, 2020.
- [13] D. Reitze *et al.*, “Cosmic Explorer: The U.S. Contribution to Gravitational-Wave Astronomy beyond LIGO,” *Bull. Am. Astron. Soc.*, vol. 51, no. 7, p. 035, 2019.
- [14] W. Baade and F. Zwicky, “Remarks on super-novae and cosmic rays,” *Phys. Rev.*, vol. 46, pp. 76–77, Jul 1934.
- [15] D. G. Yakovlev, P. Haensel, G. Baym, and C. Pethick, “Lev Landau and the concept of neutron stars,” *Physics Uspekhi*, vol. 56, pp. 289–295, Mar. 2013.
- [16] T. Gold, “Rotating neutron stars as the origin of the pulsating radio sources,” *Nature*, vol. 218, pp. 731–732, 1968.
- [17] I. S. Shklovsky, “On the Nature of the Source of X-Ray Emissions of Sco XR-1.,” *Astrophys. J. Lett.*, vol. 148, p. L1, Apr. 1967.
- [18] R. N. Manchester, G. B. Hobbs, A. Teoh, and M. Hobbs, “The Australia Telescope National Facility Pulsar Catalogue,” *Astrophys. J.*, vol. 129, pp. 1993–2006, Apr. 2005.
- [19] G. B. Hobbs, R. N. Manchester, and L. Toomey, “ATNF Pulsar Catalogue v1.70.” <https://www.atnf.csiro.au/research/pulsar/psrcat/>, 2005. Accessed: 2023-03-24.

- [20] D. R. Lorimer and M. Kramer, *Handbook of Pulsar Astronomy*. 2004. ISBN: 978-0521535342.
- [21] J. M. Lattimer, “The nuclear equation of state and neutron star masses,” *Annu. Rev. Nucl. Part. Sci.*, vol. 62, no. 1, pp. 485–515, 2012.
- [22] K. Glampedakis and L. Gualtieri, “Gravitational waves from single neutron stars: an advanced detector era survey,” *Astrophys. Space Sci. Libr.*, vol. 457, pp. 673–736, 2018.
- [23] S. Bonazzola and E. Gourgoulhon, “Gravitational waves from pulsars: Emission by the magnetic field induced distortion,” *Astron. Astrophys.*, vol. 312, p. 675, 1996.
- [24] B. Haskell, L. Samuelsson, K. Glampedakis, and N. Andersson, “Modelling magnetically deformed neutron stars,” *Mon. Not. Roy. Astron. Soc.*, vol. 385, pp. 531–542, 02 2008.
- [25] R. Ciolfi, V. Ferrari, and L. Gualtieri, “Structure and deformations of strongly magnetized neutron stars with twisted-torus configurations,” *Mon. Not. Roy. Astron. Soc.*, vol. 406, pp. 2540–2548, 08 2010.
- [26] G. Ushomirsky, C. Cutler, and L. Bildsten, “Deformations of accreting neutron star crusts and gravitational wave emission,” *Mon. Not. Roy. Astron. Soc.*, vol. 319, p. 902, 2000.
- [27] T. J. Hutchins and D. I. Jones, “Gravitational radiation from thermal mountains on accreting neutron stars: sources of temperature non-axisymmetry.” arXiv:2212.07452 [astro-ph.HE].
- [28] F. Gittins and N. Andersson, “Modelling neutron star mountains in relativity,” *Mon. Not. Roy. Astron. Soc.*, vol. 507, no. 1, pp. 116–128, 2021.
- [29] J. A. Morales and C. J. Horowitz, “Neutron star crust can support a large ellipticity,” *Mon. Not. Roy. Astron. Soc.*, vol. 517, no. 4, pp. 5610–5616, 2022.
- [30] B. J. Owen, “Maximum elastic deformations of compact stars with exotic equations of state,” *Phys. Rev. Lett.*, vol. 95, p. 211101, 2005.
- [31] R. Narayan, “The Galactic Pulsar Population and Neutron Star Birth,” in *The Origin and Evolution of Neutron Stars*, vol. 125, p. 67, Jan. 1987.
- [32] R. Narayan, “The Birthrate and Initial Spin Period of Single Radio Pulsars,” *Astrophys. J.*, vol. 319, p. 162, Aug. 1987.
- [33] K. Rajwade, J. Chennamangalam, D. Lorimer, and A. Karastergiou, “The Galactic halo pulsar population,” *Mon. Not. Roy. Astron. Soc.*, vol. 479, no. 3, pp. 3094–3100, 2018.
- [34] J. Lazio, “The Square Kilometre Array,” in *Panoramic Radio Astronomy: Wide-field 1-2 GHz Research on Galaxy Evolution*, p. 58, Jan. 2009.
- [35] C. Palomba, “Simulation of a population of isolated neutron stars evolving through the emission of gravitational waves,” *Mon. Not. Roy. Astron. Soc.*, vol. 359, pp. 1150–1164, 2005.
- [36] N. Andersson, “A gravitational-wave perspective on neutron-star seismology,” *Universe*, vol. 7, no. 4, 2021.
- [37] R. V. Wagoner, “Gravitational radiation from accreting neutron stars,” *Astrophys. J.*, vol. 278, pp. 345–348, Mar. 1984.
- [38] G. Ushomirsky, L. Bildsten, and C. Cutler, “Gravitational waves from low mass X-ray binaries: A Status report,” *AIP Conf. Proc.*, vol. 523, no. 1, pp. 65–74, 2000.
- [39] G. Ushomirsky, C. Cutler, and L. Bildsten, “Deformations of accreting neutron star crusts and gravitational wave emission,” *Mon. Not. Roy. Astron. Soc.*, vol. 319, pp. 902–932, 12 2000.
- [40] A. Watts, B. Krishnan, L. Bildsten, and B. F. Schutz, “Detecting gravitational wave emission from the known accreting neutron stars,” *Mon. Not. Roy. Astron. Soc.*, vol. 389, pp. 839–868, 2008.

- [41] B. Haskell, J. L. Zdunik, M. Fortin, M. Bejger, R. Wijnands, and A. Patruno, “Fundamental physics and the absence of sub-millisecond pulsars,” *Astron. Astrophys.*, vol. 620, p. A69, Dec. 2018.
- [42] M. Maggiore, *Gravitational Waves: Volume 1: Theory and Experiments*. Gravitational Waves, OUP Oxford, 2008. ISBN: 9780198570745.
- [43] A. Einstein, “Explanation of the Perihelion Motion of Mercury from the General Theory of Relativity,” *Sitzungsber. Preuss. Akad. Wiss. Berlin (Math. Phys.)*, vol. 1915, pp. 831–839, 1915.
- [44] J. M. Weisberg and J. H. Taylor, “Relativistic binary pulsar B1913+16: Thirty years of observations and analysis,” *ASP Conf. Ser.*, vol. 328, p. 25, 2005.
- [45] M. Zimmermann and E. Szedenits, “Gravitational waves from rotating and precessing rigid bodies: Simple models and applications to pulsars,” *Phys. Rev. D*, vol. 20, pp. 351–355, Jul 1979.
- [46] M. Sieniawska and M. Bejger, “Continuous gravitational waves from neutron stars: current status and prospects,” *Universe*, vol. 5, no. 11, p. 217, 2019.
- [47] D. I. Jones, “Learning from the Frequency Content of Continuous Gravitational Wave Signals.” arXiv:2111.08561 [astro-ph.HE], 11 2021.
- [48] P. D. Lasky, “Gravitational Waves from Neutron Stars: A Review,” *Publ. Astron. Soc. Austral.*, vol. 32, p. e034, 2015.
- [49] B. Krishnan, A. M. Sintes, M. A. Papa, B. F. Schutz, S. Frasca, and C. Palomba, “The Hough transform search for continuous gravitational waves,” *Phys. Rev. D*, vol. 70, p. 082001, 2004.
- [50] K. Wette *et al.*, “Searching for gravitational waves from Cassiopeia A with LIGO,” *Class. Quant. Grav.*, vol. 25, p. 235011, 2008.
- [51] G. L. Bretthorst, *Bayesian Spectrum Analysis and Parameter Estimation*. Springer New York, 1988. DOI: 10.1007/978-1-4684-9399-3.
- [52] R. Prix, “The F-statistic and its implementation in ComputeFStatistic\_v2.” <https://dcc.ligo.org/LIGO-T0900149/public>.
- [53] R. J. Dupuis and G. Woan, “Bayesian estimation of pulsar parameters from gravitational wave data,” *Phys. Rev. D*, vol. 72, p. 102002, 2005.
- [54] C. Rover, R. Meyer, and N. Christensen, “Modelling coloured residual noise in gravitational-wave signal processing,” *Class. Quant. Grav.*, vol. 28, p. 015010, 2011.
- [55] C. Rover, “A Student-t based filter for robust signal detection,” *Phys. Rev. D*, vol. 84, p. 122004, 2011.
- [56] L. S. Finn, “Detection, measurement and gravitational radiation,” *Phys. Rev. D*, vol. 46, pp. 5236–5249, 1992.
- [57] G. Ashton and R. Prix, “Hierarchical multistage MCMC follow-up of continuous gravitational wave candidates,” *Phys. Rev. D*, vol. 97, no. 10, p. 103020, 2018.
- [58] P. Jaranowski, A. Krolak, and B. F. Schutz, “Data analysis of gravitational - wave signals from spinning neutron stars. 1. The Signal and its detection,” *Phys. Rev. D*, vol. 58, p. 063001, 1998.
- [59] R. Prix and B. Krishnan, “Targeted search for continuous gravitational waves: Bayesian versus maximum-likelihood statistics,” *Class. Quant. Grav.*, vol. 26, p. 204013, 2009.
- [60] D. Keitel, R. Prix, M. A. Papa, P. Leaci, and M. Siddiqi, “Search for continuous gravitational waves: Improving robustness versus instrumental artifacts,” *Phys. Rev. D*, vol. 89, no. 6, p. 064023, 2014.
- [61] E. T. Jaynes, *Probability Theory: The Logic of Science*. Cambridge University Press, 2003. DOI: 10.1017/CBO9780511790423.

- [62] J. Neyman and E. S. Pearson, “On the Problem of the Most Efficient Tests of Statistical Hypotheses,” *Phil. Trans. Roy. Soc. Lond. A*, vol. 231, no. 694-706, pp. 289–337, 1933.
- [63] A. C. Searle, “Monte-Carlo and Bayesian techniques in gravitational wave burst data analysis,” in *12th Gravitational Wave Data Analysis Workshop*, 4 2008.
- [64] R. Tenorio, D. Keitel, and A. M. Sintes, “Application of a hierarchical MCMC follow-up to Advanced LIGO continuous gravitational-wave candidates,” *Phys. Rev. D*, vol. 104, no. 8, p. 084012, 2021.
- [65] R. Tenorio, L. M. Modafferi, D. Keitel, and A. M. Sintes, “Empirically estimating the distribution of the loudest candidate from a gravitational-wave search,” *Phys. Rev. D*, vol. 105, no. 4, p. 044029, 2022.
- [66] R. Prix, “Search for continuous gravitational waves: Metric of the multi-detector F-statistic,” *Phys. Rev. D*, vol. 75, p. 023004, 2007. [Erratum: *Phys.Rev.D* 75, 069901 (2007)].
- [67] R. Prix, “Template-based searches for gravitational waves: Efficient lattice covering of flat parameter spaces,” *Class. Quant. Grav.*, vol. 24, pp. S481–S490, 2007.
- [68] B. Allen, “Spherical ansatz for parameter-space metrics,” *Phys. Rev. D*, vol. 100, no. 12, p. 124004, 2019.
- [69] P. R. Brady, T. Creighton, C. Cutler, and B. F. Schutz, “Searching for periodic sources with LIGO,” *Phys. Rev. D*, vol. 57, pp. 2101–2116, 1998.
- [70] K. Wette, “Estimating the sensitivity of wide-parameter-space searches for gravitational-wave pulsars,” *Phys. Rev. D*, vol. 85, p. 042003, 2012.
- [71] C. Dreissigacker, R. Prix, and K. Wette, “Fast and Accurate Sensitivity Estimation for Continuous-Gravitational-Wave Searches,” *Phys. Rev. D*, vol. 98, no. 8, p. 084058, 2018.
- [72] K. Wette, “Lattice template placement for coherent all-sky searches for gravitational-wave pulsars,” *Phys. Rev. D*, vol. 90, no. 12, p. 122010, 2014.
- [73] P. Leaci and R. Prix, “Directed searches for continuous gravitational waves from binary systems: parameter-space metrics and optimal Scorpius X-1 sensitivity,” *Phys. Rev. D*, vol. 91, no. 10, p. 102003, 2015.
- [74] P. B. Covas *et al.*, “Identification and mitigation of narrow spectral artifacts that degrade searches for persistent gravitational waves in the first two observing runs of Advanced LIGO,” *Phys. Rev. D*, vol. 97, p. 082002, Apr. 2018.
- [75] D. Davis *et al.*, “LIGO detector characterization in the second and third observing runs,” *Class. Quant. Grav.*, vol. 38, no. 13, p. 135014, 2021.
- [76] P. Leaci, “Methods to filter out spurious disturbances in continuous-wave searches from gravitational-wave detectors,” *Phys. Scripta*, vol. 90, no. 12, p. 125001, 2015.
- [77] D. Keitel, “Robust semicoherent searches for continuous gravitational waves with noise and signal models including hours to days long transients,” *Phys. Rev. D*, vol. 93, no. 8, p. 084024, 2016.
- [78] C. Palomba, P. Astone, and S. Frasca, “Adaptive hough transform for the search of periodic sources,” *Class. Quant. Grav.*, vol. 22, p. S1255, sep 2005.
- [79] B. Krishnan and A. M. Sintes, “Hough search with improved sensitivity.” <https://dcc.ligo.org/LIGO-T070124/public>, 2007.
- [80] P. B. Covas, “Searching for continuous gravitational waves with Advanced LIGO.” PhD Thesis, U. of the Balearic Islands, 2020.
- [81] C. Cutler and B. F. Schutz, “The Generalized F-statistic: Multiple detectors and multiple GW pulsars,” *Phys. Rev. D*, vol. 72, p. 063006, 2005.

## Chapter 2

# Search methods for continuous gravitational-waves from unknown sources in the advanced-detector era

This chapter is an adaptation of the material presented in

*Search methods for continuous gravitational-wave signals from unknown sources in the advanced-detector era*

Rodrigo Tenorio, David Keitel, Alicia M. Sintes

*Universe* 2021, 7(12), 474 – [arXiv:2111.12575 \[gr-qc\]](https://arxiv.org/abs/2111.12575)

DOI: 10.3390/universe7120474

## 2.1 Introduction

The search for continuous gravitational-wave signals (CWs), long-duration forms of gravitational radiation, is one of the endeavours of gravitational-wave astronomy. These signals are produced by long-standing quadrupolar variations, such as rapidly-spinning neutron star (NSs) sustaining a crustal deformation, undergoing an r-mode instability or in free precession [1, 2, 3, 4], as well as more exotic sources such as the annihilation of ultra-light boson clouds around spinning black holes [5, 6, 7, 8] or compact dark matter objects (CDOs) in the Solar System [9].

Detecting a CW signal could shed some light on NS physics, as well as open a new channel to test general relativity (extra polarizations, Lorentz violations) or detect dark matter [10, 11, 12, 13, 14]. No confident CW detection has been reported to date. The current product of CW searches are source-agnostic upper limits on the nominal CW amplitude  $h_0$ . These results can then be mapped into different astrophysical scenarios, such as the ellipticity of nearby NSs, the mass of ultralight bosons around black holes [15, 16], or the nearby population of planetary-mass primordial black hole binaries [17, 18].

The present document reviews search methods and pipelines employed to look for CW signals in the observing runs performed by the second generation of ground-based interferometric detectors (advanced detectors). Reviews on the physical mechanisms of CW emission by NSs can be found in [1, 2, 3, 4]. Basic data analysis techniques are discussed in [19]. Finally, [20] discusses the main results of previous CW searches up to 2017.

The standard CW signal model consists of a quasi-monochromatic source emitting gravitational waves at a certain frequency  $f_0$ . For the case of a NS sustaining a certain ellipticity,  $f_0$  corresponds to twice the spinning frequency of the star. Further time derivatives of the frequency arise due to different physical mechanisms affecting the source, such as energy emission as gravitational or electromagnetic radiation. In the case of sources in binary systems, the orbital motion induces a Doppler modulation.

Regardless of the specific intrinsic frequency modulation of the source, CW signals as seen from Earth are Doppler-modulated due to the detector motion around the Solar System barycenter (SSB). For a source with a given intrinsic frequency evolution  $\hat{f}$ , the detector-frame frequency is given by

$$f(t; \lambda) = \hat{f}(t) \left[ 1 + \frac{\vec{v}(t)}{c} \cdot \vec{n} \right], \quad (2.1)$$




FIGURE 2.1: Data spectrogram showing the response of an Earth-bound detector to a passing CW signal emitted by an isolated source. The left panel displays the one-year evolution of the detector response, while the right panel zooms in on a month-long period. Frequency modulations correspond to the yearly translation of Earth around the SSB (the daily Doppler modulation is contained within a frequency bin). Amplitude modulations correspond to the change in antenna-pattern functions throughout a day.

where the phase-evolution parameters  $\lambda$  include the CW frequency  $f_0$  and the sky position of the source  $\vec{n}$ , as well as any other parameter describing the intrinsic frequency evolution of the signal.  $\vec{v}(t)/c$  refers to the detector velocity expressed as a fraction of the speed of light, and contains both the daily and yearly motion of Earth around the SSB, of orders  $\mathcal{O}(10^{-6})$  and  $\mathcal{O}(10^{-4})$ , respectively. The yearly modulations can be clearly seen in the left panel of Fig. 2.1; the daily modulation is contained within a frequency bin.

The amplitude of a CW signal is described using four parameters, namely the initial CW phase  $\phi_0$ , the spherical angles describing the orientation of the source  $\{\psi, \cos \iota\}$ , and the nominal gravitational wave amplitude  $h_0$ . The response of a ground-based detector to a passing CW is better described in terms of the so called JKS representation [21]

$$h(t; \lambda, \mathcal{A}) = \sum_{\mu=0}^3 \mathcal{A}^\mu h_\mu(t; \lambda), \quad (2.2)$$

where the four time-independent  $\mathcal{A}^\mu$  depend on the four amplitude parameters  $\mathcal{A} = \{\phi_0, \psi, \cos \iota, h_0\}$  and the antenna-pattern response of the detector is contained in the four quadratures  $h_\mu(t; \lambda)$ , which only depend on the phase-evolution parameters. The basic effect of the antenna-pattern response on a CW signal is a daily amplitude modulation, clearly visible in Fig. 2.1.

As opposed to the short signals produced by compact binary coalescences (CBCs), with typical durations between a few minutes and less than a second for current detectors, CW signals are expected to last for years, spanning several observing runs of the current and future generation of ground-based interferometric detectors [22, 23, 24, 25, 26]. This difference in duration is crucial in terms of detecting and estimating the parameters of a CW signal.

Modelled searches are usually performed using *matched filtering* [27], comparing the datastream to a set of templates in order to find a high correlation. Due to the typical duration of a CW signal (spanning the entire observing run), the required number of templates to perform a blind search is prohibitively high even for current computing standards [28, 29, 30].

The standard strategy, in a broad sense, is to reduce the effective length of the datastream by performing matched filtering over shorter segments; the segment-wise results can then be combined into a final statistic. These kind of schemes are usually referred to as *semicoherent* searches [19]: the segment-wise analysis is typically referred to as *coherent*, as it compares the phase evolution of a signal with the datastream throughout a coherence time  $T_{\text{coh}}$ . The resulting coherent filters are then combined incoherently (i. e. maintaining a consistent frequency evolution but allowing for deviations in the amplitude parameters) into a final detection statistic. This incoherent combination allows to recover part of the sensitivity lost due to the split of the initial datastream; the final sensitivity, however, is lower than that of a fully coherent search. Figure 2.2 illustrates the principle of operation of semicoherent






FIGURE 2.2: Effect of different values of  $T_{\text{coh}}$  on one of the standard detection statistics for CW searches, the  $\mathcal{F}$ -statistic (see Sec. 2.2.1). Longer coherence times impose a greater penalization to deviations with respect to the signal model; as a result, peaks tend to become narrower around the true signal parameters.

searches. By comparing shorter streams of data, a looser constraint is imposed when comparing phase-evolution templates. As a result, the characteristic width of detection-statistic peaks widens, reducing the required number of templates to ensure a good covering of the parameter space. This strategy is at the core of multi-stage approaches such as those discussed in Sec. 2.4.

Parameter estimation, on the other hand, is positively affected by the long signal durations. Typical frequency resolutions are under a mHz for initial stages, achieving nHz resolution at fully-coherent follow-ups using a year of data. Sky localization is also significantly improved. If we think of CBC sky localization, neglecting contributions from antenna-pattern amplitude modulations and higher GW modes, the problem is basically that of determining the direction a GW pulse came from, which can be solved by means of measuring the pulse from  $N$  detectors and finding the overlapping sky-positions [31]. The case of CW signals is much simpler, as they are not just single pulses, but continuously arrive at the detector as it moves around the SSB. Due to this movement, a *single* interferometric detector receiving a CW signal at different positions with respect to the SSB is essentially equivalent to an arbitrary large set of *different* detectors receiving the *same* pulse from a source. Hence, for CW sources, sharp sky localization can be achieved *using a single detector* by simply extending the duration of an observing run.

CW searches require a very fine parameter-space resolution in order not to miss a signal, increasing the computing cost of a matched-filtered search up to unaffordable figures [28]. As a result, searches for CW signals from unknown sources tend to follow a hierarchical approach [32, 33, 34, 35, 36, 37, 38]: wide parameter-space regions are analyzed using a less-constraining statistic so that a coarser template bank can be used, see Fig. 2.2. Interesting regions are then typically small enough to follow up using a more sensitive statistic. A quantitative description of this strategy can be found in [39].

The structure of this work goes as follows: in Sec. 2.2 we review the main search methods and pipelines employed to search for CW signals from unknown sources during the era of the advanced detectors. Section 2.3 discusses different post-processing stages employed by said searches; these include both specific prescriptions to select interesting candidates (e.g. clusterings) and consistency arguments to assess the overlap of a specific set of templates with instrumental disturbances. Section 2.4 reviews follow-up strategies to further analyze interesting candidates. So far, no search has claimed a confident CW detection, reporting instead constraints on the maximum detectable amplitude achieved. Different approaches to construct said constraints are listed in Sec. 2.5. Finally, a summary of the reviewed methods is presented in Sec. 2.6.

## 2.2 Wide parameter-space search pipelines

We present a review of methods and pipelines employed to date to search for CW signals from unknown sources during the era of the advanced detectors. Specifically, we consider three kinds of searches for unknown sources: (i) blind searches, with weak prior assumptions about possible source parameters [40, 41, 42, 43, 44, 45, 46, 47, 48, 49, 50, 51, 52, 53, 54]; (ii) spot-light searches, focused at sky regions

harbouring an interesting population of objects whose exact frequency is unknown, such as globular clusters or the Galactic Center (GC) [55, 56]; and (iii) directed searches, targeting specific celestial objects compatible with a CW source, such as supernova remnants (SNRs) or low-mass X-ray binaries (LMXBs) [57, 58, 59, 60, 61, 62, 63, 64, 65, 66, 67, 68, 69].

### 2.2.1 $\mathcal{F}$ -statistic searches

The  $\mathcal{F}$ -statistic is a standard detection statistic for CW signals. Initially derived as a maximum-likelihood estimator with respect to amplitude parameters [21, 70], it was later rederived in a Bayesian context as a Bayes factor, gauging the presence (or lack) of a signal in a Gaussian noise data stream, in which amplitude parameters are marginalized using a rather *unphysical* set of amplitude priors [71, 72, 73, 74, 75, 76]. This detection statistic can be extended for more generic types of sources, such as binary white-dwarf systems [77] or the inspiral phase of binary black-hole coalescences [78].

Semicoherent searches balance sensitivity and computing cost by choosing a suitable number of coherent segments whose combination into a semicoherent quantity can be performed in an efficient manner. The basic “stack-slide” procedure used for many  $\mathcal{F}$ -statistic semicoherent searches [79, 32, 33, 34] is to set up a template bank in each coherent segment to compute the segment-wise coherent detection statistic. Then, in the semicoherent stage, the semicoherent detection statistic is computed on a *finer* template bank by combining results from the segment-wise coherent detection statistics. This template bank refinement can be seen as a consequence of using multiple coherent segments: the higher the number of coherent templates to be combined, the higher the resulting number of distinct semicoherent templates.

We discuss three different implementations of semicoherent  $\mathcal{F}$ -statistic searches, namely the Global Correlation Transform hierarchical search (GCT), Weave and Time-domain  $\mathcal{F}$ -statistic. They mainly diverge in the manner of constructing semicoherent quantities, taking different trade-offs in terms of computing cost, memory requirement and robustness to non-Gaussianities. We note, however, that fully-coherent searches for targets at a specific sky position, such as supernova remnants [59, 66], are still performed nowadays.

#### GCT hierarchical search

Detection statistics (and the  $\mathcal{F}$ -statistic in particular) present a set of characteristic correlations across the CW parameter-space due to some level of degeneracy present in CW signals [80]. Understanding the structure of said correlations offers a simple method to reuse coherent  $\mathcal{F}$ -statistic values to construct a semicoherent statistic, reducing the overall computing cost of the search. The GCT introduces a set of coordinates defined by the intersection of parameter-space correlation surfaces to identify nearby parameter-space points where the  $\mathcal{F}$ -statistic achieves a high value due to the presence of a signal [81, 82, 83].

Due to its interesting trade-off between sensitivity and computational cost [84], semicoherent GCT searches have been used by the Einstein@Home project to perform deep searches throughout different observing runs [40, 49, 61, 62, 69]. Einstein@Home searches distribute the computational load of a search across a volunteer-computing network using BOINC [85] to analyze a high number of parameter-space candidates.

At the core of the GCT search there are combinations of  $\mathcal{F}$ -statistic values computed at different coherent segments containing data from one or multiple detectors. These are combined into the semicoherent  $\mathcal{F}$ -statistic using a finer template bank with refinement only on the spindown parameters [82]. This allows for the easy implementation of Bayesian extensions over the semicoherent  $\mathcal{F}$ -statistic, such as the “line-robust”  $B_{S/GL}$  or  $B_{S/GLtL}$  statistics [86, 87, 88, 89], which combine  $\mathcal{F}$ -statistics from different detectors to suppress single-detector artifacts.

No general computing model is available for the GCT search, requiring extensive software-injection campaigns in order to numerically tune the sensitivity of a search to the available computing resources [36]. As discussed in Sec. IV B of [90], this is due to a core assumption in [32, 82, 83] neglecting refinements in the sky parameter-space. Further developments in the field, discussed in Sec. 2.2.1, proposed a new strategy to solve this problem. For the case of directed searches, however, optimal setup strategies are available [91].

### Weave

Local parameter-space structure can be understood in terms of the *mismatch*  $\mu$  [92, 93]. Given a signal parameterized by a set of parameters  $\lambda_0$ , the parameter-space mismatch  $\mu$  quantifies the fractional loss of (squared) signal-to-noise ratio  $\rho^2$  produced by an arbitrary parameter-space off-set  $\Delta\lambda$  with respect to the true signal parameters:

$$\mu(\Delta\lambda; \lambda_0) = 1 - \frac{\rho^2(\lambda_0 + \Delta\lambda)}{\rho^2(\lambda_0)}. \quad (2.3)$$

In a neighbourhood of  $\lambda_0$ , the mismatch can be expressed in terms of a quadratic form as

$$\mu(\Delta\lambda; \lambda_0) \simeq \Delta\lambda \cdot \mathbf{g} \cdot \Delta\lambda, \quad (2.4)$$

where the symmetric 2-rank tensor  $\mathbf{g}$  plays the role of a Riemannian metric in the parameter-space.

It was quickly realized that, for the case of flat parameter spaces, the metric  $\mathbf{g}$  offers a complete description of the parameter-space, allowing for the easy construction of optimal setups [29, 94]; alas, the standard CW parameter-space (specifically, the sky-position subspace) presents a non-trivial structure resulting in a curved parameter-space.

**Weave** [90] represents the coming together of a number of search strategies. It implements, for the first time, a semicoherent search with a well-understood computational model based on a flat parameter-space metric. The setup is capable of constructing template banks at the suitable resolution to achieve an optimal computing cost [95, 30, 90, 96]. Setting up **Weave** requires only a list of time stamps delimiting semicoherent segments, a coherent mismatch  $\tilde{\mu}$ , with which single-segment template banks will be constructed, and a semicoherent mismatch  $\hat{\mu}$ , to setup the semicoherent template bank. The identification of a template from the semicoherent bank to its corresponding ensemble of coherent templates is handled by the **Weave** code using the coherent parameter-space metric to identify the nearest neighbor in each segment. In a sense, **Weave** retains the general characteristics of the GCT search, mainly being an engine to combine coherent  $\mathcal{F}$ -statistics, but, as opposed to it, **Weave** requires no extensive numerical calibration to be deployed using an optimal setup.

The success of **Weave** as an all-sky search, however, is related to the characteristics of the  $\mathcal{F}$ -statistic's parameter-space structure. As discussed in depth in [96, 97], the optimal setup for a realistic computing budget tends to yield semicoherent mismatch values well beyond the validity of the metric approximation (i. e.  $\mu \gg 1$ ). Nevertheless, empirical studies [96] show that in this regime the  $\mathcal{F}$ -statistic actually falls off more slowly than predicted by the metric, meaning the resulting template bank will contain more templates than strictly required. An alternative approach circumventing the empirical characterization of the  $\mu \gg 1$  regime is discussed in [98, 94, 99].

Despite achieving a better sensitivity than the GCT search at a fixed computing cost, the increased memory requirements of **Weave** make it so far unsuitable for its deployment on **Einstein@Home** [84]. Nonetheless, its sensitivity and setup flexibility have already been proved in the implementation of a novel CW search strategy [51].

### Time-domain $\mathcal{F}$ -statistic

Both GCT and **Weave** searches rely on a common implementation of the  $\mathcal{F}$ -statistic [100], publicly available under **LALSuite** [101]. The **Time-domain  $\mathcal{F}$ -statistic** pipeline [102] uses a different implementation, based on the  $\mathcal{F}$ -statistic's time-dependent behavior throughout the observing run. Instead of computing semicoherent quantities, it focuses on significant parameter-space points at coherent-segment level, looking for coincident candidates across different time segments and detectors. Using this coincidence criterion (which we review in more depth in Sec. 2.3) the pipeline automatically becomes robust to strong instrumental features, since they tend to overlap with different parameter-space regions as an observing run progresses.

This particular implementation of the  $\mathcal{F}$ -statistic uses its own template bank setup in order to optimize the number of fast Fourier transform (FFT) computations [103, 104], which normally takes a significant part of the overall computing cost of the search. Further improvements at parameter-estimation level are optimization algorithms to resolve the characteristic frequency multi-modality of the  $\mathcal{F}$ -statistic [105] and the inclusion of machine-learning algorithms to filter out non-astrophysical candidates [106].

### 2.2.2 Fourier-transform-based searches

The efficiency of  $\mathcal{F}$ -statistic searches stems from the marginalization with respect to amplitude parameters, removing four parameter-space dimensions from the search [29]. Semicoherent pipelines, moreover, assume amplitude parameters to be independent across different coherent segments. This approach makes it difficult to search for amplitude parameters, such as specific CW polarizations.

An alternative family of methods use Fourier transforms of short data segments (Short Fourier Transforms, SFTs) as the basic unit of operation. The duration of an SFT is typically such that CW signals are contained within a single frequency bin. For an isolated NS, this length is about 30 minutes [107], although the exact value depends on the considered frequency range. We note that SFTs themselves can, in some situations, correspond to coherent segments of a semicoherent search; in such cases, the effective coherent length is proportional to the SFT length.

In order to construct a detection statistic  $\mathcal{S}$ , Fourier transform amplitudes  $\tilde{x}$  are combined using a set of weights (effectively a kernel)  $\mathcal{K}$  taking into account source polarization, antenna-pattern amplitude modulations, Doppler modulations and relative phase deviations across different detectors:

$$\mathcal{S} = \sum_{t,t'} \tilde{x}^*(t)\tilde{x}(t')\mathcal{K}(t,t'), \quad (2.5)$$

where parameters other than time dependency have been kept implicit for the sake of simplicity. Frequency modulations are assumed to be contained within the specific set of Fourier amplitudes being combined, although  $\mathcal{K}$  can be configured to increase robustness against different kinds of spectral leakage by including neighboring frequency bins into the kernel [108, 39].

We review two families of searches stemming from Eq. (2.5), depending on whether their primary target is to increase the robustness against deviations from the intended CW model (**PowerFlux** & **Falcon**) or to improve sensitivity by increasing the effective amount of data used (**Cross-Correlation**).

#### PowerFlux & Falcon

The **PowerFlux** pipeline [109, 110, 111, 112] estimates power from a CW source depending on its sky position. The basic implementation [41, 42, 54] uses a diagonal kernel  $\mathcal{K}$  to combine Fourier power from each SFT. In this case, the role of  $\mathcal{K}$  is to diminish the contribution of unfavourable frequency bins (due to high noise floors or a low antenna-pattern response) to the total weighted Fourier power. Since the final statistic ignores relative phase shifts between SFTs, this corresponds to a semicoherent search for a specific CW polarization.

Loosely coherent methods [39, 113, 114] exploit the flexibility of Eq. (2.5) to set up a kernel  $\mathcal{K}$  to account for unmodeled phase shifts, be it due to parameter-space mismatches or unaccounted physics such as small binary orbital modulations [115]. To do so, phases are allowed to drift at most by a specific amount  $\delta$  across contiguous data segments, obtaining as a result [39]

$$\mathcal{K}(t,t';\delta) = \left(\frac{\sin \delta}{\delta}\right)^{|t-t'|}, \quad (2.6)$$

where data is assumed to come from a single detector for the sake of simplicity. (The framework presented in [39] is flexible enough to treat data from multiple detectors.) This simple kernel illustrates the principle of operation of loosely coherent searches: if  $\delta = \pi$ ,  $\mathcal{K}$  behaves like a delta function and phases among contiguous segments are uncorrelated, effectively performing a semicoherent search; if  $\delta = 0$ ,  $\mathcal{K} = 1$  and phases are correlated throughout the full data stream, performing a fully-coherent search. Tuning  $\delta$  to intermediate value allows to trade sensitivity and computing cost: low  $\delta$  values involve simpler kernels (less non-zero terms), easing the computing cost of Eq. (2.5) but imposing tighter constraints with respect to the chosen CW model. A discussion on physically relevant  $\delta$  values can be found in [39], whereas an efficient implementation of low- $\delta$  kernels can be found in [113, 114].

Due to its high computing cost, the initial implementation of loosely coherent methods on **PowerFlux** was mainly used on directed searches, such as spot-light surveys [55] or follow-up stages (see Sec. 2.4). An efficient implementation, **Falcon**, was later developed to perform all-sky searches throughout O1 [43, 44] and O2 [46, 47, 48], setting competitive constraints on the nearby population of Galactic neutron stars.

### Cross-Correlation

The **Cross-Correlation** search [116, 117, 118], implemented in the LALSuite library [101], closes the sensitivity gap between semicoherent and coherent searches by increasing the effective amount of data using the correlation of Fourier amplitudes at different times.

Although initially designed to look for stochastic GW backgrounds, CW signals are a perfect target to be looked for using **Cross-Correlation**, as they are long-lasting and deterministic. This allows to use not only the cross-correlation of different instruments during a certain period of time, but also the cross-correlation of data segments at different periods of time. In the context of CW searches, **Cross-Correlation** searches have been employed to look for specific high-priority targets, such as CW emission from the LMXB system Scorpius X-1 [58, 67]. Complementarily, searches for stochastic GW backgrounds, such as the **Radiometer** search [119, 120, 121], are able also to constrain CW amplitudes from specific sky locations, such as those of SNRs, LMXBs such as Scorpius X-1 or the Galactic Center [122, 123, 124].

To do so, the kernel  $\mathcal{K}$  is constructed following the expected correlation of a signal across different stretches of data at different times and detectors. Specifically, only amplitudes within a certain time range  $T_{\max}$  are combined together. As opposed to **PowerFlux** and **Falcon**, however, polarization angles are averaged out using uniform priors. The  $T_{\max}$  parameter plays a similar role to that of  $\delta$  (see Sec. 2.2.2) in terms of trading computing cost and sensitivity: longer  $T_{\max}$  increases the number of cross-correlations to perform, but also imposes a tighter constraint with respect to the specified CW model. Latest developments on this pipeline include the use of re-sampling techniques to accelerate its evaluation [125] and the use of a shear parameter-space coordinate transform to optimize the setup of template banks [118].

### 2.2.3 Hough-transform semicoherent searches

Loud instrumental features in the data tend to saturate detection statistics due to their strong resemblance to CW signals along short periods of time. This problem pushed forward the development of methods capable of suppressing narrow-band features in the data, such as the application of the Hough transform to the search for CW signals.

The basic idea is to limit the contribution of each semicoherent segment to a bounded quantity, which contributes by a limited amount to the overall statistic. This is done by binarizing a power-like quantity (whitened Fourier power formatted as SFTs, as discussed in Sec. 2.2.2, or the  $\mathcal{F}$ -statistic) into ones and zeroes using a predetermined threshold [107]. Such a binarization, however, comes at the cost of ignoring noise-floor variations and antenna-pattern amplitude modulations, meaning highly-contaminated frequency bands around low-sensitivity sky positions will contribute the same amount as clean frequency bands at the most favored sky positions. This problem is usually solved introducing a set of weights into the detection statistic [126, 127].

The Hough transform [128] can be used to identify shapes in binarized images. Given a parameter-space parametrizing a family of curves, each set of parameters is assigned a score, the number count, proportional to the number of pixels in the image consistent with the corresponding curve. In CW searches, the parameter space is typically the set of phase-evolution parameters, according to which the data spectrogram is traversed adding (weighted) ones and zeroes depending on whether each frequency bin contains excess power or not.

Hough-transform based searches have been widely used during the latest all-sky searches [41, 42, 45, 50, 52, 53]. We discuss the two main implementations of the Hough transform for CW searches, **SkyHough** and **FrequencyHough**.

#### SkyHough

The **SkyHough** pipeline was the initial implementation of the Hough transform to the search for CW signals [107, 129]. Due to the Earth’s movement around the SSB, a CW signal arriving at the detector at a certain time  $t$  with a given frequency  $f_0$  has a very specific set of sky positions from which it could have originated. Said sky positions take the shape of thick annuli (“circles in the sky”) [107, 80, 130] which, for consistency, can be described themselves as one/zero regions in the sky patch. The binarization of the data spectrogram, thus, gets mapped into the selection and summation of sky patches containing different selected annuli.

The “circles in the sky” are weakly affected by local changes in  $f_0$ ; hence, once computed for a specific frequency, these structures can be saved into a look-up table (LUT) to analyze several frequency bins. The use of LUTs, on the other hand, affects the maximum sensitivity of `SkyHough` since only an approximated frequency-evolution track is being used. Further developments on the pipeline include the re-analysis of interesting candidates using the exact frequency-evolution track and using more sensitive statistics in order to improve parameter estimation [131, 132]. Also, the combination of LUTs can be easily accelerated using GPU parallelization [133].

The use of LUTs depends on a basic set of CW parameters, namely frequency and sky position, meaning it can be arbitrarily extended to look for different types of CW signals as long as the source’s intrinsic and extrinsic frequency evolutions are uncoupled from the Earth’s Doppler modulation. This includes transient-like CW signals, such as those produced by newborn NSs [134], or NSs in binary systems [133].

### FrequencyHough

`FrequencyHough` is another pipeline based on the Hough transform [135, 136]. As opposed to `SkyHough`, binarized spectrograms are mapped onto the frequency and spindown subspace  $(f_0, f_1)$ . In this case, sky position is *fixed* for all the analyzed templates. This presents two main advantages with respect to `SkyHough`. First, it allows to increase frequency resolution independently of other parameters, resulting in a better parameter recovery. Second, all candidates analyzed at once are related to the same sky position. As a result, sky regions where templates tend to overlap with instrumental features can be dealt with in a simpler manner [137, 138]. This different mapping, however, limits the generalization possibilities of `FrequencyHough` to linear relations between parameters.

This search is usually combined with different data formats: wide parameter-space searches are usually run using the so called Short Fourier Data Base (SFDB) [139], which includes time-domain cleaning of raw data to reduce the effect of transient noise. For searches at a specific parameter-space region at hand, such as supernova remnants, the Galactic center or a specific outlier from a wider search, Band Sampled Data (BSD) [140, 141] is used to efficiently apply heterodyning filters, reducing the computing cost of analyzing such a local parameter-space region.

Several generalizations of these methods have been developed in order to look for other kinds of source. Power-law frequency evolution was covered using a generalization of `FrequencyHough` [142], allowing to search for binary neutron star merger remnants [143]. A study on the suitability of `FrequencyHough` to probe planetary-mass primordial black holes was presented in [17]. A method to conduct all-sky searches for CWs from the evaporation of boson clouds around spinning black holes was developed in [144], using the BSD framework [140, 141]. A related method to place constraints on dark-photon dark matter interacting directly with GW interferometers was developed in [145] and applied in [146].

### 2.2.4 Viterbi searches

CW signals could contain stochastic contributions (spin wandering) whose behaviour would not be well represented by the standard model introduced in Eq. (2.1). This could affect objects in both isolated and binary systems [147, 148], and should be taken into account in order to describe the underlying physics.

A simple approach is to describe the frequency-evolution model itself as a stochastic process, namely a Markov Chain (MC), and infer the most likely instantaneous frequency of a signal from the data. This is usually developed under the framework of Hidden Markov Models (HMM) [149, 150, 151, 152, 153], but an equivalent description can be done using Bayesian probability [154]. As discussed in [12], the search for ultralight boson-cloud evaporation around spinning black holes could also benefit from this kind of approach. A specific example of such a search, directed toward Cygnus X-1 using Advanced LIGO O2 data, was presented in [16].

Given a set of measurements at discrete times  $x = \{x_j, j = 0, \dots, N - 1\}$ , we want to infer the instantaneous frequency of a signal in the data  $f = \{f_j, j = 0, \dots, N - 1\}$ . This problem can be readily expressed as an inference one on  $f$  by means of Bayes theorem [155]

$$P(f|x) \propto P(f) \times \prod_{j=0}^{N-1} P(x_j|f_j) \quad (2.7)$$

Method	$\ln P(x_j f_j)$	Searches
Viterbi 1.0 [149]	(Bessel-weighted) $\mathcal{F}$ -statistic	[57, 16, 63, 156]
Viterbi 2.0 [150]	$\mathcal{J}$ -statistic	[60, 64, 157]
Viterbi SNR [151]	$\mathcal{F}$ -statistic	[65, 68]
Dual-harmonic Viterbi [152]	$\mathcal{F}$ -statistic	[68]
Transient Viterbi [158]	Norm. Fourier power	[143, 158]
SOAP [154]	Line-aware statistic	—
Viterbi 3.0 [153]	$\mathcal{B}$ -statistic	—

TABLE 2.1: Summary of CW searches based on a MC implemented via the Viterbi algorithm. Transition probabilities (second column) and other details of each specific pipeline are discussed in the text.

where the sampling distribution of different measurements has been (conservatively) factored assuming logical independence. The prior probability distribution on the instantaneous frequency  $P(f)$  is usually specified in terms of the initial frequency  $P(f_0)$  and the transition probabilities of the MC  $P(f_j|f_{j-1})$

$$P(f) = P(f_0) \times \prod_{j=1}^{N-1} P(f_j|f_{j-1}). \quad (2.8)$$

$P(f_0)$  is generally taken as a uniform distribution over the searched frequency band; the choice of transition probabilities (thus prior probabilities)  $P(f_j|f_{j-1})$  and sampling distribution  $P(x_j|f_j)$  is dependent upon the method and source of interest, as summarized in Table 2.1. Regardless of the astrophysical scope, searches using a MC evolution model [Eq. (2.7)] obtain the most likely (maximum-posterior) frequency-evolution path  $f^*$  using the Viterbi algorithm [159]

$$f^* = \arg \max_f P(f|x). \quad (2.9)$$

Two basic choices exist for transition probabilities [152], depending on whether the dominant frequency-drift time-scale is given by the source’s spin wandering or secular spindown. The former allows transitions to any neighbouring frequency bin at each time step. The latter uses a so-called biased HMM: given a frequency at bin  $j$ , the following frequency bin must be equal or lower,  $f_{j+1} \leq f_j$ ; the specific number of bins is usually between two and three [65, 156].

Since no template bank is involved, Viterbi searches are able to benefit from a wide variety of detection statistics at a small tuning cost. Specifically, binary modulations can be easily folded in using the  $\mathcal{J}$ -statistic, which improves over the  $\mathcal{C}$ -statistic [160] by combining frequency side-bands using complex weights. In order to track phase information, Viterbi 3.0 uses an efficient implementation of the  $\mathcal{B}$ -statistic [71, 72] first proposed in [113].

Finally, the SOAP pipeline [154] uses Bayesian spectral analysis to avoid relying on the  $\mathcal{F}$ -statistic, looking for CW signals displaying a sinusoidal behaviour during short periods of time. To do so, the sampling distribution is taken proportional to the data’s Schuster periodogram [161], adding an extra hypothesis to increase the robustness against strong monochromatic features in the data.

## 2.2.5 Machine learning

The use of machine-learning (ML) techniques in the search for CW signals follows one of two trends: either to classify and summarize the results of a search’s main stage, as discussed in Sec. 2.3.2, or to substitute the search step itself, acting as a detection statistic. A recent review on ML applications to GW data analysis in general can be found in [162].

A first approach is to train a classifier directly over Fourier-transformed raw time-series data to distinguish the presence of a signal within background noise. The specific data format is dependant upon the signals being looked for: for CW signals, which last for long periods of time over narrow frequency bands. In [163, 164], a convolutional neural network (CNN) is trained on the real and imaginary parts of short-time Fourier transforms. Slight variations on this proposal were employed to look for postmerger signals. These signals, compared to CW signals, last a shorter period of time over a wider frequency

band [165]. In order to reflect the non-trivial time dependency of the signal, Ref. [166] took Fourier transforms using shorter time durations, finally using the data spectrogram as input for, again a CNN.

The second approach takes a less radical point of view. Instead of starting from raw data, machine learning algorithms are applied on the output of a search pipeline to construct a new detection statistic [167, 168]. This approach could be beneficial, as the output of search pipelines typically enhances signal features across the output parameter space. The specific format with which the output data is better represented, however, remains a point of discussion.

Further applications of ML to post-process the output of a search [106, 169, 170] will be covered in Sec. 2.3.2.

## 2.3 Post-processing strategies

The main stage of a wide parameter-space search usually returns the loudest templates in terms of a specific detection statistic. A good portion of these templates are correlated, either because their corresponding time-frequency evolution tracks sweep over similar data or because of the presence of non-Gaussianities in the data, producing broad parameter-space artifacts [88, 171, 172]. The idea behind post-processing stages is to reduce the number of candidates into an affordable quantity to be followed up.

Complementarily, veto strategies can be applied in order to reduce further the number of candidates resulting from a search. A veto is a simple method to assess the consistency of a CW candidate with respect to a signal hypothesis. The outcome is usually a boolean answer; as opposed to a post-processing or follow-up stage, the primary objective is to quickly reject an inconsistent CW candidate at a low computing cost.

The following subsections summarize common post-processing strategies employed in contemporary CW searches. After a brief overview of coincidence and clustering steps, we review four families of vetoes. Other techniques employed in previous searches can be found in [137].

### 2.3.1 Coincidences

In a network of gravitational-wave detectors with a comparable level of sensitivity, a CW signal is expected to produce significant candidates in the analysis of every detector’s data. Imposing a coincidence criterion, that is, focusing on common parameter-space regions highlighted in every single dataset, reduces the false-alarm probability of the search, as noise fluctuations are less likely to be coincident across different detectors than CW signals [102]. This approach obeys a robustness versus sensitivity trade-off, as combining data from different detectors into a single analysis increases the sensitivity of a search without increasing the required number of templates to be evaluated, keeping the computing budget under control [93].

This approach has been widely employed in `FrequencyHough` and `SkyHough` searches. In both cases, a parameter-space distance, based on an Euclidean ansatz, is used to identify closeby candidates in each detector’s results. Other searches, such as `Time-domain  $\mathcal{F}$ -statistic`, `PowerFlux` or `Falcon`, impose coincidence criteria based on the overlap of enhanced parameter-space regions across both datasets, rather than using a parameter-space distance.

In particular, as discussed in Sec. 2.2, `Time-domain  $\mathcal{F}$ -statistic` is the most prominent user of this strategy [102]. As opposed to other semicoherent searches, it does not combine coherent segments into a semicoherent statistic, but looks for coincident candidates across multiple segments (including different detectors), reducing the overall false-alarm probability of the search.

### 2.3.2 Parameter-space clustering

Another option is to group together nearby templates according to some notion of distance. This process, usually referred to as *clustering* in the data analysis literature, has been extensively employed by `PowerFlux`, `Einstein@Home`, `FrequencyHough`, and `SkyHough` using different implementations, both in terms of clustering strategy and parameter-space distance [41, 42, 40, 45, 49, 50, 52]. Clustering is typically implemented using an unsupervised approach [173, 130]: the parameter-space and the clustering algorithm itself act as prior information to construct meaningful groupings of the resulting template bank.



Supervised ML approaches [106, 169, 170], aimed at identifying specific parameter-space structures, have also been proposed.

Unsupervised approaches try to unveil structure from a data set. Prior information is encoded both in terms of the distance used to compare nearby candidates and the linkage criteria. The resulting clusters are sieved through selection criteria which can take into account parameters like the maximum significance in the cluster or its number of elements. We focus our exposition on the choice of parameter-space distance; a complete description of the clustering algorithms themselves is given in the corresponding references.

Initial implementations, used by `FrequencyHough` and `SkyHough` [41, 42, 50] assumed a Euclidean parameter-space distance on the CW signal parameters  $\lambda = \{f_0, f_1, \dots, \vec{n}, \dots\}$

$$d(\lambda, \lambda_*) = \sqrt{\sum_i \left( \frac{\lambda^{(i)} - \lambda_*^{(i)}}{\delta\lambda^{(i)}} \right)^2}, \quad (2.10)$$

pairing candidates within a certain distance threshold to form the final clusters. The exact parameters  $\lambda^{(i)}$  and parameter-space resolutions  $\delta\lambda^{(i)}$  are search-dependent.

A more informative approach, still based on a Euclidean ansatz, was proposed in [173]. In this case, clusters are classified according to topographic parameters by projecting detection statistics over planes, namely over the frequency-spindown plane and the ecliptic plane. Instead of using the basic CW parameters, distance was computed after projecting the sky position of the candidate onto the ecliptic plane; thus allowing a greater variance around the ecliptic, where sky localization tends to become more uncertain [107, 80].

These distance measures are effective for local analyses, but quickly become unreliable whenever more involved parameter-space structures such as correlated parameters with periodic boundaries come into play, as is the case for signals from sources in binary systems [174]. A parameter-space distance for a generic, quasi-monochromatic CW signal was proposed in [130] using the instantaneous detector-frame frequency associated to a CW template  $f(t; \lambda)$ . Concretely, the distance between two templates  $\lambda$  and  $\lambda_*$  can be defined as the average mismatch between their corresponding detector-frame frequency tracks throughout an observing run

$$d(\lambda, \lambda_*) \propto \frac{1}{T} \int^T dt |f(t; \lambda) - f(t; \lambda_*)|. \quad (2.11)$$

This prescription is consistent with the  $\mathcal{F}$ -statistic's parameter-space correlations and can be simplified into a discrete sum for a faster implementation [130, 52, 53].

Candidate post-processing, and clustering in specific, is also a suitable step in which machine learning strategies are able to deliver an improvement of sensitivity. As opposed to raw data, on which CW signals are typically a subdominant contribution, the structures produced by different features in the data on the parameter space of a search are suitable to be classified using a supervised approach, as long as a clear classification of the features at hand is available. To date, two approaches have been proposed.

The first one [106], framed within the `Time-domain  $\mathcal{F}$ -statistic` pipeline, uses a convolutional network to classify different representations of the main search's output into three possible classes, namely noisy bands containing Gaussian-like noise, narrow spectral artifacts, and CW signals. A similar (albeit more manual) approach was reported in [40].

The second approach is developed as an alternative to `Einstein@Home` searches' clustering algorithm [169, 170]. In this case, a neural network is trained to recognize signal-induced patterns on a certain projection of the parameter space in order to identify typical structures associated to CW signals. The training set is produced by manually identifying software-injected signals using an image editing tool.

### 2.3.3 Detector-consistency vetoes

The first family of vetoes tests the consistency of a CW candidate across the network of GW detectors. The basic implementation is formulated as follows: given a CW candidate with parameters  $p$  and  $N$  detectors, a detection statistic  $\mathcal{S}(p)$  is computed using data from each detector alone  $\{\mathcal{S}_1(p), \dots, \mathcal{S}_N(p)\}$  and combining the datasets from all detectors at once  $\mathcal{S}_M(p)$ . Then, a function of single-detector statistics

$F\{\mathcal{S}_1(p), \dots, \mathcal{S}_N(p)\}$  is compared to the multi-detector statistics in order to decide whether the candidate behaves consistently with a CW signal or not. The decision boundary, usually expressed in terms of the difference  $F\{\mathcal{S}_1(p), \dots, \mathcal{S}_N(p)\} - \mathcal{S}_M(p)$ , can be calibrated by means of a software-injection campaign. As discussed in [86, 87, 88, 89], the  $B_{S/GL}$  and  $B_{S/GLtL}$  detection statistics are a Bayesian approach to this implementation of the detector-consistency veto. In their implementation [175, 40, 49], however, the information is processed during search-time, discarding inconsistent candidates at an earlier stage and potentially improving the detection of marginally significant signals.

An example of detector-consistency vetoes can be found in the **SkyHough** contribution to [41], where  $F$  was taken to be the expected multi-detector statistic computed from the single-detector statistics including the varying noise floors but ignoring the antenna pattern modulations. Another example is found in the **Weave** search [51], where  $F$  is simply the *maximum* detection statistic over all the involved detectors. A slight variation was employed in the **BinarySkyHough** analysis of early O3 data [52, 53], where the detection statistic of one of the LIGO detectors was compared against the other one. This was motivated due to the asymmetric behaviour of said detectors during the third observing run, with H1 more affected by noise disturbances than L1.

### 2.3.4 $\chi^2$ vetoes

The second family of vetoes considers the behaviour of a putative CW signal within a particular dataset, namely, whether the detection statistic accumulates throughout the observing run in a way which is more consistent with an instrumental artifact than an astrophysical signal. This is the idea behind the  $\chi^2$  veto, initially proposed in [176] for CBC signals and later implemented for CW signals in [177]. In this case, the dataset is partitioned into  $p$  segments over which a detection statistic is computed  $\{\mathcal{S}_1, \dots, \mathcal{S}_p\}$ . These segment-wise statistics are then compared to the *expected* segment-wise statistics under the presence of a signal in Gaussian noise, usually characterized by a mean and standard deviation  $\mu$  and  $\sigma$ , and combined into a chi-squared discriminant

$$\chi^2\{\mathcal{S}_1, \dots, \mathcal{S}_p\} = \sum_{i=1}^p \left( \frac{\mathcal{S}_i - \mu_i}{\sigma_i} \right)^2. \quad (2.12)$$

Under the assumption of Gaussian noise, this discriminant follows a  $\chi^2$  distribution with  $p - 1$  degrees of freedom; real-data applications, however, must calibrate this test using a suitable injection campaign [41]. In [52, 53], extreme deviations of the segment-wise detection statistic were used to identify stretches of data in which the CW template showed a high degree of overlap with an instrumental feature; this approach is equivalent to using a  $\chi^2$  discriminant in the limit of a very strong sample  $\mathcal{S}_j - \mu_j \gg \sigma_j$ .

### 2.3.5 Vetoing narrow spectral features

The third family of vetoes relies on detector characterization to identify frequency bands in which an instrumental feature is present. CW searches integrate long periods of time looking for quasi-monochromatic signals concentrated around a fraction of a Hertz. Quite often, those narrow bands are populated by narrow spectral features (lines) due to instrumental or environmental disturbances (defective power supplies, blinking LEDs, wind blowing, local fauna interacting with the detector...) which, under very general conditions, are able to mimic the effect produced by a CW signal in the detector, usually producing a high number of candidates in a search [178, 171, 179]. Catalogs listing narrow spectral features and their cause (if known) for the latest runs of the advanced detectors are publicly available [180, 181, 182, 183, 184, 185, 186]. This kind of features are generally not a problem for CBC searches, as those signals sweep wide frequency bands in a relatively short time duration, although noise subtraction techniques are applied in severe cases [187, 188].

A common approach, usually referred to as the (known) *line veto* (see e.g. [41, 42, 45, 52]), is to check whether the frequency evolution of a CW candidate overlaps with any frequency band containing such instrumental features, in which case the candidate is discarded. This requires a high degree of manual intervention, as line catalogs must be created and properly understood, and incurs the risk of removing a genuine signal candidate due to an unfortunate line crossing at a potentially insignificant period of the run.

The DM-off veto was proposed in [189] as a hypothesis-test version of the line veto: it compares the significance of a CW candidate, which includes Doppler modulations due to the Earth’s movement around the SSB, versus the significance obtained after analyzing its surrounding frequency band using an unmodulated template bank (i.e. *without* Doppler Modulation, hence DM-off). This veto was applied with great success in [40]. As happens with the detector-consistency veto and the  $B_{S/GL}$  statistic [88], the DM-off veto can be refactored as a detection statistic in a Bayesian framework to construct a proper line-robust statistic: instead of testing against a single-detector artifact, the hypothetical  $B_{S/GMU}$  statistic would test against an ensemble of *monochromatic* and *unmodulated* signals in any number of detectors [178].

Alternatively, as performed in [40, 190, 49], frequency bins in the data containing lines can be replaced by Gaussian noise drawn from the distribution of neighboring bins to suppress the presence of candidates, preventing any candidate of instrumental origin from polluting the search results. This process is generally performed on SFT data before starting a search.

Short-duration loud instrumental glitches also present a problem to CW searches, as they tend to degrade the noise floor across a wide frequency band. This sort of artifacts, however, are typically dealt with before starting a search using a cleaning procedure such as gating [139, 191, 179, 192].

### 2.3.6 Null-hypothesis vetoes

The fourth family of vetoes are essentially a reformulation of the standard null-hypothesis test, in which a CW candidate is deemed as uninteresting if it is consistent enough with respect to the background noise distribution. A simple proposal, usually referred to as *off-sourcing* [64, 193], is to evaluate a CW candidate on nearly independent noise realizations by shifting its sky position away. This is based on the fact that detector artifacts tend to imprint wider parameter-space regions with significant templates than CW signals. Off-sourced time-frequency tracks are able to break signal-induced correlations while still being affected by instrumental disturbances. The resulting distribution is a proxy of the noise hypothesis’ sampling distribution and can be used to construct significance arguments about the CW candidate of interest.

This veto was adapted by [51] to evaluate the final surviving candidate of a search. In their use case, however, they considered the distribution of the *loudest* candidate from a CW search, that is, considering the number of trials performed by evaluating a template bank on a data stream. Such a distribution has been previously studied in the CW literature [194, 195], but it was not until recently that a method applicable to a generic CW search based on extreme value theory was proposed [172].

Steps towards a fully Bayesian treatment of loudest-candidate null-hypothesis vetoes were taken in [38], which proposed a Bayes factor to evaluate the loudest candidate of a CW search,  $\mathcal{B}_{S/G}^*$ , whose noise-hypothesis component was constructed fitting a Gumbel distribution to the loudest outliers of off-sourced template banks. This approach was used to develop a complete hierarchical follow-up framework, discussed in Sec. 2.4.2.

## 2.4 Follow-up

Follow-up stages are contextualized within a hierarchical search, as discussed in the introduction. They improve the parameter estimation of a CW candidate by imposing tighter constraints on its expected behaviour. This leads to the factual use of simple follow-up stages as signal-consistency veto strategies. Base search stages construct less-sensitive detection statistics by effectually using less-constraining CW signal models. For example, a semicoherent search, in which phase information is contained in discrete, non-overlapping segments, is insensitive to arbitrary phase jumps between coherence segments. In this sense, the sensitivity loss with respect to a fully-coherent search is due to the increased trials factor of this looser family of signal models [43].

There are two ways in which follow-ups may be performed. Following the notation established in [33], we refer to them as *fresh data mode* (FDM) and *recycling data mode* (RDM). As their names suggest, FDM looks for a CW candidate in a new dataset containing brand new information; RDM, on the other hand, re-analyzes the same dataset using a different method. The typical example of FDM is to look for a CW candidate obtained in a certain observing run using data from subsequent observing runs [40, 62, 54]. Examples of RDM include, for example, multi-stage semicoherent or loosely-coherent searches aiming towards a fully-coherent search in a restricted parameter space region [39, 36, 37, 38].

Most CW searches are conducted focusing on a single dataset, usually the latest available observing run of the advanced detectors. Hence, follow-up strategies operate on RDM (even though sometimes FDM assumptions are used for simplicity, as they turn out more conservative and simpler to implement [195, 196]). Nevertheless, this sort of strategies could be detrimental towards detecting a CW signal, as standard RDM operations leading to effectively longer coherence time may lose candidates affected by some form of unmodelled behaviour, such as glitches [197, 198] or accretion-induced spin wandering [148]. Strict FDM must be used in order for a search to follow up a CW candidate with a consistent signal model. Looking into previous observing runs, on the other hand, runs into a sensitivity problem, as marginal candidates may end up completely lost due to the lower quality of the detectors. Examples of searches in which a FDM follow-up looked into a posterior observing run include [40, 55].

### 2.4.1 Single-stage follow-up

The simplest follow-up strategy calibrates a threshold on a different (more sensitive) detection statistic according to some prescription and compares it to the score returned by reanalyzing the CW candidate. As opposed to a veto, this approach points towards evaluating the consistency (or discrepancy) of the CW candidate with respect to a certain population of signals.

The standard `FrequencyHough` follow-up as employed in [41, 45] belongs to this category: baseline Fourier transform length is increased, imposing tighter constraints on the signal model and discarding short-duration candidates. Later searches also make use of the BSD framework [140, 141]. Similar strategies, in this case using the *fully-coherent*  $\mathcal{F}$ -statistic, were proposed in [199, 200].

In order to overcome the curse of dimensionality, `BinarySkyHough` searches [50, 52, 53] employ an MCMC-based follow-up implemented in `PyFstat` [201, 37, 38]. The multi-detector  $\mathcal{F}$ -statistic [21, 70] is used to allow for arbitrarily long coherence times. In this sense, following up a CW candidate is equivalent to sampling the posterior probability distribution of the phase-evolution parameters  $\lambda$  given a data stream [37]

$$P(\lambda|x) \propto e^{\mathcal{F}(\lambda;x)} \cdot P(\lambda), \quad (2.13)$$

where the prior probability distribution  $P(\lambda)$  represents the parameter-space region of interest identified by a search. The result of this grid-less approach is the  $\mathcal{F}$ -statistic evaluated at the loudest candidate of the parameter space at a negligible mismatch. This approach generalizes that of [199, 200], which specializes in the single-stage fully-coherent follow-up of CW candidates.

As discussed in Sec. 2.4.2, the use of MCMC methods simplifies the setup of a generic multi-stage follow-up, as no calibration of parameter-space grids are required. The onus in this case is on the search pipeline to deliver a small-enough prior support for the MCMC to converge. This can always be achieved by starting from a shorter coherence time [37], at the expense of increasing the number of follow-up stages.

A similar strategy is used by the `Time-domain`  $\mathcal{F}$ -statistic follow-up, focusing on the optimization aspect of the procedure. In this case, a max-finding algorithm such as [202] is employed to travel around the parameter space. The result, as with `PyFstat`, are the most favored signal parameters, corresponding to the ones reporting the loudest  $\mathcal{F}$ -statistic value.

### 2.4.2 Multi-stage follow-up

As discussed during Sec. 2.1, multi-stage follow-ups are the natural continuation to a wide parameter-space search after identifying interesting parameter-space regions. Given a data stream, each subsequent follow-up stage operates in RDM, gradually increasing the coherence time with respect to previous stages and, as a consequence, imposing a tighter version of the selected signal model. The effect on a CW candidate is twofold: first, non-astrophysical CW-like artifacts tend to get rejected as coherence time increases (though, as earlier discussed, this could have detrimental effects on more complex CW signals too); second, increasing coherence time results in a refinement in parameter-space resolution, improving the parameter estimation of the candidate at hand. This last effect must be considered carefully, as it also implies an increased number of templates to analyze, quickly becoming unaffordable if parameter-space regions are not gradually narrowed down.

We start by discussing the follow-up strategy introduced in [36], which was applied as a generic follow-up to multiple CW searches with minor modifications [41, 42, 40, 49]. This example is paradigmatic in the sense that it fully exposes the two main challenges of a multi-stage setup. A similar approach,

albeit at much smaller scale, was employed to follow-up `SkyHough` results in [45]. First, one must set up a proper set of parameter-space grids such that CW signals are not lost due to a bad parameter-space covering. If an analytical model of the follow-up method at hand is not available, as is the case in [36], one must resort to an extensive software injection campaign to construct a suitable setup. Second, a criterion must be set up to select/reject CW candidates after performing different stages. If the detection statistic's behaviour across different stages is well understood (see e.g. [196]), an analytical criterion can be derived from first principles; otherwise, a software injection campaign must be used to calibrate a rejection criterion.

Latest developments on the follow-up of CW candidates [37, 38] are able to simplify the setup for the  $\mathcal{F}$ -statistic, although further work is required for its application to generic detection statistics. As discussed in Sec. 2.4.1, the use of MCMC samplers simplifies the setup due to their lack of grids: if a CW signal is within the prior support, parameter-space samplers can get arbitrarily close to the injection parameters given enough time to wander around the prior volume. This argument can be posed quantitatively in terms of the so-called coherence time *ladder* [37], which make use of the parameter-space metric to increase the parameter-space resolution in a controlled manner. Since the follow-up is typically a local analysis, rough estimates of the number of templates using typical parameter-space resolutions are usually a valid approximation [198].

Proper comparison of detection statistics from different stages in a Bayesian framework is currently restricted to the  $\mathcal{F}$ -statistic [196], as the sampling distribution under a signal hypothesis must be known given a (squared) signal-to-noise ratio  $\rho^2$ . This result was used in [38] to propose  $\mathcal{B}_{S/N}^*$ , a (meta) Bayes factor (as the  $\mathcal{F}$ -statistic itself is a Bayes factor) evaluating the result of a multi-stage follow-up, pushing forward the development of a fully-Bayesian follow-up of CW candidates. In this case, the probability under the noise hypothesis was derived from a combination of off-sourcing and extreme value theory [172]. The use of detection statistics (Bayes factors) as data proxies to construct Bayesian arguments is also discussed in [203, 195].

An alternative family of follow-up methods were developed under the name of *loose coherence* [39, 113, 114], already introduced in Sec. 2.2.2. In this case, instead of following the ad hoc recipe of increasing coherence time until a fully-coherent search is achieved, phase information across neighbouring time segments is gradually correlated in a controlled manner by combining (complex) Fourier amplitudes. As opposed to semicoherent methods, which allow for arbitrary phase jumps at the border of a segment, loosely coherent methods allow for phase shifts within pre-specified ranges, depending on the required robustness of the method. Under this framework, semicoherent methods are rediscovered imposing delta-correlation between phases at consecutive time segments. This follow-up approach is fully integrated within the `PowerFlux` and `Falcon` searches [41, 42, 55, 43, 44, 46, 47]

## 2.5 Upper bounds on $h_0$

Since no CW detection has been reported to date, the main data product of CW searches are bounds on the nominal gravitational-wave amplitude  $h_0$  produced by a population of sources consistent with the target signal model. As discussed in Sec. 2.1, astrophysical information can be extracted from this quantity by taking different assumptions, such as the maximum allowed ellipticity from a galactic neutron star at a certain distance from the detector.

Two basic approaches are pursued at production level to derive said upper bounds, depending on whether the aim is for a strict frequentist upper limit or population-based sensitivity estimations. The calibration and establishment of upper bounds of any kind usually involves an extensive software injection campaign in real data with a non-negligible computing cost; due to this, a common approach lies in between both extrema, quoting a proper estimation at a definite set of representative frequency bands and interpolating the results across the rest of the spectrum.

Population-based sensitivity estimations are based on estimating the false-dismissal probability of a search given a certain setup (be it a threshold at a fixed false-alarm probability or a more intricate procedure). The  $p\%$  detection probability amplitude  $h_0^{p\%}$  corresponds then to the amplitude  $h_0$  associated to a false dismissal of  $(100 - p)\%$  after properly marginalizing with respect to other amplitude parameters using a set of priors reflecting the studied source distribution [194]. In practise [41, 42, 40, 45, 49, 52, 50], detection probabilities are usually estimated numerically by means of an injection recovery campaign.

Strict frequentist upper limits, on the other hand, return a conservative estimate of the upper bound, in the sense that false-dismissal probability is *at most*  $(100 - p)\%$ . An example of this kind is the universal statistic procedure [204], which is able to construct strict frequentist upper limits regardless of the underlying noise distribution. This procedure has been extensively combined with the **PowerFlux** and **Falcon** pipelines to efficiently produce robust upper limits under different GW polarization assumptions [41, 42, 55, 43, 44, 46, 47].

To date, all wide parameter-space searches have made use of one of these two upper bounds to report on their results. These upper bounds, however, describe the probability of detecting a signal given an ensemble of equivalent noise realizations, rather than the range where a signal could be found given the data stream at hand. Work towards reporting the latter, Bayesian upper bounds, for wide parameter-space searches was developed in [203, 195].

## 2.6 Summary

We reviewed the methods employed by current wide parameter-space searches for continuous gravitational waves from unknown sources conducted on advanced-detector data. The most widespread approach consists of a hierarchical setup in which parameter-space regions are analyzed using more sensitive (and consequently more expensive) methods as they are gradually narrowed-down. Detecting a CW signal requires both an instrumental and computational effort to confidently unveil such a weak signal using the current generation of gravitational-wave detectors. The use of multiple methods taking different tradeoffs in sensitivity and robustness against instrumental artifacts provides an ideal environment to pursue new strategies towards CW detection and parameter estimation. For the sake of completeness, Table 2.2 provides a comprehensive summary of the search methods reviewed during the present work.

Search	Pipeline	References
All-sky O1	Einstein@Home	[40]
	Falcon	[43, 44]
	FrequencyHough	[41]
	PowerFlux	[41, 42]
	SkyHough	[41, 42]
	Time-domain $\mathcal{F}$ -statistic	[41, 42]
All-sky O2	BinarySkyHough	[50]
	Einstein@Home	[49]
	Falcon	[46, 47, 48]
	FrequencyHough	[45]
	SkyHough	[45]
All-sky O3a	Time-domain $\mathcal{F}$ -statistic	[45]
	BinarySkyHough	[52]
All-sky O3a	PowerFlux	[54]
	Deep exploration O2	Weave
GC O1	PowerFlux	[55]
GC O2	FrequencyHough + BSD	[56]
SNR O1	Einstein@Home	[61]
	Fully-coherent $\mathcal{F}$ -statistic	[59]
SNR O2	Einstein@Home	[62, 69]
	Fully-coherent $\mathcal{F}$ -statistic	[66]
	Viterbi 1.0	[63]
SNR O3a	Viterbi SNR	[65]
	FrequencyHough + BSD	[68]
	Dual-harmonic Viterbi	[68]
	Viterbi SNR	[68]
CDOs in the Solar System O2	Excess power	[9]
Cygnus X-1 O2	Viterbi 1.0	[16]
Scorpius X-1 O1	Cross-Correlation	[58]
	Viterbi 1.0	[57]
Scorpius X-1 O2	Cross-Correlation	[67]
	Viterbi 2.0	[60]
LMXBs O2	Viterbi 2.0	[64]

TABLE 2.2: Summary of CW search methods covered by the present review, grouped by scope and observing run.





# Bibliography

- [1] P. D. Lasky, “Gravitational Waves from Neutron Stars: A Review,” *Publ. Astron. Soc. Austral.*, vol. 32, p. e034, 2015.
- [2] K. Glampedakis and L. Gualtieri, “Gravitational waves from single neutron stars: an advanced detector era survey,” *Astrophys. Space Sci. Libr.*, vol. 457, pp. 673–736, 2018.
- [3] M. Sieniawska and M. Bejger, “Continuous gravitational waves from neutron stars: current status and prospects,” *Universe*, vol. 5, no. 11, p. 217, 2019.
- [4] B. Haskell and K. Schwenzer, “Gravitational waves from isolated neutron stars.” arXiv:2104.03137 [gr-qc].
- [5] R. Essig *et al.*, “Working Group Report: New Light Weakly Coupled Particles,” in *Community Summer Study 2013: Snowmass on the Mississippi*, 10 2013.
- [6] R. Brito, S. Ghosh, E. Barausse, E. Berti, V. Cardoso, I. Dvorkin, A. Klein, and P. Pani, “Gravitational wave searches for ultralight bosons with LIGO and LISA,” *Phys. Rev. D*, vol. 96, no. 6, p. 064050, 2017.
- [7] R. Brito, S. Ghosh, E. Barausse, E. Berti, V. Cardoso, I. Dvorkin, A. Klein, and P. Pani, “Stochastic and resolvable gravitational waves from ultralight bosons,” *Phys. Rev. Lett.*, vol. 119, no. 13, p. 131101, 2017.
- [8] S. J. Zhu, M. Baryakhtar, M. A. Papa, D. Tsuna, N. Kawanaka, and H.-B. Eggenstein, “Characterizing the continuous gravitational-wave signal from boson clouds around Galactic isolated black holes,” *Phys. Rev. D*, vol. 102, no. 6, p. 063020, 2020.
- [9] C. J. Horowitz, M. A. Papa, and S. Reddy, “Gravitational waves from compact dark matter objects in the solar system,” *Phys. Lett. B*, vol. 800, p. 135072, 2020.
- [10] M. Isi, A. J. Weinstein, C. Mead, and M. Pitkin, “Detecting Beyond-Einstein Polarizations of Continuous Gravitational Waves,” *Phys. Rev. D*, vol. 91, no. 8, p. 082002, 2015.
- [11] M. Isi, M. Pitkin, and A. J. Weinstein, “Probing Dynamical Gravity with the Polarization of Continuous Gravitational Waves,” *Phys. Rev. D*, vol. 96, no. 4, p. 042001, 2017.
- [12] M. Isi, L. Sun, R. Brito, and A. Melatos, “Directed searches for gravitational waves from ultralight bosons,” *Phys. Rev. D*, vol. 99, no. 8, p. 084042, 2019. [Erratum: *Phys.Rev.D* 102, 049901 (2020)].
- [13] R. Xu, Y. Gao, and L. Shao, “Precession of spheroids under Lorentz violation and observational consequences for neutron stars,” *Phys. Rev. D*, vol. 103, no. 8, p. 084028, 2021.
- [14] R. Xu, Y. Gao, and L. Shao, “Signatures of Lorentz Violation in Continuous Gravitational-Wave Spectra of Ellipsoidal Neutron Stars,” *Galaxies*, vol. 9, no. 1, p. 12, 2021.
- [15] C. Palomba *et al.*, “Direct constraints on ultra-light boson mass from searches for continuous gravitational waves,” *Phys. Rev. Lett.*, vol. 123, p. 171101, 2019.
- [16] L. Sun, R. Brito, and M. Isi, “Search for ultralight bosons in Cygnus X-1 with Advanced LIGO,” *Phys. Rev. D*, vol. 101, no. 6, p. 063020, 2020. [Erratum: *Phys.Rev.D* 102, 089902 (2020)].
- [17] A. L. Miller, S. Clesse, F. De Lillo, G. Bruno, A. Depasse, and A. Tanasijczuk, “Probing planetary-mass primordial black holes with continuous gravitational waves,” *Phys. Dark Univ.*, vol. 32, p. 100836, 2021.

- [18] A. L. Miller, N. Aggarwal, S. Clesse, and F. De Lillo, “Constraints on planetary and asteroid-mass primordial black holes from continuous gravitational-wave searches.” arXiv:2110.06188 [gr-qc].
- [19] R. Prix, *Gravitational Waves from Spinning Neutron Stars*, vol. 357 of *Astrophys. Space Sci. Lib.*, ch. 24, pp. 651–685. Berlin Heidelberg: Springer, 2009.
- [20] K. Riles, “Recent searches for continuous gravitational waves,” *Mod. Phys. Lett. A*, vol. 32, no. 39, p. 1730035, 2017.
- [21] P. Jaranowski, A. Krolak, and B. F. Schutz, “Data analysis of gravitational - wave signals from spinning neutron stars. 1. The Signal and its detection,” *Phys. Rev. D*, vol. 58, p. 063001, 1998.
- [22] J. Aasi *et al.*, “Advanced LIGO,” *Class. Quant. Grav.*, vol. 32, p. 074001, 2015.
- [23] F. Acernese *et al.*, “Advanced Virgo: a second-generation interferometric gravitational wave detector,” *Class. Quant. Grav.*, vol. 32, no. 2, p. 024001, 2015.
- [24] T. Akutsu *et al.*, “KAGRA: 2.5 Generation Interferometric Gravitational Wave Detector,” *Nature Astron.*, vol. 3, no. 1, pp. 35–40, 2019.
- [25] M. Maggiore *et al.*, “Science Case for the Einstein Telescope,” *JCAP*, vol. 03, p. 050, 2020.
- [26] D. Reitze *et al.*, “Cosmic Explorer: The U.S. Contribution to Gravitational-Wave Astronomy beyond LIGO,” *Bull. Am. Astron. Soc.*, vol. 51, no. 7, p. 035, 2019.
- [27] B. P. Abbott *et al.*, “A guide to LIGO–Virgo detector noise and extraction of transient gravitational-wave signals,” *Class. Quant. Grav.*, vol. 37, no. 5, p. 055002, 2020.
- [28] P. R. Brady, T. Creighton, C. Cutler, and B. F. Schutz, “Searching for periodic sources with LIGO,” *Phys. Rev. D*, vol. 57, pp. 2101–2116, 1998.
- [29] R. Prix, “Template-based searches for gravitational waves: Efficient lattice covering of flat parameter spaces,” *Class. Quant. Grav.*, vol. 24, pp. S481–S490, 2007.
- [30] K. Wette, “Lattice template placement for coherent all-sky searches for gravitational-wave pulsars,” *Phys. Rev. D*, vol. 90, no. 12, p. 122010, 2014.
- [31] S. Fairhurst, “Source localization with an advanced gravitational wave detector network,” *Class. Quant. Grav.*, vol. 28, p. 105021, 2011.
- [32] P. R. Brady and T. Creighton, “Searching for periodic sources with LIGO II: Hierarchical searches,” *Phys. Rev. D*, vol. 61, p. 082001, 2000.
- [33] C. Cutler, I. Gholami, and B. Krishnan, “Improved stack-slide searches for gravitational-wave pulsars,” *Phys. Rev. D*, vol. 72, p. 042004, 2005.
- [34] R. Prix and M. Shaltev, “Search for Continuous Gravitational Waves: Optimal StackSlide method at fixed computing cost,” *Phys. Rev. D*, vol. 85, p. 084010, 2012.
- [35] M. Shaltev, “Optimizing the StackSlide setup and data selection for continuous-gravitational-wave searches in realistic detector data,” *Phys. Rev. D*, vol. 93, no. 4, p. 044058, 2016.
- [36] M. A. Papa *et al.*, “Hierarchical follow-up of subthreshold candidates of an all-sky Einstein@Home search for continuous gravitational waves on LIGO sixth science run data,” *Phys. Rev. D*, vol. 94, no. 12, p. 122006, 2016.
- [37] G. Ashton and R. Prix, “Hierarchical multistage MCMC follow-up of continuous gravitational wave candidates,” *Phys. Rev. D*, vol. 97, no. 10, p. 103020, 2018.
- [38] R. Tenorio, D. Keitel, and A. M. Sintes, “Application of a hierarchical MCMC follow-up to Advanced LIGO continuous gravitational-wave candidates,” *Phys. Rev. D*, vol. 104, no. 8, p. 084012, 2021.

- [39] V. Dergachev, “On blind searches for noise dominated signals: a loosely coherent approach,” *Class. Quant. Grav.*, vol. 27, p. 205017, 2010.
- [40] B. P. Abbott *et al.*, “First low-frequency Einstein@Home all-sky search for continuous gravitational waves in Advanced LIGO data,” *Phys. Rev. D*, vol. 96, no. 12, p. 122004, 2017.
- [41] B. P. Abbott *et al.*, “All-sky Search for Periodic Gravitational Waves in the O1 LIGO Data,” *Phys. Rev. D*, vol. 96, no. 6, p. 062002, 2017.
- [42] B. P. Abbott *et al.*, “Full Band All-sky Search for Periodic Gravitational Waves in the O1 LIGO Data,” *Phys. Rev. D*, vol. 97, no. 10, p. 102003, 2018.
- [43] V. Dergachev and M. A. Papa, “Sensitivity improvements in the search for periodic gravitational waves using O1 LIGO data,” *Phys. Rev. Lett.*, vol. 123, no. 10, p. 101101, 2019.
- [44] V. Dergachev and M. A. Papa, “Results from an Extended Falcon All-Sky Survey for Continuous Gravitational Waves,” *Phys. Rev. D*, vol. 101, no. 2, p. 022001, 2020.
- [45] B. P. Abbott *et al.*, “All-sky search for continuous gravitational waves from isolated neutron stars using Advanced LIGO O2 data,” *Phys. Rev. D*, vol. 100, no. 2, p. 024004, 2019.
- [46] V. Dergachev and M. A. Papa, “Results from the First All-Sky Search for Continuous Gravitational Waves from Small-Ellipticity Sources,” *Phys. Rev. Lett.*, vol. 125, no. 17, p. 171101, 2020.
- [47] V. Dergachev and M. A. Papa, “Results from high-frequency all-sky search for continuous gravitational waves from small-ellipticity sources,” *Phys. Rev. D*, vol. 103, no. 6, p. 063019, 2021.
- [48] V. Dergachev and M. A. Papa, “Search for continuous gravitational waves from small-ellipticity sources at low frequencies,” *Phys. Rev. D*, vol. 104, no. 4, p. 043003, 2021.
- [49] B. Steltner, M. A. Papa, H. B. Eggenstein, B. Allen, V. Dergachev, R. Prix, B. Machenschalk, S. Walsh, S. J. Zhu, and S. Kwang, “Einstein@Home All-sky Search for Continuous Gravitational Waves in LIGO O2 Public Data,” *Astrophys. J.*, vol. 909, no. 1, p. 79, 2021.
- [50] P. B. Covas and A. M. Sintes, “First all-sky search for continuous gravitational-wave signals from unknown neutron stars in binary systems using Advanced LIGO data,” *Phys. Rev. Lett.*, vol. 124, no. 19, p. 191102, 2020.
- [51] K. Wette, L. Dunn, P. Clearwater, and A. Melatos, “Deep exploration for continuous gravitational waves at 171–172 Hz in LIGO second observing run data,” *Phys. Rev. D*, vol. 103, no. 8, p. 083020, 2021.
- [52] R. Abbott *et al.*, “All-sky search in early O3 LIGO data for continuous gravitational-wave signals from unknown neutron stars in binary systems,” *Phys. Rev. D*, vol. 103, no. 6, p. 064017, 2021.
- [53] R. Tenorio, “An all-sky search in early O3 LIGO data for continuous gravitational-wave signals from unknown neutron stars in binary systems,” in *55th Rencontres de Moriond on Gravitation*, 2021.
- [54] R. Abbott *et al.*, “All-sky search for continuous gravitational waves from isolated neutron stars in the early O3 LIGO data,” *Phys. Rev. D*, vol. 104, no. 8, p. 082004, 2021.
- [55] V. Dergachev, M. A. Papa, B. Steltner, and H.-B. Eggenstein, “Loosely coherent search in LIGO O1 data for continuous gravitational waves from Terzan 5 and the galactic center,” *Phys. Rev. D*, vol. 99, no. 8, p. 084048, 2019.
- [56] O. J. Piccinni, P. Astone, S. D’Antonio, S. Frasca, G. Intini, I. La Rosa, P. Leaci, S. Mastrogiovanni, A. Miller, and C. Palomba, “Directed search for continuous gravitational-wave signals from the Galactic Center in the Advanced LIGO second observing run,” *Phys. Rev. D*, vol. 101, no. 8, p. 082004, 2020.
- [57] B. P. Abbott *et al.*, “Search for gravitational waves from Scorpius X-1 in the first Advanced LIGO observing run with a hidden Markov model,” *Phys. Rev. D*, vol. 95, no. 12, p. 122003, 2017.

- [58] B. P. Abbott *et al.*, “Upper Limits on Gravitational Waves from Scorpius X-1 from a Model-Based Cross-Correlation Search in Advanced LIGO Data,” *Astrophys. J.*, vol. 847, no. 1, p. 47, 2017.
- [59] B. P. Abbott *et al.*, “Searches for Continuous Gravitational Waves from 15 Supernova Remnants and Fomalhaut b with Advanced LIGO,” *Astrophys. J.*, vol. 875, no. 2, p. 122, 2019. [Erratum: *Astrophys. J.* 918, 91 (2021)].
- [60] B. P. Abbott *et al.*, “Search for gravitational waves from Scorpius X-1 in the second Advanced LIGO observing run with an improved hidden Markov model,” *Phys. Rev. D*, vol. 100, no. 12, p. 122002, 2019.
- [61] J. Ming *et al.*, “Results from an Einstein@Home search for continuous gravitational waves from Cassiopeia A, Vela Jr. and G347.3,” *Phys. Rev. D*, vol. 100, no. 2, p. 024063, 2019.
- [62] M. A. Papa, J. Ming, E. V. Gotthelf, B. Allen, R. Prix, V. Dergachev, H.-B. Eggenstein, A. Singh, and S. J. Zhu, “Search for Continuous Gravitational Waves from the Central Compact Objects in Supernova Remnants Cassiopeia A, Vela Jr., and G347.3–0.5,” *Astrophysical Journal*, vol. 897, no. 1, p. 22, 2020.
- [63] M. Millhouse, L. Strang, and A. Melatos, “Search for gravitational waves from 12 young supernova remnants with a hidden Markov model in Advanced LIGO’s second observing run,” *Phys. Rev. D*, vol. 102, no. 8, p. 083025, 2020.
- [64] H. Middleton, P. Clearwater, A. Melatos, and L. Dunn, “Search for gravitational waves from five low mass X-ray binaries in the second Advanced LIGO observing run with an improved hidden Markov model,” *Phys. Rev. D*, vol. 102, no. 2, p. 023006, 2020.
- [65] D. Jones and L. Sun, “Search for continuous gravitational waves from Fomalhaut b in the second Advanced LIGO observing run with a hidden Markov model,” *Phys. Rev. D*, vol. 103, no. 2, p. 023020, 2021.
- [66] L. Lindblom and B. J. Owen, “Directed searches for continuous gravitational waves from twelve supernova remnants in data from Advanced LIGO’s second observing run,” *Phys. Rev. D*, vol. 101, no. 8, p. 083023, 2020.
- [67] Y. Zhang, M. A. Papa, B. Krishnan, and A. L. Watts, “Search for Continuous Gravitational Waves from Scorpius X-1 in LIGO O2 Data,” *Astrophys. J. Lett.*, vol. 906, no. 2, p. L14, 2021.
- [68] R. Abbott *et al.*, “Searches for continuous gravitational waves from young supernova remnants in the early third observing run of Advanced LIGO and Virgo.” arXiv:2105.11641 [astro-ph.HE].
- [69] J. Ming, M. A. Papa, H.-B. Eggenstein, B. Machenschalk, B. Steltner, R. Prix, B. Allen, and O. Behnke, “Results from an Einstein@Home search for continuous gravitational waves from G347.3 at low frequencies in LIGO O2 data.” arXiv:2108.02808 [gr-qc].
- [70] C. Cutler and B. F. Schutz, “The Generalized F-statistic: Multiple detectors and multiple GW pulsars,” *Phys. Rev. D*, vol. 72, p. 063006, 2005.
- [71] R. Prix and B. Krishnan, “Targeted search for continuous gravitational waves: Bayesian versus maximum-likelihood statistics,” *Class. Quant. Grav.*, vol. 26, p. 204013, 2009.
- [72] J. T. Whelan, R. Prix, C. J. Cutler, and J. L. Willis, “New Coordinates for the Amplitude Parameter Space of Continuous Gravitational Waves,” *Class. Quant. Grav.*, vol. 31, p. 065002, 2014.
- [73] R. Prix, S. Giampanis, and C. Messenger, “Search method for long-duration gravitational-wave transients from neutron stars,” *Phys. Rev. D*, vol. 84, p. 023007, 2011.
- [74] S. Dhurandhar, B. Krishnan, and J. L. Willis, “Marginalizing the likelihood function for modeled gravitational wave searches.” arXiv:1707.08163 [gr-qc].
- [75] J. J. Bero and J. T. Whelan, “An Analytic Approximation to the Bayesian Detection Statistic for Continuous Gravitational Waves,” *Class. Quant. Grav.*, vol. 36, no. 1, p. 015013, 2019. [Erratum: *Class. Quant. Grav.* 36, 049601 (2019)].

- [76] K. Wette, “Geometric Approach to Analytic Marginalisation of the Likelihood Ratio for Continuous Gravitational Wave Searches,” *Universe*, vol. 7, no. 6, p. 174.
- [77] R. Prix and J. T. Whelan, “F-statistic search for white-dwarf binaries in the first Mock LISA Data Challenge,” *Class. Quant. Grav.*, vol. 24, no. 19, pp. S565–S574, 2007.
- [78] D. Keppel, “The multi-detector F-statistic metric for short-duration non-precessing inspiral gravitational-wave signals,” *Phys. Rev. D*, vol. 86, p. 123010, 2012.
- [79] S. B. Anderson, *A study of recycled pulsars in globular clusters*. PhD thesis, California Institute of Technology, 1993.
- [80] R. Prix and Y. Itoh, “Global parameter-space correlations of coherent searches for continuous gravitational waves,” *Class. Quant. Grav.*, vol. 22, pp. S1003–S1012, 2005.
- [81] H. J. Pletsch, “Parameter-space correlations of the optimal statistic for continuous gravitational-wave detection,” *Phys. Rev. D*, vol. 78, p. 102005, 2008.
- [82] H. J. Pletsch and B. Allen, “Exploiting global correlations to detect continuous gravitational waves,” *Phys. Rev. Lett.*, vol. 103, p. 181102, 2009.
- [83] H. J. Pletsch, “Parameter-space metric of semicoherent searches for continuous gravitational waves,” *Phys. Rev. D*, vol. 82, p. 042002, 2010.
- [84] S. Walsh, K. Wette, M. A. Papa, and R. Prix, “Optimizing the choice of analysis method for all-sky searches for continuous gravitational waves with Einstein@Home,” *Phys. Rev. D*, vol. 99, no. 8, p. 082004, 2019.
- [85] D. P. Anderson, “BOINC: A Platform for Volunteer Computing.” arXiv:1903.01699 [cs.DC].
- [86] D. Keitel, R. Prix, M. A. Papa, P. Leaci, and M. Siddiqi, “Search for continuous gravitational waves: Improving robustness versus instrumental artifacts,” *Phys. Rev. D*, vol. 89, no. 6, p. 064023, 2014.
- [87] D. Keitel and R. Prix, “Line-robust statistics for continuous gravitational waves: safety in the case of unequal detector sensitivities,” *Class. Quant. Grav.*, vol. 32, no. 3, p. 035004, 2015.
- [88] D. Keitel, “Robust semicoherent searches for continuous gravitational waves with noise and signal models including hours to days long transients,” *Phys. Rev. D*, vol. 93, no. 8, p. 084024, 2016.
- [89] D. Keitel, “Distinguishing transient signals and instrumental disturbances in semi-coherent searches for continuous gravitational waves with line-robust statistics,” *J. Phys. Conf. Ser.*, vol. 716, no. 1, p. 012003, 2016.
- [90] K. Wette, “Parameter-space metric for all-sky semicoherent searches for gravitational-wave pulsars,” *Phys. Rev. D*, vol. 92, no. 8, p. 082003, 2015.
- [91] J. Ming, B. Krishnan, M. A. Papa, C. Aulbert, and H. Fehrmann, “Optimal directed searches for continuous gravitational waves,” *Phys. Rev. D*, vol. 93, no. 6, p. 064011, 2016.
- [92] B. J. Owen, “Search templates for gravitational waves from inspiraling binaries: Choice of template spacing,” *Phys. Rev. D*, vol. 53, pp. 6749–6761, 1996.
- [93] R. Prix, “Search for continuous gravitational waves: Metric of the multi-detector F-statistic,” *Phys. Rev. D*, vol. 75, p. 023004, 2007. [Erratum: *Phys.Rev.D* 75, 069901 (2007)].
- [94] B. Allen, “Optimal template banks,” *Phys. Rev. D*, vol. 104, no. 4, p. 042005, 2021.
- [95] K. Wette and R. Prix, “Flat parameter-space metric for all-sky searches for gravitational-wave pulsars,” *Phys. Rev. D*, vol. 88, no. 12, p. 123005, 2013.
- [96] K. Wette, “Empirically extending the range of validity of parameter-space metrics for all-sky searches for gravitational-wave pulsars,” *Phys. Rev. D*, vol. 94, no. 12, p. 122002, 2016.

- [97] K. Wette, S. Walsh, R. Prix, and M. A. Papa, “Implementing a semicoherent search for continuous gravitational waves using optimally-constructed template banks,” *Phys. Rev. D*, vol. 97, no. 12, p. 123016, 2018.
- [98] B. Allen, “Spherical ansatz for parameter-space metrics,” *Phys. Rev. D*, vol. 100, no. 12, p. 124004, 2019.
- [99] B. Allen and A. A. Shoom, “Template banks based on  $\mathbb{Z}^n$  and  $A_n^*$  lattices.” arXiv:2102.11631 [astro-ph.IM].
- [100] R. Prix, “The F-statistic and its implementation in ComputeFStatistic\_v2.” <https://dcc.ligo.org/LIGO-T0900149/public>, 2006.
- [101] LIGO Scientific Collaboration, “LIGO Algorithm Library - LALSuite.” free software (GPL), 2018.
- [102] J. Aasi *et al.*, “Implementation of an  $\mathcal{F}$ -statistic all-sky search for continuous gravitational waves in Virgo VSR1 data,” *Class. Quant. Grav.*, vol. 31, p. 165014, 2014.
- [103] A. Piskarski and P. Jaranowski, “Banks of templates for all-sky narrow-band searches of gravitational waves from spinning neutron stars,” *Class. Quant. Grav.*, vol. 32, no. 14, p. 145014, 2015.
- [104] G. Poghosyan, S. Matta, A. Streit, M. Bejger, and A. Królak, “Architecture, implementation and parallelization of the software to search for periodic gravitational wave signals,” *Comput. Phys. Commun.*, vol. 188, pp. 167–176, 2015.
- [105] M. Sieniawska, M. Bejger, and A. Krolak, “Follow-up procedure for gravitational wave searches from isolated neutron stars using the time-domain  $\mathcal{F}$ -statistic method,” *Class. Quant. Grav.*, vol. 36, p. 225008, Nov. 2019.
- [106] F. Morawski, M. Bejger, and P. Ciecielag, “Convolutional neural network classifier for the output of the time-domain F-statistic all-sky search for continuous gravitational waves,” *Machine Learning: Science and Technology*, p. arXiv:1907.06917, 2019.
- [107] B. Krishnan, A. M. Sintes, M. A. Papa, B. F. Schutz, S. Frasca, and C. Palomba, “The Hough transform search for continuous gravitational waves,” *Phys. Rev. D*, vol. 70, p. 082001, 2004.
- [108] B. Allen, M. A. Papa, and B. F. Schutz, “Optimal strategies for sinusoidal signal detection,” *Phys. Rev. D*, vol. 66, p. 102003, 2002.
- [109] V. Dergachev, “Description of PowerFlux algorithms and implementation.” <https://dcc.ligo.org/LIGO-T050186/public>, 2006.
- [110] V. Dergachev, “Description of PowerFlux 2 algorithms and implementation.” <https://dcc.ligo.org/LIGO-T1000272/public>, 2011.
- [111] V. Dergachev and K. Riles, “PowerFlux polarization analysis.” <https://dcc.ligo.org/LIGO-T050187/public>, 2006.
- [112] G. Mendell and K. Wette, “Using generalized PowerFlux methods to estimate the parameters of periodic gravitational waves,” *Class. Quant. Grav.*, vol. 25, p. 114044, 2008.
- [113] V. Dergachev, “Loosely coherent searches for sets of well-modeled signals,” *Phys. Rev. D*, vol. 85, p. 062003, 2012.
- [114] V. Dergachev, “Loosely coherent searches for medium scale coherence lengths,” *arXiv e-print*, 2018.
- [115] A. Singh, M. A. Papa, and V. Dergachev, “Characterizing the sensitivity of isolated continuous gravitational wave searches to binary orbits,” *Phys. Rev. D*, vol. 100, no. 2, p. 024058, 2019.
- [116] S. Dhurandhar, B. Krishnan, H. Mukhopadhyay, and J. T. Whelan, “Cross-correlation search for periodic gravitational waves,” *Phys. Rev. D*, vol. 77, p. 082001, 2008.
- [117] J. T. Whelan, S. Sundaesan, Y. Zhang, and P. Peiris, “Model-Based Cross-Correlation Search for Gravitational Waves from Scorpius X-1,” *Phys. Rev. D*, vol. 91, p. 102005, 2015.

- [118] K. J. Wagner, J. T. Whelan, J. K. Wofford, and K. Wette, “Template Lattices for a Cross-Correlation Search for Gravitational Waves from Scorpius X-1.” arXiv:2106.16142 [gr-qc].
- [119] S. Mitra, S. Dhurandhar, T. Souradeep, A. Lazzarini, V. Mandic, S. Bose, and S. Ballmer, “Gravitational wave radiometry: Mapping a stochastic gravitational wave background,” *Phys. Rev. D*, vol. 77, p. 042002, 2008.
- [120] A. Ain, P. Dalvi, and S. Mitra, “Fast Gravitational Wave Radiometry using Data Folding,” *Phys. Rev. D*, vol. 92, no. 2, p. 022003, 2015.
- [121] A. Ain, J. Suresh, and S. Mitra, “Very fast stochastic gravitational wave background map making using folded data,” *Phys. Rev. D*, vol. 98, no. 2, p. 024001, 2018.
- [122] B. P. Abbott *et al.*, “Directional Limits on Persistent Gravitational Waves from Advanced LIGO’s First Observing Run,” *Phys. Rev. Lett.*, vol. 118, no. 12, p. 121102, 2017.
- [123] B. P. Abbott *et al.*, “Directional limits on persistent gravitational waves using data from Advanced LIGO’s first two observing runs,” *Phys. Rev. D*, vol. 100, no. 6, p. 062001, 2019.
- [124] R. Abbott *et al.*, “All-sky, all-frequency directional search for persistent gravitational-waves from Advanced LIGO’s and Advanced Virgo’s first three observing runs.” arXiv:2110.09834 [gr-qc].
- [125] G. D. Meadors, B. Krishnan, M. A. Papa, J. T. Whelan, and Y. Zhang, “Resampling to accelerate cross-correlation searches for continuous gravitational waves from binary systems,” *Phys. Rev. D*, vol. 97, p. 044017, Feb. 2018.
- [126] C. Palomba, P. Astone, and S. Frasca, “Adaptive Hough transform for the search of periodic sources,” *Class. Quant. Grav.*, vol. 22, pp. S1255–S1264, 2005.
- [127] B. Krishnan and A. M. Sintes, “Hough search with improved sensitivity.” <https://dcc.ligo.org/LIGO-T070124-x0/public>, 2007.
- [128] P. V. Hough, “Method and means for recognizing complex patterns,” *U. S. Patent*, 12 1962.
- [129] A. M. Sintes and B. Krishnan, “Improved hough search for gravitational wave pulsars,” *J. Phys. Conf. Ser.*, vol. 32, pp. 206–211, 2006.
- [130] R. Tenorio, D. Keitel, and A. M. Sintes, “Time-frequency track distance for comparing continuous gravitational wave signals,” *Phys. Rev. D*, vol. 103, no. 6, p. 064053, 2021.
- [131] L. Sancho de la Jordana, “Hierarchical Hough all-sky search for periodic gravitational waves in LIGO S5 data,” *J. Phys. Conf. Ser.*, vol. 228, p. 012004, 2010.
- [132] P. B. Covas, *Searching for continuous gravitational waves with advanced LIGO*. PhD thesis, Universitat de les Illes Balears, 2020.
- [133] P. B. Covas and A. M. Sintes, “New method to search for continuous gravitational waves from unknown neutron stars in binary systems,” *Phys. Rev. D*, vol. 99, no. 12, p. 124019, 2019.
- [134] M. Oliver, D. Keitel, and A. M. Sintes, “Adaptive transient Hough method for long-duration gravitational wave transients,” *Phys. Rev. D*, vol. 99, no. 10, p. 104067, 2019.
- [135] F. Antonucci, P. Astone, S. D’Antonio, S. Frasca, and C. Palomba, “Detection of periodic gravitational wave sources by Hough transform in the  $f$  versus  $f(\cdot)$  plane,” *Class. Quant. Grav.*, vol. 25, p. 184015, 2008.
- [136] P. Astone, A. Colla, S. D’Antonio, S. Frasca, and C. Palomba, “Method for all-sky searches of continuous gravitational wave signals using the frequency-Hough transform,” *Phys. Rev. D*, vol. 90, no. 4, p. 042002, 2014.
- [137] P. Leaci, “Methods to filter out spurious disturbances in continuous-wave searches from gravitational-wave detectors,” vol. 90, p. 125001, nov 2015.

- [138] G. Intini, P. Leaci, P. Astone, S. D. Antonio, S. Frasca, I. L. Rosa, A. Miller, C. Palomba, and O. Piccinni, “A doppler-modulation based veto to discard false continuous gravitational-wave candidates,” vol. 37, p. 225007, oct 2020.
- [139] P. Astone, S. Frasca, and C. Palomba, “The short FFT database and the peak map for the hierarchical search of periodic sources,” *Class. Quant. Grav.*, vol. 22, no. 18, pp. S1197–S1210, 2005.
- [140] O. J. Piccinni, P. Astone, S. D’Antonio, S. Frasca, G. Intini, P. Leaci, S. Mastrogiovanni, A. Miller, C. Palomba, and A. Singhal, “A new data analysis framework for the search of continuous gravitational wave signals,” *Class. Quant. Grav.*, vol. 36, no. 1, p. 015008, 2019.
- [141] O. J. Piccinni and S. Frasca, “The band-sampled-data collection for the search of continuous gravitational wave signals,” in *2018 26th European Signal Processing Conference (EUSIPCO)*, pp. 2653–2657, 2018.
- [142] A. Miller *et al.*, “Method to search for long duration gravitational wave transients from isolated neutron stars using the generalized frequency-Hough transform,” *Phys. Rev. D*, vol. 98, no. 10, p. 102004, 2018.
- [143] B. P. Abbott *et al.*, “Search for gravitational waves from a long-lived remnant of the binary neutron star merger GW170817,” *Astrophys. J.*, vol. 875, no. 2, p. 160, 2019.
- [144] S. D’Antonio *et al.*, “Semicoherent analysis method to search for continuous gravitational waves emitted by ultralight boson clouds around spinning black holes,” *Phys. Rev. D*, vol. 98, no. 10, p. 103017, 2018.
- [145] A. L. Miller *et al.*, “Probing new light gauge bosons with gravitational-wave interferometers using an adapted semicoherent method,” *Phys. Rev. D*, vol. 103, no. 10, p. 103002, 2021.
- [146] R. Abbott *et al.*, “Constraints on dark photon dark matter using data from LIGO’s and Virgo’s third observing run.” arXiv:2105.13085 [astro-ph.CO].
- [147] G. Ashton, D. I. Jones, and R. Prix, “Effect of timing noise on targeted and narrow-band coherent searches for continuous gravitational waves from pulsars,” *Phys. Rev. D*, vol. 91, p. 062009, Mar. 2015.
- [148] A. Mukherjee, C. Messenger, and K. Riles, “Accretion-induced spin-wandering effects on the neutron star in Scorpius X-1: Implications for continuous gravitational wave searches,” *Phys. Rev. D*, vol. 97, no. 4, p. 043016, 2018.
- [149] S. Suvorova, L. Sun, A. Melatos, W. Moran, and R. J. Evans, “Hidden Markov model tracking of continuous gravitational waves from a neutron star with wandering spin,” *Phys. Rev. D*, vol. 93, no. 12, p. 123009, 2016.
- [150] S. Suvorova, P. Clearwater, A. Melatos, L. Sun, W. Moran, and R. J. Evans, “Hidden Markov model tracking of continuous gravitational waves from a binary neutron star with wandering spin. II. Binary orbital phase tracking,” *Phys. Rev. D*, vol. 96, p. 102006, Nov. 2017.
- [151] L. Sun, A. Melatos, S. Suvorova, W. Moran, and R. J. Evans, “Hidden Markov model tracking of continuous gravitational waves from young supernova remnants,” *Phys. Rev. D*, vol. 97, p. 043013, Feb. 2018.
- [152] L. Sun, A. Melatos, and P. D. Lasky, “Tracking continuous gravitational waves from a neutron star at once and twice the spin frequency with a hidden Markov model,” *Phys. Rev. D*, vol. 99, p. 123010, June 2019.
- [153] A. Melatos, P. Clearwater, S. Suvorova, L. Sun, W. Moran, and R. J. Evans, “Hidden Markov model tracking of continuous gravitational waves from a binary neutron star with wandering spin. III. Rotational phase tracking,” *Phys. Rev. D*, vol. 104, no. 4, p. 042003, 2021.



- [154] J. Bayley, C. Messenger, and G. Woan, “Generalized application of the Viterbi algorithm to searches for continuous gravitational-wave signals,” *Phys. Rev. D*, vol. 100, no. 2, p. 023006, 2019.
- [155] E. T. Jaynes, *Probability Theory: The Logic of Science*. Cambridge University Press, 2003.
- [156] D. Beniwal, P. Clearwater, L. Dunn, A. Melatos, and D. Ottaway, “Search for continuous gravitational waves from ten H.E.S.S. sources using a hidden Markov model,” *Phys. Rev. D*, vol. 103, no. 8, p. 083009, 2021.
- [157] R. Abbott *et al.*, “Search for continuous gravitational waves from 20 accreting millisecond X-ray pulsars in O3 LIGO data.” arXiv:2109.09255 [astro-ph.HE].
- [158] L. Sun and A. Melatos, “Application of hidden Markov model tracking to the search for long-duration transient gravitational waves from the remnant of the binary neutron star merger GW170817,” *Phys. Rev. D*, vol. 99, p. 123003, June 2019.
- [159] A. Viterbi, “Error bounds for convolutional codes and an asymptotically optimum decoding algorithm,” *IEEE Transactions on Information Theory*, vol. 13, no. 2, pp. 260–269, 1967.
- [160] L. Sammut, C. Messenger, A. Melatos, and B. J. Owen, “Implementation of the frequency-modulated sideband search method for gravitational waves from low mass X-ray binaries,” *Phys. Rev. D*, vol. 89, no. 4, p. 043001, 2014.
- [161] G. L. Bretthorst, *Bayesian spectrum Analysis and parameter estimation*. Springer-Verlag Berlin Heidelberg, 1988.
- [162] E. Cuoco *et al.*, “Enhancing Gravitational-Wave Science with Machine Learning,” *Mach. Learn. Sci. Tech.*, vol. 2, no. 1, p. 011002, 2021.
- [163] C. Dreissigacker, R. Sharma, C. Messenger, R. Zhao, and R. Prix, “Deep-Learning Continuous Gravitational Waves,” *Phys. Rev. D*, vol. 100, no. 4, p. 044009, 2019.
- [164] C. Dreissigacker and R. Prix, “Deep-Learning Continuous Gravitational Waves: Multiple detectors and realistic noise,” *Phys. Rev. D*, vol. 102, no. 2, p. 022005, 2020.
- [165] N. Sarin and P. D. Lasky, “The evolution of binary neutron star post-merger remnants: a review,” *Gen. Rel. Grav.*, vol. 53, no. 6, p. 59, 2021.
- [166] A. L. Miller *et al.*, “How effective is machine learning to detect long transient gravitational waves from neutron stars in a real search?,” *Phys. Rev. D*, vol. 100, no. 6, p. 062005, 2019.
- [167] T. S. Yamamoto and T. Tanaka, “Use of an excess power method and a convolutional neural network in an all-sky search for continuous gravitational waves,” *Phys. Rev. D*, vol. 103, no. 8, p. 084049, 2021.
- [168] J. Bayley, C. Messenger, and G. Woan, “Robust machine learning algorithm to search for continuous gravitational waves,” *Phys. Rev. D*, vol. 102, no. 8, p. 083024, 2020.
- [169] B. Beheshtipour and M. A. Papa, “Deep learning for clustering of continuous gravitational wave candidates,” *Phys. Rev. D*, vol. 101, no. 6, p. 064009, 2020.
- [170] B. Beheshtipour and M. A. Papa, “Deep learning for clustering of continuous gravitational wave candidates II: identification of low-SNR candidates,” *Phys. Rev. D*, vol. 103, no. 6, p. 064027, 2021.
- [171] P. B. Covas *et al.*, “Identification and mitigation of narrow spectral artifacts that degrade searches for persistent gravitational waves in the first two observing runs of Advanced LIGO,” *Phys. Rev. D*, vol. 97, no. 8, p. 082002, 2018.
- [172] R. Tenorio, L. M. Modafferi, D. Keitel, and A. M. Sintes, “Empirically estimating the distribution of the loudest candidate from a gravitational-wave search,”
- [173] A. Singh, M. A. Papa, H.-B. Eggenstein, and S. Walsh, “Adaptive clustering procedure for continuous gravitational wave searches,” *Phys. Rev. D*, vol. 96, no. 8, p. 082003, 2017.

- [174] P. Leaci and R. Prix, “Directed searches for continuous gravitational waves from binary systems: parameter-space metrics and optimal Scorpius X-1 sensitivity,” *Phys. Rev. D*, vol. 91, no. 10, p. 102003, 2015.
- [175] B. P. Abbott *et al.*, “Results of the deepest all-sky survey for continuous gravitational waves on LIGO S6 data running on the Einstein@Home volunteer distributed computing project,” *Phys. Rev. D*, vol. 94, no. 10, p. 102002, 2016.
- [176] B. Allen, “ $\chi^2$  time-frequency discriminator for gravitational wave detection,” *Phys. Rev. D*, vol. 71, no. 6, p. 062001, 2005.
- [177] L. Sancho de la Jordana and A. M. Sintes, “A  $\chi^2$  veto for continuous gravitational wave searches,” *Class. Quant. Grav.*, vol. 25, no. 18, p. 184014, 2008.
- [178] D. Keitel, *Improving robustness of continuous-gravitational-wave searches against signal-like instrumental artefacts and a concept for an octahedral gravitational-wave detector in space*. PhD thesis, Leibniz U., Hannover, 2014.
- [179] D. Davis *et al.*, “LIGO detector characterization in the second and third observing runs,” *Class. Quant. Grav.*, vol. 38, no. 13, p. 135014, 2021.
- [180] “GWOSC - O1 Instrumental Lines.” <https://www.gw-openscience.org/o1speclines/>, 2021.
- [181] “GWOSC - O2 Instrumental Lines.” <https://www.gw-openscience.org/o2speclines/>, 2021.
- [182] “GWOSC - O3a Instrumental Lines.” <https://www.gw-openscience.org/O3/o3aspeclines/>, 2021.
- [183] E. Goetz, A. Neunzert, K. Riles, A. Mateas, S. Kandhasamy, J. Tasson, *et al.*, “O3a lines and combs in C00 data.” <https://dcc.ligo.org/T2000719/public>, 2020.
- [184] E. Goetz, A. Neunzert, K. Riles, A. Mateas, S. Kandhasamy, J. Tasson, *et al.*, “Unidentified O3a lines and combs in C00 data.” <https://dcc.ligo.org/T2000720/public>, 2020.
- [185] E. Goetz, A. Neunzert, K. Riles, A. Mateas, S. Kandhasamy, J. Tasson, *et al.*, “O3 lines and combs found in self-gated C01 data.” <https://dcc.ligo.org/T2100200/public>, 2021.
- [186] E. Goetz, A. Neunzert, K. Riles, A. Mateas, S. Kandhasamy, J. Tasson, *et al.*, “Unidentified O3 lines and combs found in self-gated C01 data.” <https://dcc.ligo.org/T2100200/public>, 2021.
- [187] J. C. Driggers *et al.*, “Improving astrophysical parameter estimation via offline noise subtraction for Advanced LIGO,” *Phys. Rev. D*, vol. 99, no. 4, p. 042001, 2019.
- [188] D. Davis, T. J. Massinger, A. P. Lundgren, J. C. Driggers, A. L. Urban, and L. K. Nuttall, “Improving the Sensitivity of Advanced LIGO Using Noise Subtraction,” *Class. Quant. Grav.*, vol. 36, no. 5, p. 055011, 2019.
- [189] S. J. Zhu, M. A. Papa, and S. Walsh, “New veto for continuous gravitational wave searches,” *Phys. Rev. D*, vol. 96, no. 12, p. 124007, 2017.
- [190] D. Keitel, G. Woan, M. Pitkin, C. Schumacher, B. Pearlstone, K. Riles, A. G. Lyne, J. Palfreyman, B. Stappers, and P. Weltevrede, “First search for long-duration transient gravitational waves after glitches in the Vela and Crab pulsars,” *Phys. Rev. D*, vol. 100, no. 6, p. 064058, 2019.
- [191] J. Zweizig and K. Riles, “Information on self-gating of  $h(t)$  used in O3 continuous-wave and stochastic searches.” <https://dcc.ligo.org/LIGO-T2000384/public>, 2021.
- [192] B. Steltner, M. A. Papa, and H.-B. Eggenstein, “Identification and removal of non-Gaussian noise transients for gravitational wave searches.”
- [193] M. Isi, S. Mastroianni, M. Pitkin, and O. J. Piccinni, “Establishing the significance of continuous gravitational-wave detections from known pulsars,” *Phys. Rev. D*, vol. 102, no. 12, p. 123027, 2020.

- [194] K. Wette, “Estimating the sensitivity of wide-parameter-space searches for gravitational-wave pulsars,” *Phys. Rev. D*, vol. 85, p. 042003, 2012.
- [195] C. Dreissigacker, R. Prix, and K. Wette, “Fast and Accurate Sensitivity Estimation for Continuous-Gravitational-Wave Searches,” *Phys. Rev. D*, vol. 98, no. 8, p. 084058, 2018.
- [196] Prix, Reinhard, “Coherent  $\mathcal{F}$ -statistic on semi-coherent candidate.” <https://dcc.ligo.org/LIGO-T1700236/public>, 2019.
- [197] G. Ashton, R. Prix, and D. I. Jones, “Statistical characterization of pulsar glitches and their potential impact on searches for continuous gravitational waves,” *Phys. Rev. D*, vol. 96, no. 6, p. 063004, 2017.
- [198] G. Ashton, R. Prix, and D. I. Jones, “A semicoherent glitch-robust continuous-gravitational-wave search method,” *Phys. Rev. D*, vol. 98, no. 6, p. 063011, 2018.
- [199] M. Shaltev and R. Prix, “Fully coherent follow-up of continuous gravitational-wave candidates,” *Phys. Rev. D*, vol. 87, no. 8, p. 084057, 2013.
- [200] M. Shaltev, P. Leaci, M. A. Papa, and R. Prix, “Fully coherent follow-up of continuous gravitational-wave candidates: an application to Einstein@Home results,” *Phys. Rev. D*, vol. 89, no. 12, p. 124030, 2014.
- [201] D. Keitel, R. Tenorio, G. Ashton, and R. Prix, “PyFstat: a Python package for continuous gravitational-wave data analysis,” *J. Open Source Softw.*, vol. 6, no. 60, p. 3000, 2021.
- [202] J. A. Nelder and R. Mead, “A simplex method for function minimization,” *Comput. J.*, vol. 7, pp. 308–313, 1965.
- [203] C. Rover, C. Messenger, and R. Prix, “Bayesian versus frequentist upper limits,” in *PHYSTAT 2011*, (Geneva), pp. 158–163, CERN, 2011.
- [204] V. Dergachev, “Novel universal statistic for computing upper limits in an ill-behaved background,” *Phys. Rev. D*, vol. 87, no. 6, p. 062001, 2013.



## Chapter 3

# Time-frequency track distance for comparing continuous gravitational wave signals

This chapter is an adaptation of the material presented in

*Time-frequency track distance for comparing continuous gravitational wave signals*

Rodrigo Tenorio, David Keitel, Alicia M. Sintes

[Phys. Rev. D 103, 064053 \(2021\)](#) – [arXiv:2012.05752 \[gr-qc\]](#)

DOI: [10.1103/PhysRevD.103.064053](#)

### 3.1 Introduction

Continuous gravitational waves (CWs) are long-lasting, quasi-monochromatic, gravitational wave (GW) signals emitted by sustained quadrupolar variations of a mass distribution. As opposed to GWs from compact binary coalescences [1], CWs require long periods of data to be integrated in order to accumulate a significant signal-to-noise ratio (SNR), as expected signals are very faint. Focusing on current ground-based detectors like Advanced LIGO [2] and Advanced Virgo [3], possible sources include rapidly-spinning non-axisymmetric neutron stars, which could emit GW radiation through a variety of mechanisms (see [4] for a recent review), as well as more exotic alternatives such as ultralight boson clouds around spinning black holes (see [5] and references therein).

Several methods have been proposed to search for CWs, depending on the available information about the targeted sources. On one end, targeted searches aim at known pulsars, assuming a tight phase locking between electromagnetic and gravitational radiation; on the other end, all-sky searches try to uncover unknown sources, imposing fewer constraints on possible signals.

Narrow parameter-space regions, as considered by targeted and other well-constrained searches, can be effectively analyzed using coherent methods which compare the full data stream to a bank of phase-evolution templates. The number of templates to be considered in a coherent search scales with a strong power of the observing time [6]. As a result, these methods can not be directly applied to wider parameter-space surveys, such as all-sky searches.

Semicoherent methods (e.g. [7, 8, 9, 10, 11]), on the other hand, divide the data stream into a discrete set of segments, each delivering a coherent statistic taking the full phase evolution into account. The final detection statistic is then computed incoherently from these coherent statistics by following the frequency-evolution track associated to each parameter-space candidate across segments. Since phase and frequency evolution are related by time differentiation, a single frequency-evolution template is related to an ensemble of phase-evolution templates, effectively reducing the number of required templates to cover a parameter space region (see e.g. [12, 13]). Also, this renders semicoherent methods insensitive to discrete phase jumps between coherent segments. Consequently, these types of searches are susceptible to deliver more spurious candidates than coherent searches, which then require further investigation [14].

Typical semicoherent searches follow a hierarchical scheme, applying increasingly more sensitive (and more computationally expensive) stages as interesting parameter-space regions are narrowed down [15]. Both the setup of such hierarchical searches and the post processing of any outlier candidates depend on an understanding of the parameter-space structure. For example, many searches [16, 17, 18, 19, 20, 21]

use clustering algorithms to reduce the number of correlated candidates that proceed from one stage to the next. Most clustering algorithms depend on a meaningful measure of parameter-space distance to represent said correlations within the algorithm.

Much of the understanding about CW parameter-space correlations has been built upon the  $\mathcal{F}$ -statistic, a maximum likelihood estimator accounting for the SNR in a coherent CW search [22, 23]. For a signal present in some data, the best-matching template will return a high  $\mathcal{F}$ -statistic value. Neighboring templates, however, will also show enhanced (albeit smaller)  $\mathcal{F}$ -statistic values. The relative loss due to mismatched templates was used to propose a family of local parameter-space metrics [6]. These metrics can be used to set up optimal parameter-space grids, placing templates in such a way that any signal within the considered parameter space is not mismatched by more than a certain fixed amount [24, 25, 26].

Optimal template-placement methods have been successfully generalized to semicoherent methods [12, 27, 28], improving the overall effectiveness of hierarchical schemes. However, there is currently still a lack of literature on the definition of distance measures for the post-processing of candidates from semicoherent search stages.

In this paper we propose a new parameter-space distance suited for the post-processing of semicoherent CW searches by comparing the average mismatch along the frequency evolution tracks associated to a pair of parameter-space points. This novel approach can be applied to any kind of quasi-monochromatic CW signal as long as a model of its frequency evolution is available. We provide a theoretical justification for this proposal by relating it to the  $\mathcal{F}$ -statistic. As a demonstration of its usefulness, we apply it to a clustering algorithm, comparing its effectiveness with respect to the results obtained in a previous semicoherent search.

In Sec. 3.2 we review how parameter-space structure is taken into account in current post-processing approaches. We introduce and discuss the properties of our proposal in Sec. 3.3. In Sec. 3.4 we discuss a clustering algorithm as a first application of the new distance and give a numerical assessment of its effectiveness in Sec. 3.5 by improving the sensitivity of an all-sky search using open data from the second observing run of the Advanced LIGO detectors. We summarize our proposal and conclusions in Sec. 3.6.

## 3.2 Parameter-space distances and clustering algorithms

The main application of parameter-space distances in the context of semicoherent searches is to reduce the number of outliers by grouping them into clusters, thus reducing the computational cost and human effort of additional follow-up stages. This can also shift the focus from individual candidates to correlated parameter-space regions, and help associate these to specific origins such as instrumental noise or astrophysical signals.

To understand the potential benefits of our new proposed distance measure, here we briefly review the ways in which parameter-space structure is taken into account by clustering algorithms currently in practical use for CW all-sky searches. Those can be divided into two families according to the way in which they take parameter-space information into account, either based on fixed grids or on machine learning.

### 3.2.1 Grid-based clustering methods

The first family of methods relies on the definition of a distance in terms of a parameter-space grid. This approach measures *how many grid steps away* two given parameter-space points are from each other, making it well suited to estimate the computing costs of further follow-up stages.

Examples of this approach are the all-sky searches for CWs from isolated neutron stars in LIGO O1 and O2 data [17, 18, 19]. The signals covered by these searches can be described by four parameters, namely initial frequency, spin-down and sky position:  $\lambda = \{f_0, f_1, \vec{n}\}$ .

The Euclidean parameter-space distance is defined as

$$d(\lambda, \lambda_*) = \sqrt{\left(\frac{\Delta f_0}{\delta f_0}\right)^2 + \left(\frac{\Delta f_1}{\delta f_1}\right)^2 + \left(\frac{\Delta \theta}{\delta \theta}\right)^2}, \quad (3.1)$$

where  $\Delta$  quantities represent one-dimensional differences between parameter-space points and  $\delta$  quantities represent grid resolutions, which are constructed according to the criteria of each search pipeline.

While the frequency and spin-down contributions are treated rather consistently, different pipelines have opted for different prescriptions for the sky contribution  $\theta$ . A common characteristic amongst O1 searches was the choice of the *great circle distance* to account for the sky contribution in Eq. (3.1).

As discussed in [7], the structure of the sky subspace is highly anisotropic, as CW signals located close to the ecliptic plane tend to be more delocalized than those coming from the ecliptic poles. This is due to the Doppler modulation induced by the Earth’s motion, the dominant component of which is given by the yearly motion around the Solar System Barycenter (SSB). Thus, an alternative ansatz was proposed in [29] by projecting sky positions onto the ecliptic plane

$$\Delta\theta^2 = \Delta x^2 + \Delta y^2, \quad (3.2)$$

where  $\Delta x, \Delta y$  represent the difference in Cartesian ecliptic coordinates. A convenient anisotropic sky grid is implicitly set up after neglecting contributions from the third Cartesian component  $\Delta z$ : The sky grid around the poles remains untouched, while the region around the ecliptic plane becomes *wider* along the polar direction, in qualitative agreement with the described sky structure. This idea was applied to the SkyHoughpipeline during the LIGO O2 all-sky search [19]. However the great circle distance was still used to estimate local sky distances for the setup of follow-up stages [19, 20].

Extensions to Eq. (3.1) to account for the additional modulations experienced by signals from sources in circular binary orbits were proposed in [20], adapting this type of methods to the search for CWs from neutron stars in binary systems. The proposal considered the three parameters required to describe a circular binary orbit, namely the projected semi-major axis  $a_p$ , orbital frequency  $\Omega$ , and time of passage through the ascending node  $t_{\text{asc}}$ , adding them in quadrature as an extra term to Eq. (3.1):

$$d_{\text{Binary}}(\lambda, \lambda_*) = \sqrt{\left(\frac{\Delta a_p}{\delta a_p}\right)^2 + \left(\frac{\Delta \Omega}{\delta \Omega}\right)^2 + \left(\frac{\Delta t_{\text{asc}}}{\delta t_{\text{asc}}}\right)^2}. \quad (3.3)$$

Current clustering applications to CW searches then typically group candidates located within a certain distance  $d^{\text{th}}$  of each other, with  $d^{\text{th}}$  calibrated to maximize the detection efficiency for a population of artificial signals injected into the data. For the searches referenced above, this was chosen as  $d^{\text{th}} \in [3, 4]$  parameter-space bins.

The effectiveness of these clustering methods has been demonstrated by application in the cited searches. It is unclear, however, how to extend them towards a more general description of the parameter space. So far, distance proposals have assumed parameters to be uncorrelated. But this is generally not true, with correlations highly dependent on the parameter space under analysis [12, 13, 27]. Also regarding the actual functional form in which each parameter-space dimension enters into the equation, existing grid-based proposals turned out to be effective for “simple” cases, such as isolated sources (e.g. Eq. (3.2)). However, more complex types of signals, such as those modulated by binary orbits, pose a challenge to this approach.

### 3.2.2 Machine Learning-based clustering

The second family of clustering methods uses machine learning (ML) techniques to classify parameter-space regions and identify those potentially containing CW signals [30, 31, 32]. The input is usually a graphical representation of the search outputs over the parameter space, e.g. the detection statistic values obtained during the main stage of the search. Algorithms then learn to classify these through the usual train-test-validate ML workflow (see e.g. [33]).

This approach presents several advantages with respect to methods from Sec. 3.2.1, as ML classifiers are able to learn parameter-space correlations through the provision of representative realizations of parameter-space regions during the training stage of the algorithm. These training examples are relatively easy to produce, making it computationally feasible to train ML classifiers using a rich set of correlation examples. In this sense, the parameter-space structure is no longer described by an explicit distance function. Still, one has to carefully consider the way in which search results are represented in order to properly train the algorithm and hence to fully exploit the potential of ML techniques.

### 3.3 A new parameter space distance

We introduce a new parameter space distance with the intent of overcoming the limitations identified during the previous section. We employ the usual description of CWs in terms of the instantaneous frequency evolution of a signal, which is in turn described in terms of the physical parameters of the source. By considering frequency evolution tracks rather than a particular choice of parameters, we are able to describe a general distance that relates directly to the same information content exploited by semicoherent CW searches.

#### 3.3.1 Deriving a distance from the $\mathcal{F}$ -statistic

CWs can be characterized by two sets of parameters, namely the phase parameters  $\lambda$  and the amplitude parameters  $\mathcal{A}$ . The former contain information about the phase and frequency evolution of a signal, which is due to both physical mechanisms intrinsic to the source (e.g. binary orbital motions) and relative motions of a GW detector with respect to the SSB; the latter describe the amplitude of a signal, encoding information about the GW polarization, as well as the antenna response of the detector.

The detector response to a CW can be expressed as a linear superposition of four filters, obtained by projecting each GW polarization onto both detector response functions:

$$h_{\mathcal{A},\lambda}(t) = \sum_{\mu=0}^3 \mathcal{A}^\mu h_\mu(t; \lambda) \quad (3.4)$$

where each component  $\mathcal{A}^\mu$  is a time-independent expression depending only on the amplitude parameters and  $h_\mu$  depends on the phase parameters only [22, 34].

Let us now assume a time series  $x(t)$  consisting of additive Gaussian noise and a CW signal  $h_{\mathcal{A},\lambda}(t)$ . Then, following the Neyman-Pearson criterion [22, 35], the optimal detection statistic would be given by the likelihood ratio

$$\log \Lambda(\mathcal{A}, \lambda) \propto \langle x, h_{\mathcal{A},\lambda} \rangle, \quad (3.5)$$

where  $\langle \cdot, \cdot \rangle$  represents a noise-weighted scalar product comparing data to CW templates (high values corresponding to good alignment between data stream and the proposed template). Since amplitude parameters enter as linear coefficients in Eq. (3.4), analytic maximization can be performed in Eq. (3.5), obtaining the so-called  $\mathcal{F}$ -statistic [22, 23]

$$\mathcal{F}(\lambda) \equiv \max_{\mathcal{A}} \log \Lambda(\mathcal{A}, \lambda). \quad (3.6)$$

However, as first pointed out in [36], this statistic presents a certain degeneracy with respect to phase parameters; in other words, the locus of phase parameters for which Eq. (3.6) attains (close to) a maximum value is not simply a point, but a finite parameter-space region.

A study of this property was done in [37] using a simplified version of Eq. (3.4) by neglecting amplitude modulations and solely focusing on the phase parameter dependency

$$h_{\mathcal{A},\lambda}^*(t) = A \sin \Phi_\lambda(t) + B \cos \Phi_\lambda(t), \quad (3.7)$$

where  $\Phi_\lambda$  is the phase evolution associated to  $\lambda$ . The validity of this model and its agreement with Eq. (3.6) was justified in [38]; in turn, it implies a corresponding  $\mathcal{F}^*$  statistic by maximizing with respect to the (now constant) amplitude parameters:

$$\mathcal{F}^*(\lambda) \equiv \max_{A,B} \log \Lambda^* \propto |\mathcal{X}(\lambda)|^2, \quad (3.8)$$

where the following notation for the simplified filter was introduced:

$$\mathcal{X}(\lambda) = \langle x, e^{-i\Phi_\lambda} \rangle. \quad (3.9)$$

Let us now assume a CW signal described by a set of parameters  $\lambda_*$  is present in our data. If we neglect rapidly oscillating terms in Eq. (3.9), we can express the  $\mathcal{F}^*$  statistic in terms of the phase



difference between the template parameters  $\lambda$  and those of the signal,  $\Delta\Phi(t) = \Phi_{\lambda_*}(t) - \Phi_\lambda(t)$ , namely

$$\mathcal{X}(\lambda) \propto \int_0^T dt e^{i\Delta\Phi(t)}. \quad (3.10)$$

As discussed in [37], regions where  $\mathcal{F}^*$  attains a maximum value will approximately correspond to those for  $\mathcal{F}$ . It is clear from Eq. (3.8) that a maximum value of  $\mathcal{F}^*$  corresponds to a maximum value of  $|\mathcal{X}|$ , which is achieved whenever

$$\Delta\Phi(t) = 0 \quad \forall t. \quad (3.11)$$

We can re-state this condition in terms of the instantaneous frequency of said pair of templates  $f_\lambda(t)$

$$\frac{\partial\Delta\Phi}{\partial t} = 0 \rightarrow f_{\lambda_*}(t) - f_\lambda(t) = 0 \quad \forall t, \quad (3.12)$$

which is to say the statistic will be maximal for those pairs of templates whose frequency evolution coincides over the observing time. We also refer to the frequency evolution of a candidate or template as its *time-frequency track*, or simply *track*.

With Eq. (3.12) we start our new proposal: the information a search recovers is best described not just by the parameter-space point corresponding to a candidate, but by its actual frequency evolution. As a result, we can relate different parameter-space points by comparing their frequency evolution tracks using an arbitrary functional distance

$$d(\lambda_*, \lambda) \equiv \mathfrak{D} [f_{\lambda_*} - f_\lambda]. \quad (3.13)$$

This means the distance amongst parameter-space candidates is measured in a consistent way with semicoherent searches, where the statistical significance is defined as an integral along a time-frequency track. In the same sense a frequency-evolution mismatch would produce a loss in significance, it accounts for the amount of distance from one candidate to another.

For the remainder of this work we choose to focus on the  $L_1$  distance in order to reduce the amount of implicit weights involved in our analysis, noting that any other proposal from the  $L_p$  family should yield similar results.

Before concluding this exposition, it is worth noticing there were no assumptions on the actual form of  $f_\lambda$  except for the fact it had to be a CW signal (i.e. long-lasting and quasi-monochromatic). This means that Eq. (3.13) could be used for any kind of quasi-monochromatic CW signal regardless of its parametrization as long as we are able to describe the frequency evolution of the source in terms of a set of parameters  $\lambda \in \mathbb{P}$ .

### 3.3.2 Distance implementation

We proceed to explain how to adapt this newly introduced idea to the particularities of practical CW searches, reserving a discussion on the actual effect of this distance as well as a comparison to previous approaches for sections 3.4 and 3.5.

As already discussed, a first proposal of parameter-space distance can be constructed by comparing the time-frequency tracks of a pair of candidates through the  $L_1$  distance

$$d(\lambda, \lambda_*) \propto \int dt |f_\lambda(t) - f_{\lambda_*}(t)|. \quad (3.14)$$

While useful on its own, it can be improved by taking the following considerations into account.

First, the data products which all-sky CW searches deal with are discretized in a specific way. For instance, most all-sky semicoherent searches work on Fourier-transformed data segments (so-called *Short Fourier Transforms*, SFTs) with a time baseline  $T_{\text{SFT}}$ , which in turn imposes a natural frequency discretization  $\delta f = T_{\text{SFT}}^{-1}$ . Each  $T_{\text{SFT}}$ -long segment  $\alpha$  is labeled by a starting *timestamp*  $t_\alpha$ , meaning we can get a discretized version of Eq. (3.14) by introducing the substitution

$$\int dt \rightarrow \frac{1}{\delta f} \sum_\alpha, \quad (3.15)$$

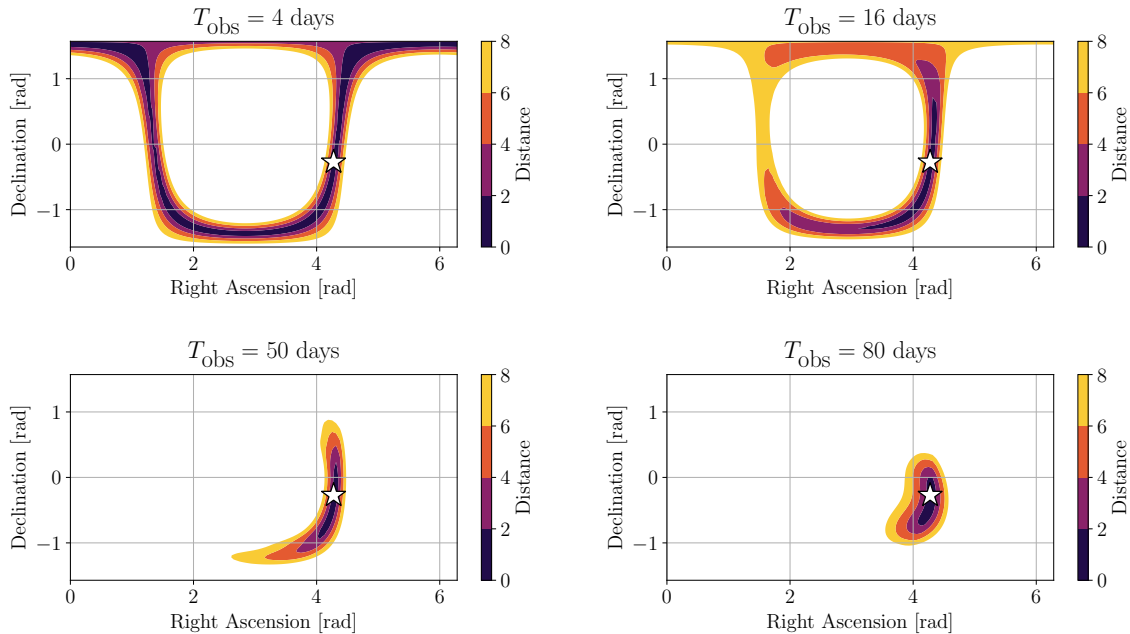


FIGURE 3.1: Evaluation of the distance from Eq. (3.16) across the sky with respect to a reference sky position  $(\alpha, \delta) = (4.27, -0.27)$  rad for different observing times.

where the multiplying factors are kept in such a way that the distance stays dimensionless, i.e. the distance is measured in the corresponding (fractional) number of frequency bins  $\delta f$ .

Second, consider two templates  $\lambda, \lambda_* \in \mathbb{P}$  such that their time-frequency tracks are, on average,  $\tilde{\Delta}$  bins apart. If we assume a set of  $N$  SFTs, then from Eqs. (3.14) and (3.15) the distance is given by

$$d(\lambda, \lambda_*) \propto \sum_{\alpha=0}^{N-1} |f_{\lambda}(t_{\alpha}) - f_{\lambda_*}(t_{\alpha})| \simeq \sum_{\alpha=0}^{N-1} \tilde{\Delta} = N\tilde{\Delta},$$

i.e. longer observing times would be able to build up higher distance values. We can factor this dependency out by choosing a proper normalization, effectively constructing a parameter-space distance whose properties will be maintained across different observing runs:

$$d(\lambda, \lambda_*) \equiv \frac{1}{N\delta f} \sum_{\alpha} |f_{\lambda}(t_{\alpha}) - f_{\lambda_*}(t_{\alpha})|. \quad (3.16)$$

This last expression can be interpreted as the *average mismatch* among the time-frequency tracks associated to a pair of templates.

Also, in the same way as search methods are able to produce sound results for data streams containing gaps due to down time of the detectors, Eq. (3.16) can be computed using a sub-set of timestamps  $\bar{\alpha}$ , reducing the amount of required computations

$$d(\lambda, \lambda_*) \equiv T_{\text{SFT}} \cdot \langle |f_{\lambda} - f_{\lambda_*}| \rangle, \quad (3.17)$$

where  $\langle \cdot \rangle$  represents an average over a certain set of timestamps  $\alpha \in \bar{\alpha}$ .

The following subsection introduces the main properties of Eq. (3.16) and compares them to those of Eq. (3.17). After showing their quantitative equivalence, we will simply refer to Eq. (3.17) as the *Time-frequency Track Distance* (TTD).

### 3.3.3 Discussion

To illustrate the properties of our new distance proposal from Eq. (3.16), we first consider the simplest case of a CW signal, consisting of a source emitting GWs at a certain fixed frequency  $f_0$ . Neglecting

any proper motion of the source, the detected signal only experiences the Doppler modulation from the detector's movement with respect to the solar system barycenter [7]:

$$f_\lambda(t) = f_0 \cdot \left( 1 + \frac{\vec{v}(t) \cdot \vec{n}}{c} \right), \quad (3.18)$$

where  $v/c \simeq 10^{-4}$  is the detector velocity expressed in natural units and  $\vec{n}$  represents the sky position of the source. Focusing on sky position effects will deliver useful results for a general discussion of the implications of our distance, since this is the only common contribution to every kind of CW search (as long as considering ground-based detectors).

We will make use of open data from the O2 observing run of the Advanced LIGO detectors [39, 40] to extract timestamps and detector velocity vectors using [41]. Both LIGO detectors are taken into account by extending Eq. (3.16) to a multi-detector setup: track mismatches are computed for each detector, computing an overall average over the timestamps of all used SFTs. This approach is possible due to the usage of track mismatch rather than SNR loss to compare parameter-space points, as the former is unaffected by amplitude modulations due to the different antenna pattern functions of each detector.

The parameter-space structure of this model is highly dependent on the length of the data stream under study. As first pointed out in [7], if we measured the frequency of a CW signal only at a certain moment in time, we would not be able to fully locate its corresponding sky position, since the Earth-induced Doppler shift would point us to a *circle* on the celestial sphere. It is through the integration over multiple timestamps that the sky position of the source becomes well resolved.

This type of structure is well captured by Eq. (3.16), as shown in Fig. 3.1: For short observing times, not being able to resolve sky position corresponds to the locus of close parameter-space points (as measured by the distance) being broadly extended along a circle on the celestial sphere. As observing time increases, this degeneracy starts to break down, with the set of close templates finally narrowing down to a compact neighborhood around the actual sky position of the source. This example demonstrates how our proposal is able to automatically capture an underlying parameter-space structure. We note that the improved sky localization with increasing observing times is a generic feature of CW searches, as it is the result of identifying sky regions consistent with the long-term Doppler modulation induced by the Earth's movement with respect to the SSB [7]. Longer observing times also increase search sensitivity by increasing the amount of SNR that can be accumulated by candidates.

The behavior of Eq. (3.16) should be compared to those used in previous searches, as discussed in section 3.2. In particular, we compare to the great circle distance to provide us with an estimation of the typical numerical scale of the new distance. We express the great circle distance according to the natural sky resolution induced by the Doppler modulation (so called *sky bins*) [7]:

$$\delta\theta = \left[ \frac{v}{c} \cdot T_{\text{SFT}} \cdot f_0 \right]^{-1}. \quad (3.19)$$

Both distances are compared in Fig. 3.2, using the same reference sky position as in Fig. 3.1 (the actual choice of sky position plays a minor role, as we are focusing on close neighborhoods in terms of the great circle distance). Given a certain distance value, higher frequencies tend to correspond to longer circle lengths, as the bin resolution from Eq. (3.19) becomes thinner.

We note a distance of  $\mathcal{O}(1)$  is enough to cover arc lengths corresponding to less than 4 parameter-space bins. This will allow us to compare the performance of this distance to that of the Euclidean ansatz in the context of a clustering algorithm in section 3.5.

Next we test the validity of the reduced Eq. (3.17) with respect to Eq. (3.16) by computing the absolute deviation produced by computing the distance using a subset of the available timestamps. Concretely, we select 500 timestamps along the O2 observing run data out of the initial  $\mathcal{O}(10^4)$ . In Fig. 3.3 we show the distribution of absolute deviations between Eq. (3.17) and Eq. (3.16), using CW signals with binary modulations as in [20]. Absolute deviations are two orders of magnitude smaller than the exact distance value within the useful regime; hence we can safely use this approach to save on computing cost. This is consistent with the expected behavior of time-frequency tracks: sufficiently close parameter-space points would produce similar time-frequency evolutions, meaning their mismatch during a particular time span would be similar to the average value across the whole run; hence, only a small fraction of the timestamps are required to be taken into account. This discussion is easily




FIGURE 3.2: Distance comparison with respect to the great circle distance using  $10^5$  neighbouring sky positions, at two representative signal frequencies. Markers represent the average distance value for sky positions located at the specified number of bins  $\pm 0.25$ , and envelopes represent the maximum and minimum values within each bracket. The vertical dashed line marks a value of  $\sqrt{14}$  bins, the reference value employed by [20] to cluster parameter-space points.

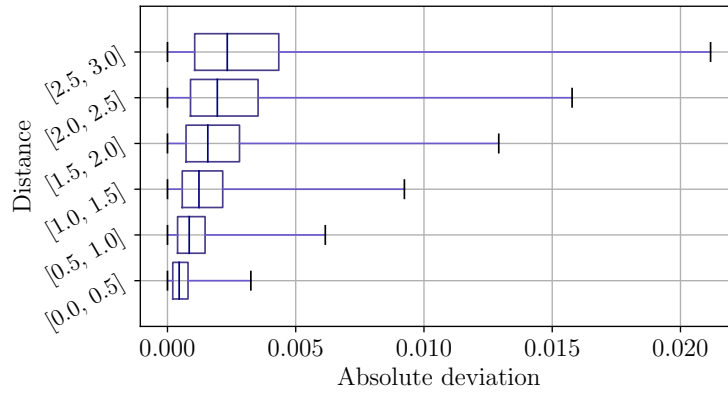


FIGURE 3.3: Distance comparison between Eqs. (3.16) and (3.17) using  $3 \times 10^5$  pairs of parameter-space points corresponding to the binary parameter space of [20]. Vertical lines show the median of the distribution of values bracketed by the specified interval; boxes mark the first and third quartiles of the distribution; whiskers denote the minimum and maximum values within each of the brackets.

extendible to different types of sources, and can be taken at face value if relative motion effects are smaller than Earth-induced Doppler effects (which correspond to a relative frequency variation of  $10^{-4}$ ).

### 3.4 Application to clustering for an all-sky binary CW search

We now apply TTD to improve the clustering of candidates from CW searches. This first example application will prove the benefits of using track-based parameter-space distances in CW searches. Other possible uses of this proposal, such as the setup of parameter-space grids for the initial search stage, or for the follow-up of candidates, are left for future work.

For the rest of this work, we will focus on the semicoherent `BinarySkyHough` pipeline [19, 20], which is able to perform all-sky searches for CWs from unknown neutron stars in binary systems. The pipeline selects a narrow frequency band (typically on the order of 0.1 Hz) and a certain binary parameter-space region, performing an all-sky search and returning the most significant candidates according to a set of statistics.

We use the new distance in a clustering algorithm that is functionally equivalent to the one in [19, 20], which originally used the distance from Eq. (3.2); from now on we refer to that reference algorithm as “the O2 clustering”. We comment on minor implementation differences in subsection 3.4.1, ensuring the fairness of the comparison; a numerical assessment is postponed to Sec. 3.5.

The general starting point for clustering is a toplist  $\mathcal{T}$  containing the most significant candidates of a certain parameter-space region according to a search algorithm. We define a candidate as the pair  $c^{(i)} = (\lambda^{(i)}, \mathfrak{S}^{(i)})$ , where  $\lambda^{(i)} \in \mathbb{P}$  represents the parameters associated to said candidate and  $\mathfrak{S}^{(i)} \in \mathbb{R}$  is the significance, represented by some statistic.

#### 3.4.1 Algorithm overview

The clustering algorithm consists of these steps:

1. *Construct reachable sets*: select candidates located in neighboring frequency bins to avoid having to compute irrelevant comparisons later on.
2. *Construct coincidental sets*: compute TTD from Eq. (3.17) among the selected set of candidates and pair them according to a coincidence threshold  $d^{\text{th}}$ .<sup>1</sup>
3. *Construct clusters*: group coincidental pairs into connected components. Each connected component represents a *cluster*.
4. *Identify cluster centers*: retrieve the loudest candidate from each cluster according to a ranking statistic  $\mathfrak{S}$ .
5. *Select clusters to follow up*: rank the clusters according to the statistic of their centers, selecting a certain number  $N_c$  to follow up using a more sensitive method.

The first conceptual difference with respect to the O2 clustering is in step 1: While we focus on the neighboring frequency bins of a candidate, the O2 clustering chains together neighboring candidates until a frequency gap wider than a certain threshold is encountered. The main effect of such a strategy is to create wider reachable sets than strictly required, meaning the expected results are consistent with our proposal.

The second difference is related to the ranking criterion in step 5, as the O2 clustering ranks clusters according to the *sum* of the individual statistics of their elements. In the presence of strong instrumental disturbances, individual candidates in the cluster are strong enough to be selected regardless of the exact criterion; in the case of signals (expected to be weak), however, there is an explicit dependence on the number of toplist elements used, as selecting more candidates could lead to an inflation of the cluster population due to the presence of candidates produced by pure noise. Selecting the *loudest* candidate is safer in this regard, as the number of toplist elements plays no role. As discussed in [11], the effect

<sup>1</sup>Unlike in many other contexts of gravitational-wave data analysis, by “coincidence” here we are not referring to a comparison between data from different detectors. Instead, the point is whether two candidates coincide in their time-frequency evolution sufficiently closely.

of this difference will be negligible in our comparison as we will be focusing on estimating a sensitivity corresponding to the 95% detection efficiency level.

Finally, the O2 clustering adds a minimum population to the selection criteria in step 5: clusters must contain at least 3 elements, being discarded otherwise. We comment on the effects of this criterion in Sec. 3.5.

The following subsections cover implementation details of each step of the algorithm.

### 3.4.2 Reachable sets

We start by defining an auxiliary distance based only on the frequency difference (measured in bins) between two candidates

$$d_{f_0}(c, c_*) = T_{\text{SFT}} \cdot |f_0^c - f_0^{c_*}|, \quad (3.20)$$

where  $f_0^c$  represents the frequency associated to candidate  $c$ . For each candidate  $c \in \mathcal{T}$  we define its *reachable set* as those candidates closer than a certain threshold  $d_{f_0}^{\text{th}}$  according to Eq. (3.20):

$$R[c] = \{c_* \in \mathcal{T} \mid f_0^c \leq f_0^{c_*} \text{ and } d_{f_0}(c, c_*) \leq d_{f_0}^{\text{th}}\}. \quad (3.21)$$

This construction is motivated by the expected way a CW signal creates candidates across the parameter space, as the usual conventions taken by CW searches (and by `BinarySkyHough` [11] in particular) aim to reduce its wandering to less than one frequency bin over a coherent segment; hence, we can focus on a neighborhood of each candidate rather than the whole toplist. As a result, in order to group candidates it is enough to focus on reachable sets. Note that the reachable set only needs to include candidates with a higher frequency than the selected one, as lower frequencies are already considered by previous candidates due to the symmetry of the distance function.

### 3.4.3 Coincidental sets

Pair-wise distances using TTD are computed within each reachable set. Then, we impose a *coincidence threshold*  $d^{\text{th}}$  to define the *coincidental set* of a candidate  $c \in \mathcal{T}$  as those candidates within the reachable set whose distance to the reference candidate is lower than the specified threshold:

$$C[c] = \{c_* \in R[c] \mid d(c, c_*) \leq d^{\text{th}}\}. \quad (3.22)$$

Clusters are then constructed by computing the greatest groups of candidates such that any given candidate within a group is closer than a distance of  $d^{\text{th}}$  to at least one other candidate within the same group. This concludes the clustering procedure.

Each cluster possesses two attributes, namely a *center* and a value of *significance*. The center of a cluster is defined as the *loudest* point belonging to it, i.e., the one with the highest detection statistic. Consequently, the *significance* of a cluster is defined as that of the loudest point.

The resulting clusters are then selected according to their significance, using the cluster center as the representative point to be furthered followed up using a more sensitive algorithm [42, 43, 44, 45, 19, 21, 46, 47].

## 3.5 Improving the sensitivity of all-sky searches

We assess the sensitivity improvement for an all-sky search due to applying our new distance in the previously described clustering algorithm, using a set of Monte-Carlo injections. This software injection campaign was performed using data from the Advanced LIGO O2 run so as to be comparable to the results obtained in the O2 open data all-sky search for CWs from neutron stars in binary systems [20]. Clustering algorithms are functionally equivalent in both cases, meaning sensitivity improvements can be attributed to the change of parameter-space distance. We also discuss the choice of clustering parameters in detail in order to understand these improvements.

Frequency Band [Hz]	$\langle \mathcal{D}^{95\%} \rangle \pm 1\sigma$
[100, 125)	$21.0 \pm 0.7$
[125, 150)	$19.9 \pm 0.8$
[150, 200)	$19.1 \pm 1.0$
[200, 250)	$18.0 \pm 0.8$
[250, 300)	$18.0 \pm 0.8$

TABLE 3.1: Summary of the average 95% efficiency sensitivity depths obtained by the O2 open data all-sky search [19].

### 3.5.1 The O2 data software injection campaign

We reproduced the setup employed in [20] to obtain a faithful comparison to their results. We used the full O2 Advanced LIGO run, spanning nine months of data taken by the two LIGO detectors H1 (Hanford) and L1 (Louisiana) [2, 39, 40]. The employed time segments are those with the “all” tag in [48]. The data stream was divided into segments of duration  $T_{\text{SFT}} = 900$  s which were used to compute 50% overlapping Fourier transforms (the previously introduced SFTs), as discussed in [19]; data from both detectors was analyzed together. This yields a set of 14788 SFTs from H1 and 14384 SFTs from L1.

Injections were performed at different *sensitivity depth* values [49, 50]

$$\mathcal{D} = \frac{\sqrt{S_n}}{h_0}, \quad (3.23)$$

where  $h_0$  represents the CW amplitude and  $\sqrt{S_n}$  refers to an estimation of the amplitude spectral density of the detector noise; in essence, injecting at constant depth values allows us to produce consistent injection sets with respect to the underlying noise floor, allowing for comparisons at different frequency bands.

We analyzed two different sets of injections. The first set aimed to understand the general behavior of our proposal when used together with a clustering algorithm; the second one allowed us to assess the effectiveness of our proposal by obtaining an improvement in sensitivity.

The generating procedure was the same for both sets. We first selected a set of representative frequency bands and injected a certain number of CW signals into the data at different depths, drawing their parameters from uniform distributions across the analyzed parameter space. Then, we used the `BinarySkyHoughsearch` to obtain a list of the most significant candidates to be clustered. Following the criteria specified in [19], we counted an injection as *detected* if any of the  $N_c = 3$  most significant clusters containing at least three candidates was significant enough so as to appear in the all-sky search toplist and each of its parameters were no further than five parameter-space bins from the injection point.

Table 3.1 summarizes the sensitivity obtained by the O2 all-sky `BinarySkyHoughsearch` [20]. The quoted uncertainties correspond to the overall depth variance across the quoted frequency bands; these are consistent with (if not wider than) individual uncertainties obtained through the fitting procedure, meaning the comparisons we are about to carry out do not overestimate the effectiveness of the new distance.

For the first set of injections we selected two 0.1 Hz frequency bands, namely [150.1, 294.7] Hz, and two sensitivity depth values, [17.5, 21.5]  $\text{Hz}^{-1/2}$ . For each pair we injected 100 artificial CW signals. Comparing with the 95% efficiency obtained by [20], a depth of 21.5  $\text{Hz}^{-1/2}$  corresponds to a set of *weak* injections, while 17.5  $\text{Hz}^{-1/2}$  corresponds to a set of *strong* injections. Lower frequencies tend to deliver greater 95% efficiency depth values: the required parameter space resolution becomes finer as we approach greater frequencies, as smaller Doppler modulations become resolvable by the method; but in order to keep computing cost under control, coarser resolutions are used in practice, reducing the density of the grid and, in turn, the effective parameter-space fraction under analysis. This is a well-known fact affecting CW searches [24].

For the second set we used six different frequency bands, [123.2, 129.8, 146.9, 165.2, 234.0, 262.7] Hz, selecting six different depths ranging from 14  $\text{Hz}^{-1/2}$  to 29  $\text{Hz}^{-1/2}$  (including a value of 32  $\text{Hz}^{-1/2}$  for frequencies below 200Hz). Here we injected 300 artificial signals at each sensitivity depth.




FIGURE 3.4: Performance comparison for different parameter choices, in two example frequency bands. Solid lines represent the strong set of injections; dashed lines represent the weak set of injections. Horizontal lines show the efficiency results obtained by the O2 clustering. Markers represent the efficiency results obtained by using TTD. Square markers represent the results obtained by our proposal for different coincidence thresholds. Triangular markers show the results imposing a minimum cluster population of 3. Diamond markers show the results imposing a minimum population of 5.

The clustering efficiency is given by the fraction of detected injections at each sensitivity depth. We will use the sensitivity depth corresponding to the 95% efficiency, obtained through an interpolation procedure, as a figure of merit for comparison purposes. A thorough exposition of this procedure can be found in [17].

### 3.5.2 TTD: Clustering performance

First we discuss the clustering results obtained from the first set of injections. The O2 clustering employed a distance threshold of  $d_{O2}^{\text{th}} = \sqrt{14} \simeq 3.7$  parameter-space bins, optimized through the use of Monte-Carlo injections [20]; efficiencies are reported using horizontal lines in Fig. 3.4.

A fair comparison requires us to understand how each of the clustering parameters, namely the thresholds to construct reachable and coincidental sets and the minimum population threshold, interact with our proposal.

First, let us focus on the impact of changing the reachability threshold  $d_{f_0}^{\text{th}}$ . This parameter controls the number of neighboring frequency bins taken into account to compute distances between pairs of candidates: as previously discussed, a CW candidate would tend to spread only across a small set of neighboring frequency bins, either due to spectral leakage (if the CW parameters fell close to the border of a frequency bin) or parameter-space correlations; hence  $d_{f_0}^{\text{th}}$  could be kept at low values, incidentally reducing the computing cost of the overall algorithm. Indeed, running the same setup for three reachability thresholds  $d_{f_0}^{\text{th}} = (1, 2, 3)$  [ $T_{\text{SFT}}^{-1}$ ] showed no preferred value in terms of detection efficiency. As a result, it is feasible to work with  $d_{f_0}^{\text{th}} = 1$  [ $T_{\text{SFT}}^{-1}$ ], reducing the computational cost by neglecting redundant parameter-space relations.

Second, consider the interplay between a minimum population threshold (i.e. a minimum number of candidates to define a cluster as such) and the coincidence threshold. In Fig. 3.4 we show the efficiency achieved by different combinations of these two parameters. Two main features are worth of comment:

First, the impact of a minimum population threshold is noticeable when combined with low coincidence thresholds: Significant candidates are unable to group with their neighbors and, consequently, get removed due to the lack of population in the cluster. As a result, potentially detectable signals (in the sense that the main stage of the algorithm did *not* reject them) are completely discarded. This effect is more noticeable for a set of weak injections, as they are less prone to enhance detection statistics across a wide region of parameter space; higher frequencies are also affected by this more, since the density of parameter-space templates at those regions tends to be lower than at low frequencies.

Secondly, increasing the coincidence threshold partially mitigates this effect, as stand-alone significant parameter-space points are able to group together with their less significant neighbors. Low-frequency results, however, suggest a limiting factor to this strategy, as this increase could also lead to neighboring






FIGURE 3.5: Example of 95% efficiency depth interpolation at 165.2 Hz. Points represent the detected fraction of 300 software injections after running a clustering using TTD with  $d_{f_0}^{\text{th}} = 1 [T_{\text{SFT}}^{-1}]$  and  $d^{\text{th}} = 1.0$ ; the error bar on each point represents a binomial uncertainty. No minimum population is imposed. The vertical line shows the average sensitivity depth obtained by the O2 open data search at this frequency band. Shaded regions represent  $1\sigma$  and  $2\sigma$  uncertainties. The 95% efficiency depth interpolated from a sigmoid fit is represented by a star marker.

clusters merging together; given that the algorithm focuses on the *center* of a cluster only, this means a signal could be discarded if its cluster gets merged with an instrumental disturbance. This phenomenon is highly dependent on the data under analysis, as well as the properties of the CW signal in it. In particular highly disturbed bands tend to deliver highly significant clusters corresponding to instrumental artifacts, as observed in previous all-sky searches [17, 18, 19, 51]. Nevertheless, these clusters could also be isolated enough from the rest of interesting local maxima so as to avoid any mergers.

Let us impose a minimum population of three elements per cluster akin to the O2 clustering. Fig. 3.4 shows how TTD obtains improved results with respect to the Euclidean ansatz used in the original O2 algorithm, provided we select a suitable coincidence threshold. We devote the following subsection to expose the soundness of this improvement.

### 3.5.3 TTD: Sensitivity improvement

With the second set of injections we provide an assessment of the effectiveness of our proposal in terms of search sensitivity. We recall this set consisted of five frequency bands, selecting six sensitivity depths for each of them, at which we injected 300 CW signals each. We use these injections to interpolate the sensitivity depth corresponding to 95% detection efficiency,  $\mathcal{D}^{95\%}$ . This is a commonly used figure of merit in CW searches. A higher depth values means being able to detect weaker signals, as they are buried deeper in the noise. This procedure is customarily employed in CW searches in order to assess their sensitivity [52, 17, 18, 19, 20]. An example of the interpolation procedure for one band is provided in Fig. 3.5.

Our previous discussion on the implications of a population threshold can be extended using the 95% detection efficiency depth. In Fig. 3.6 we compare the relative improvement obtained with respect to the O2 clustering by simply using TTD from Eq. (3.17) rather than Eqs. (3.1), (3.2) and (3.3).

By properly tuning the coincidence threshold, the usage of our proposal under the same conditions as the O2 clustering achieves a better sensitivity. As already discussed, the application of a population threshold can be compensated by a higher coincidence threshold. The specific increase notably depends on the frequency band (and more concretely, on the template density of the parameter-space region) under study.

We summarize the improvement in sensitivity in Fig. 3.7, comparing our 95% efficiency sensitivity depth estimates against the results reported in Table 3.1. Results using TTD are maximized with respect




FIGURE 3.6: Relative sensitivity improvement at two particular frequency bands due to the usage of TTD as a function of the coincidence threshold and minimum population parameters. Squares and a solid line represent no population threshold; circles and a dashed line represent a minimum population of one; inverted triangles and a dotted line represent a minimum population of three.




FIGURE 3.7: Comparison of 95% sensitivity depths from the second set of CW software injections obtained using TTD against depths from [20]. Each point corresponds to the depth estimated from injections in one frequency band. The dashed line represents equality. Vertical uncertainties are derived from the fitting procedure as illustrated in Fig. 3.5. Horizontal uncertainties are those reported in Table 3.1.

to the employed coincidence threshold, deeming them to be insensitive to the imposition of a minimum population threshold. We report a 5 – 15% improvement by the sole use TTD of instead of an Euclidean ansatz.

### 3.6 Conclusion

We have introduced a new parameter-space distance for continuous gravitational-wave searches by relating parameter-space points to their corresponding instantaneous frequency evolution. As opposed to previous approaches, this proposal provides a simple way of comparing candidates from CW semicoherent searches in a consistent fashion with respect to the underlying parameter-space structure to which the searches themselves are sensitive.

We have demonstrated this consistency of the new distance with the well-known behavior of semicoherent searches by reproducing a well-understood parameter-space structure related to the Earth-induced Doppler modulation of ground-based detectors, the “circles in the sky”. The new distance can be computed as an average over a subset of timestamps from a given data set, significantly reducing its computational cost at negligible numerical accuracy loss.

Our proposal is effective in improving practical CW searches, as we have demonstrated by using it to improve the candidate clustering procedure from the O2 open data search for CWs in binary systems of [20] and reproducing their sensitivity estimation procedure. After illustrating the role played by different parameters in the procedure, with the best settings we have obtained a 5 – 15% sensitivity improvement from using the new distance in the clustering step.

This distance definition can be seamlessly extended to any kind of quasi-monochromatic CW signal, provided we are able to describe its associated frequency evolution, as the main operational quantity is not provided by parameter-space points, but by their embedding into the time-frequency plane according to the CW model under consideration. This makes it a valuable tool to develop new post-processing stages and improve the sensitivity of wide parameter-space searches.



# Bibliography

- [1] R. Abbott *et al.*, “GWTC-2: Compact Binary Coalescences Observed by LIGO and Virgo During the First Half of the Third Observing Run.” arXiv:2010.14527 [gr-qc].
- [2] J. Aasi *et al.*, “Advanced LIGO,” *Classical and Quantum Gravity*, vol. 32, p. 074001, mar 2015.
- [3] F. Acernese *et al.*, “Advanced virgo: a second-generation interferometric gravitational wave detector,” *Classical and Quantum Gravity*, vol. 32, p. 024001, dec 2014.
- [4] M. Sieniawska and M. Bejger, “Continuous gravitational waves from neutron stars: current status and prospects,” *Universe*, vol. 5, no. 11, p. 217, 2019.
- [5] S. J. Zhu, M. Baryakhtar, M. A. Papa, D. Tsuna, N. Kawanaka, and H.-B. Eggenstein, “Characterizing the continuous gravitational-wave signal from boson clouds around galactic isolated black holes,” *Physical Review D*, vol. 102, p. 063020, Sep 2020.
- [6] R. Prix, “Search for continuous gravitational waves: Metric of the multi-detector F-statistic,” *Physical Review D*, vol. 75, p. 023004, 2007. [Erratum: *Physical Review D* 75, 069901 (2007)].
- [7] B. Krishnan, A. M. Sintes, M. A. Papa, B. F. Schutz, S. Frasca, and C. Palomba, “Hough transform search for continuous gravitational waves,” *Physical Review D*, vol. 70, p. 082001, Oct 2004.
- [8] C. Cutler, I. Gholami, and B. Krishnan, “Improved stack-slide searches for gravitational-wave pulsars,” *Physical Review D*, vol. 72, p. 042004, 2005.
- [9] P. Astone, A. Colla, S. D’Antonio, S. Frasca, and C. Palomba, “Method for all-sky searches of continuous gravitational wave signals using the frequency-Hough transform,” *Physical Review D*, vol. 90, p. 042002, Aug. 2014.
- [10] S. D’Antonio *et al.*, “Semicoherent analysis method to search for continuous gravitational waves emitted by ultralight boson clouds around spinning black holes,” *Physical Review D*, vol. 98, no. 10, p. 103017, 2018.
- [11] P. B. Covas and A. M. Sintes, “New method to search for continuous gravitational waves from unknown neutron stars in binary systems,” *Physical Review D*, vol. 99, p. 124019, Jun 2019.
- [12] H. J. Pletsch, “Parameter-space metric of semicoherent searches for continuous gravitational waves,” *Physical Review D*, vol. 82, p. 042002, 2010.
- [13] P. Leaci and R. Prix, “Directed searches for continuous gravitational waves from binary systems: parameter-space metrics and optimal Scorpius X-1 sensitivity,” *Physical Review D*, vol. 91, no. 10, p. 102003, 2015.
- [14] V. Dergachev and M. A. Papa, “Sensitivity improvements in the search for periodic gravitational waves using O1 LIGO data,” *Physical Review Letters*, vol. 123, no. 10, p. 101101, 2019.
- [15] R. Prix and M. Shaltev, “Search for Continuous Gravitational Waves: Optimal StackSlide method at fixed computing cost,” *Physical Review D*, vol. 85, p. 084010, 2012.
- [16] B. P. Abbott *et al.*, “First low-frequency einstein@home all-sky search for continuous gravitational waves in advanced ligo data,” *Physical Review D*, vol. 96, p. 122004, Dec 2017.
- [17] B. P. Abbott *et al.*, “All-sky search for periodic gravitational waves in the o1 ligo data,” *Physical Review D*, vol. 96, p. 062002, Sep 2017.

- [18] B. P. Abbot *et al.*, “Full band all-sky search for periodic gravitational waves in the O1 LIGO data,” *Physical Review D*, vol. 97, p. 102003, May 2018.
- [19] B. P. Abbott *et al.*, “All-sky search for continuous gravitational waves from isolated neutron stars using Advanced LIGO O2 data,” *Physical Review D*, vol. 100, p. 024004, Jul 2019.
- [20] P. B. Covas and A. M. Sintes, “First all-sky search for continuous gravitational-wave signals from unknown neutron stars in binary systems using advanced ligo data,” *Physical Review Letters*, vol. 124, p. 191102, May 2020.
- [21] B. Steltner, M. A. Papa, H.-B. Eggenstein, B. Allen, V. Dergachev, R. Prix, B. Machenschalk, S. Walsh, S. Zhu, and S. Kwang, 9 2020.
- [22] P. Jaranowski, A. Królak, and B. F. Schutz, “Data analysis of gravitational-wave signals from spinning neutron stars: The signal and its detection,” *Physical Review D*, vol. 58, p. 063001, Aug 1998.
- [23] C. Cutler and B. F. Schutz, “Generalized  $\mathcal{F}$ -statistic: Multiple detectors and multiple gravitational wave pulsars,” *Physical Review D*, vol. 72, p. 063006, Sep 2005.
- [24] K. Wette and R. Prix, “Flat parameter-space metric for all-sky searches for gravitational-wave pulsars,” *Physical Review D*, vol. 88, no. 12, p. 123005, 2013.
- [25] K. Wette, “Lattice template placement for coherent all-sky searches for gravitational-wave pulsars,” *Physical Review D*, vol. 90, no. 12, p. 122010, 2014.
- [26] B. Allen, “Spherical ansatz for parameter-space metrics,” *Physical Review D*, vol. 100, no. 12, p. 124004, 2019.
- [27] K. Wette, “Parameter-space metric for all-sky semicoherent searches for gravitational-wave pulsars,” *Physical Review D*, vol. 92, no. 8, p. 082003, 2015.
- [28] K. Wette, S. Walsh, R. Prix, and M. A. Papa, “Implementing a semicoherent search for continuous gravitational waves using optimally-constructed template banks,” *Physical Review D*, vol. 97, no. 12, p. 123016, 2018.
- [29] A. Singh, M. A. Papa, H.-B. Eggenstein, and S. Walsh, “Adaptive clustering procedure for continuous gravitational wave searches,” *Physical Review D*, vol. 96, p. 082003, Oct 2017.
- [30] F. Morawski, M. Bejger, and P. Ciecielag, “Convolutional neural network classifier for the output of the time-domain f-statistic all-sky search for continuous gravitational waves,” *Machine Learning: Science and Technology*, vol. 1, p. arXiv:1907.06917, jun 2020.
- [31] B. Beheshtipour and M. A. Papa, “Deep learning for clustering of continuous gravitational wave candidates,” *Physical Review D*, vol. 101, p. 064009, Mar 2020.
- [32] B. Beheshtipour and M. A. Papa, “Deep learning for clustering of continuous gravitational wave candidates II: identification of low-SNR candidates,” *arXiv e-prints*, 12 2020.
- [33] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.
- [34] J. T. Whelan, R. Prix, C. J. Cutler, and J. L. Willis, “New Coordinates for the Amplitude Parameter Space of Continuous Gravitational Waves,” *Classical and Quantum Gravity*, vol. 31, p. 065002, 2014.
- [35] J. Neyman, E. S. Pearson, and K. Pearson, “IX. on the problem of the most efficient tests of statistical hypotheses,” *Philosophical Transactions of the Royal Society of London A*, vol. 231, no. 694-706, pp. 289–337, 1933.
- [36] R. Prix and Y. Itoh, “Global parameter-space correlations of coherent searches for continuous gravitational waves,” *Classical and Quantum Gravity*, vol. 22, pp. S1003–S1012, sep 2005.

- [37] H. J. Pletsch, “Parameter-space correlations of the optimal statistic for continuous gravitational-wave detection,” *Physical Review D*, vol. 78, p. 102005, Nov 2008.
- [38] P. Jaranowski and A. Królak, “Data analysis of gravitational-wave signals from spinning neutron stars. ii. accuracy of estimation of parameters,” *Physical Review D*, vol. 59, p. 063003, Feb 1999.
- [39] R. Abbott *et al.*, “Open data from the first and second observing runs of advanced ligo and advanced virgo,” *SoftwareX*, vol. 13, p. 100658, 2021.
- [40] LIGO Scientific Collaboration and Virgo Collaboration, “Gravitational Wave Open Science Center - Advanced LIGO O2 Data Release.” <https://www.gw-openscience.org>, 2019.
- [41] LIGO Scientific Collaboration, “LIGO Algorithm Library - LALSuite.” free software (GPL), 2018.
- [42] M. Shaltev, P. Leaci, M. A. Papa, and R. Prix, “Fully coherent follow-up of continuous gravitational-wave candidates: an application to Einstein@Home results,” *Physical Review D*, vol. 89, no. 12, p. 124030, 2014.
- [43] J. Aasi *et al.*, “First low frequency all-sky search for continuous gravitational wave signals,” *Physical Review D*, vol. 93, no. 4, p. 042007, 2016.
- [44] M. A. Papa *et al.*, “Hierarchical follow-up of subthreshold candidates of an all-sky Einstein@Home search for continuous gravitational waves on LIGO sixth science run data,” *Physical Review D*, vol. 94, no. 12, p. 122006, 2016.
- [45] M. Sieniawska, M. Bejger, P. Cieciela, and A. Królak, “Followup procedure in time-domain F-statistic searches for continuous gravitational waves,” in *XXXVIII Polish Astronomical Society Meeting* (A. Różańska, ed.), vol. 7, pp. 37–40, Aug. 2018.
- [46] G. Ashton and R. Prix, “Hierarchical multistage MCMC follow-up of continuous gravitational wave candidates,” *Physical Review D*, vol. 97, no. 10, p. 103020, 2018.
- [47] D. Keitel, R. Tenorio, G. Ashton, and R. Prix, “PyFstat: a Python package for continuous gravitational-wave data analysis.” arXiv:2101.10915 [gr-qc].
- [48] E. Goetz, “Segments used for creating standard SFTs in O2 data.” <https://dcc.ligo.org/LIGO-T1900085/public>.
- [49] B. Behnke, M. A. Papa, and R. Prix, “Postprocessing methods used in the search for continuous gravitational-wave signals from the galactic center,” *Physical Review D*, vol. 91, p. 064007, Mar 2015.
- [50] C. Dreissigacker, R. Prix, and K. Wette, “Fast and Accurate Sensitivity Estimation for Continuous-Gravitational-Wave Searches,” *Physical Review D*, vol. 98, no. 8, p. 084058, 2018.
- [51] P. B. Covas *et al.*, “Identification and mitigation of narrow spectral artifacts that degrade searches for persistent gravitational waves in the first two observing runs of Advanced LIGO,” *Physical Review D*, vol. 97, p. 082002, Apr. 2018.
- [52] S. Walsh, M. Pitkin, M. Oliver, S. D’Antonio, V. Dergachev, A. Królak, *et al.*, “Comparison of methods for the detection of gravitational waves from unknown neutron stars,” *Physical Review D*, vol. 94, no. 12, p. 124010, 2016.





## Chapter 4

# Application of a hierarchical MCMC follow-up to Advanced LIGO continuous gravitational-wave candidates

This chapter is an adaptation of the material presented in

*Application of a hierarchical MCMC follow-up to Advanced LIGO continuous gravitational-wave candidates*

Rodrigo Tenorio, David Keitel, Alicia M. Sintes

[Phys. Rev. D 104, 084012 \(2021\)](#) – [arXiv:2105.13860 \[gr-qc\]](#)

DOI: 10.1103/PhysRevD.104.084012

### 4.1 Introduction

Continuous gravitational waves (CWs) are persistent forms of gravitational radiation. These yet-to-be detected signals are orders of magnitude weaker than compact binary coalescences [1], requiring long integration times (months to years) to differentiate them from noise. Potentially detectable sources using the current generation of ground-based interferometric detectors, Advanced LIGO [2] and Advanced Virgo [3], are neutron stars (NSs) presenting some non-axisymmetry such as crustal deformations, r-mode instabilities or free precession [4], or the annihilation of ultralight boson clouds around spinning black holes [5].

Searching for a CW consists in filtering a data stream against a set of signal templates, each of which is related to a certain set of parameters describing the CW model being searched for. The number of templates required to properly cover a certain parameter space region, however, scales as a large power of observing time [6]. At a fixed computing cost, the optimal strategy is to split the data stream into segments on which the filtering is performed, and then combine the resulting statistics [7, 8]. Since phase information is only fully preserved within each of these segments, they are usually referred to as *coherent segments* spanning a certain *coherence time*.

The approach taken by current implementations of wide parameter space searches such as [9, 10, 11, 12, 13, 14] lies in the middle ground. Wide parameter space regions are analyzed using a relatively low coherence time, ranging from half an hour to a few weeks. Surviving outliers are then sieved through a suite of vetoes testing their (in)consistency with a CW signal; this includes studying their persistence over the data stream, comparing their significance in different detectors or checking whether they cross a frequency band containing known instrumental artifacts [15, 16, 17, 18, 19, 20, 21]. Other common strategies are coincidence analyses between detectors or clustering neighbouring outliers in order to relate them to a common cause [22, 23, 24, 25]. Finally, if there are any surviving outliers, various follow-up strategies use longer coherence times [26, 27, 28, 29, 30, 31, 21], either in a single stage or in a hierarchical scheme where candidates are narrowed down over a “ladder” of coherence times.

Large-scale CW searches would benefit from a simple, general hierarchical setup, as it would allow for the systematic follow-up of CW outliers using longer coherence times, imposing tighter constraints and reducing the presence of outliers due to background noise.

Here we present the first complete framework to conduct hierarchical Markov Chain Monte Carlo (MCMC) follow-ups and its application to a set of outliers obtained by different CW search pipelines on Advanced LIGO O2 data. Our work builds on top of [32], which introduced the MCMC follow-up of CW outliers and studied its performance on simulated signals in pure Gaussian noise. We propose a new hypothesis test for the presence of a signal in the data after the full follow-up procedure. The probability of the signal hypothesis is derived from first principles as proposed in [33]; the probability of the noise hypothesis is derived from the application of extreme value theory. We demonstrate the general applicability of this follow-up strategy by analyzing outliers stemming from different analysis pipelines.

Although we restrict ourselves to outliers from CW searches for unknown *isolated* sources, this framework and the corresponding software [34] can also be applied to outliers from searches for sources in binary systems [30, 21], glitching NSs [35] and long-duration gravitational-wave transients [36, 37].

The paper is organized as follows: Section 4.2 describes the basic tools of CW data analysis and overviews the application of MCMC samplers to the follow-up problem; Sec. 4.3 introduces a new statistic in terms of hypothesis testing; Sec. 4.4 introduces the O2 outliers to be analyzed and the follow-up setup. The results are presented in Sec. 4.5, concluding in Sec. 4.6. We briefly comment on the statistical properties of the maximum  $\mathcal{F}$ -statistic over correlated templates in appendix 4.A.

## 4.2 Continuous-wave data analysis: Search and follow up

A CW signal can be parametrized in terms of two families of parameters, namely the *phase-evolution parameters*  $\lambda$  and the *amplitude parameters*  $\mathcal{A}$ . This separation is motivated by the response of a GW detector to such signals

$$h(t; \lambda, \mathcal{A}) = \sum_{\mu=0}^3 \mathcal{A}^\mu h_\mu(t; \lambda), \quad (4.1)$$

where the functions  $\mathcal{A}^\mu$  are independent of time [38].

The search for a CW signal can be stated in a Bayesian framework as a hypothesis test between the noise hypothesis  $\mathcal{H}_G$ , under which the data consists of Gaussian noise  $n(t)$ , and the signal hypothesis  $\mathcal{H}_S(\lambda, \mathcal{A})$ , supporting the presence of a CW signal with a defined set of parameters within said noise  $n(t) + h(t; \lambda, \mathcal{A})$ . The support of a stream of data  $x$  for either of these hypotheses is quantified by the *Bayes factor* [39]

$$B_{S/G}(x; \lambda, \mathcal{A}) = \frac{P(x|\mathcal{H}_S(\lambda, \mathcal{A}))}{P(x|\mathcal{H}_G)}. \quad (4.2)$$

Following [40, 41], and motivated by the linear dependency of Eq. (4.1) on the amplitude functions  $\mathcal{A}^\mu$ , one can choose an appropriate set of priors  $P(\mathcal{A})$  such that Eq. (4.2) can be analytically marginalized:

$$B_{S/G}(x; \lambda) = \int d\mathcal{A} B_{S/G}(x; \lambda, \mathcal{A}) P(\mathcal{A}) \propto e^{\mathcal{F}(x; \lambda)}. \quad (4.3)$$

The  $\mathcal{F}$ -statistic, which depends only on the data and the phase parameters, was originally derived as the maximum-likelihood estimator with respect to  $\mathcal{A}$  [38, 42]. This is a general detection statistic which only relies on the waveform decomposition presented in Eq. (4.1) and hence can be applied also to variations of the CW signal model such as sources in binaries [43] and transients [36]. Furthermore, the methods developed in this work can also be applied to CW outliers from any kind of search using a different detection statistic, as long as they can be associated with a parameter-space point with a certain uncertainty.

The role of Eq. (4.3) is to update the prior probability on the phase evolution parameters  $P(\lambda)$  by means of the information conveyed by the data stream  $x$ . This can be stated in terms of Bayes' theorem as

$$P(\lambda|x, \mathcal{H}_S) \propto B_{S/G}(x; \lambda) P(\lambda). \quad (4.4)$$

We note that  $B_{S/G}$  and  $\mathcal{F}$  have the same statistical power as they are related by a strictly monotonic function. For the sake of later consistency, we will focus on  $\mathcal{F}$  from now on. We refer the reader to [40, 36, 44, 45] for a more in-depth analysis of these statistics.

The detection problem is now stated in terms of a maximization: Given a stream of data  $x$ , we are interested in finding the phase-evolution parameters  $\lambda$  (also referred to as *templates*) which maximize Eq. (4.4) or, equivalently,  $\mathcal{F}(x; \lambda)$ .

### 4.2.1 Coherent and semicoherent searches

The fully-coherent  $\mathcal{F}$ -statistic can be expressed in terms of a linear filter between the data stream and a signal template,

$$\tilde{\mathcal{F}}(\lambda) \propto |\langle x, h(\lambda) \rangle|^2, \quad (4.5)$$

where  $\langle \cdot \rangle$  represents a functional scalar product. Throughout this work, and following the convention of [46], fully-coherent quantities will be represented with a tilde; semicoherent quantities, introduced in Eq. (4.7), will be represented with a caret. The response of  $\tilde{\mathcal{F}}$  to an offset  $\Delta\lambda$  in the phase-evolution parameters  $\lambda$  is quantified using the *mismatch* [47], which can be defined in terms of a local quadratic approximation around the true signal parameters  $\lambda$  where the mismatch has a minimum:

$$\mathbf{m}(\Delta\lambda; \lambda) = \frac{\tilde{\mathcal{F}}(\lambda) - \tilde{\mathcal{F}}(\lambda + \Delta\lambda)}{\tilde{\mathcal{F}}(\lambda)} \simeq \Delta\lambda^T \cdot \bar{g} \cdot \Delta\lambda + \mathcal{O}(\Delta\lambda^3). \quad (4.6)$$

The symmetric tensor  $\bar{g}$  is referred to as the parameter-space metric, and can be used to set up parameter-space coverings, also known as template banks, at a certain mismatch level. This quadratic approximation is known to be valid up to  $\mathbf{m} \lesssim 0.3 - 0.5$ , although latest developments on this subject suggest to further extend the approximation up to  $\mathbf{m} \sim 1$  [48, 49, 50].

Maximizing Eq. (4.5) poses a computational challenge, as the number of templates to be considered in the optimization scales with a large power of the total length of the data stream [51, 6], while the sensitivity only scales as the square root of it [52, 53]. As discussed in [54, 13], such a strong scaling stems from the tight restrictions imposed by the  $\mathcal{F}$ -statistic on the signal model, requiring phase coherence over the whole duration of the data stream. A looser statistic can be constructed by imposing said coherence in a segment-wise manner. To do so, the data stream, spanning a time of  $T_{\text{obs}}$ , is divided into  $N_{\text{seg}}$  segments, each of them with a duration of  $T_{\text{coh}}$ . The *semicoherent*  $\mathcal{F}$ -statistic is then constructed by adding the coherent  $\mathcal{F}$ -statistics computed in each segment

$$\hat{\mathcal{F}}(\lambda) = \sum_{n=0}^{N_{\text{seg}}-1} \tilde{\mathcal{F}}_n(\lambda), \quad (4.7)$$

where  $\tilde{\mathcal{F}}_n$  refers to the coherent  $\mathcal{F}$ -statistic computed using only data within segment  $n$ . This approach uncorrelates the template's phase-evolution between consecutive coherent segments, loosening the constraints imposed on the data and widening  $\mathcal{F}$ -statistic peaks in the parameter space [54, 32]. In other words, given a parameter-space coordinate volume, the number of templates required to cover it at a given mismatch *decreases* with lower  $T_{\text{coh}}$ . This implies a dependency of the parameter space metric  $\bar{g}$  on  $T_{\text{coh}}$ .

The optimal strategy to sweep a wide parameter-space region under a controlled computational budget is then to use a hierarchical scheme with a varying  $T_{\text{coh}}$ : The first stage surveys a parameter-space region with  $T_{\text{coh}} \ll T_{\text{obs}}$ , using an affordable number of templates.  $\mathcal{F}$ -statistic outliers are then analyzed with an increased coherence time, further narrowing down the parameter-space region of interest. This process continues either until  $T_{\text{coh}} = T_{\text{obs}}$  or the candidate is vetoed by a complementary procedure [8, 46, 28, 32].

### 4.2.2 MCMC-based follow-ups

The follow-up of CW outliers requires to set up a template bank across the parameter-space region of interest. Typical gridded approaches use a parameter-space metric to cover the parameter space at fixed maximum mismatch [28, 27, 55]. This approach usually requires an extensive campaign of software injections to be performed in order to calibrate the optimal set up in terms of sensitivity and computing cost [28].

Alternatively, one could view the problem from the point of view of Bayesian inference. Eq. (4.4) relates the  $\mathcal{F}$ -statistic to a posterior probability distribution. This distribution can be sampled using a

MCMC method, effectively constructing an adaptive random template bank in the parameter space in which  $\mathcal{F}$ -statistic values will be more densely evaluated around high posterior probability regions. As first discussed in [32], this approach achieves close to the theoretical optimal sensitivity for signals in Gaussian noise as long as the parameter space region is small enough to ensure a good convergence of the MCMC.

For the purpose of estimating the effectiveness of an MCMC, as discussed in [32], the effective size of a parameter-space region can be computed in terms of the *number of templates*  $\mathcal{N}$  required to cover it at a mismatch of unity using a lattice with unit normalized thickness [56, 51]

$$\mathcal{N}(T_{\text{coh}}, \Delta\lambda) = \int_{\Delta\lambda} d\lambda \sqrt{g(T_{\text{coh}})}, \quad (4.8)$$

where  $g(T_{\text{coh}})$  is the determinant of the parameter-space metric, which depends on  $T_{\text{coh}}$  as explained in Sec. 4.2.1, and  $\Delta\lambda$  represents the region being followed up. The integral in Eq. (4.8) must be computed along the resolved parameter-space dimensions only; i.e., one should *not* include fractional templates, as doing so would underestimate the actual number of templates [46, 8, 43]. For a follow-up search, the parameter-space region under analysis is typically smaller than the scale of parameter-space correlations, meaning  $\sqrt{g}$  can be taken out of the integral as a constant and Eq. (4.8) simplifies to

$$\mathcal{N}(T_{\text{coh}}, \Delta\lambda) \simeq \sqrt{g(T_{\text{coh}})} \text{Vol}(\Delta\lambda) \quad (4.9)$$

where  $\text{Vol}(\Delta\lambda)$  is the coordinate volume of the region being followed up. Seminal analyses in [32] and follow-up searches performed in [30, 21] suggest that values up to  $\mathcal{N} \simeq 10^{3-4}$  are compatible with effective MCMC runs in terms of convergence.

CW outliers are identified as a parameter-space point carrying an uncertainty which depends on the pipeline used to conduct the search. Upon entering the follow-up pipeline, these uncertainties are converted into prior probability distributions to start the MCMC sampling. Ref. [32] proposed the use of bounded uniform priors in order to restrict the surveyed parameter-space region; however, such hard boundaries may prevent the successful follow-up of CW candidates whose parameters are shifted due to the presence of parameter-space correlations. We propose the use of uncorrelated Gaussian priors, which concentrate their probability density around a characteristic region while being unbounded. See Sec. 4.5 for details on the choice of Gaussian priors.

An MCMC-based follow-up is implemented in the `PyFstat` package [34] using the parallel-tempered ensemble MCMC sampler `ptemcee` [57, 58] to sample the posterior distribution Eq. (4.4) using either the coherent (Eq. 4.5) or semicoherent (Eq. 4.7)  $\mathcal{F}$ -statistic. We refer the reader to [32] for an extended discussion on the characteristics of this particular MCMC implementation. The analyses presented in this work were performed using `PyFstat` version 1.11.3 [59].

### 4.2.3 A coherence-time ladder

Early setups of hierarchical schemes were based on the optimization of computing resources in order to achieve a prescribed level of sensitivity [7, 8]. Alternatively, if no computational cost model was available, software injection campaigns were used to calibrate the number of stages [28]. For the case of an MCMC-based follow-up, one can use the quantity  $\mathcal{N}(T_{\text{coh}}, \Delta\lambda)$  to design a hierarchical scheme by imposing the proper convergence of the MCMC run at each stage [32].

Suppose a wide parameter-space semicoherent search produces an interesting outlier in the parameter-space region  $\Delta\lambda^{(0)}$ , where the exact shape is entirely dependent on the pipeline. Round-bracketed superindices denote different stages of the follow-up. To set up a first follow-up stage, we choose a coherence time  $T_{\text{coh}}^{(0)}$  such that  $\mathcal{N}(T_{\text{coh}}^{(0)}, \Delta\lambda^{(0)}) \lesssim \mathcal{N}^*$ , ensuring the effective parameter-space resolution is coarse enough for the MCMC algorithm to properly converge towards the region of interest. If successful, the resulting parameter-space region will be narrower,  $\Delta\lambda^{(1)} \leq \Delta\lambda^{(0)}$ , and a second MCMC stage using a new coherence time  $T_{\text{coh}}^{(1)}$  will be applied. This procedure is repeated until  $T_{\text{coh}} = T_{\text{obs}}$  and a final fully-coherent follow-up is performed.

In [32], a simple method was proposed to find the coherence time for a stage  $j$  given the previous stage's results. The idea is to increase the coherence time as much as possible such that the MCMC is able to converge to the target distribution. Since this convergence can be quantified in terms of a maximum number of templates within a region  $\mathcal{N}^*$ , the new coherence time  $T_{\text{coh}}^{(j)}$  can be obtained by

solving

$$\mathcal{N}(T_{\text{coh}}^{(j)}, \Delta\lambda^{(j)}) = \mathcal{N}^*. \quad (4.10)$$

This choice minimizes the number of stages in the scheme, reducing the overall computing cost, while ensuring the effectiveness of the MCMC approach. The explicit dependency of Eq. (4.10) on the parameter-space region under analysis  $\Delta\lambda^{(j)}$ , however, hinders the construction of a complete hierarchical scheme.

This dependency can be removed by noticing the inherent self-similarity of MCMC stages: A successful MCMC follow-up stage ends up with a set of samples around a prominent global maximum, the fine structure of which is *underresolved* because of the chosen coherence time. By progressing to the next stage, this fine structure gets resolved and the MCMC *zooms in* further towards the parameter-space maximum. The setup of a coherence-time ladder is simply a problem of minimizing the number of stages to reduce computing cost while maintaining sufficiently big underresolved regions for the MCMC follow-up to properly sample the region of interest. This condition can be simply expressed as  $\mathcal{N}(T_{\text{coh}}^{(j)}, \Delta\lambda^{(j+1)}) \simeq 1$ ; hence, comparing consecutive stages factors out the problematic dependency and the hierarchical scheme can be constructed by solving the recurrence

$$\mathcal{N}^* \simeq \frac{\mathcal{N}(T_{\text{coh}}^{(j+1)}, \Delta\lambda^{(j+1)})}{\mathcal{N}(T_{\text{coh}}^{(j)}, \Delta\lambda^{(j+1)})} = \frac{\sqrt{g(T_{\text{coh}}^{(j+1)})}}{\sqrt{g(T_{\text{coh}}^{(j)})}} \quad (4.11)$$

given  $T_{\text{coh}}^{(0)}$  and  $\mathcal{N}^*$ . A numerical solver for Eq. (4.11) is included in the `PyFstat` package [34]. Constructing the coherence-time ladder as proposed by [32] makes use of the so-called SuperSky metric [60, 61] to compute the parameter-space volume element. This metric is numerically well-conditioned, but requires  $T_{\text{coh}} \gtrsim 1$  day.

Alternatively, we derive an equivalent coherence-time ladder by considering the parameter-space volume reduction from one stage to the next. Let us define

$$\gamma^{(j+1)} = \frac{\text{Vol}(\Delta\lambda^{(j)})}{\text{Vol}(\Delta\lambda^{(j+1)})} \quad (4.12)$$

as the parameter-space volume *shrinkage* from stage  $j$  to stage  $j+1$ . In a practical application, this quantity can be computed by comparing the volume containing a certain amount of posterior probability from two consecutive stages.

Eq. (4.10) can now be re-expressed as

$$1 = \frac{\mathcal{N}(T_{\text{coh}}^{(j+1)}, \Delta\lambda^{(j+1)})}{\mathcal{N}(T_{\text{coh}}^{(j)}, \Delta\lambda^{(j)})}, \quad (4.13)$$

and Eq. (4.11) is generalized by including Eq. (4.12)

$$\gamma^{(j+1)} = \frac{\sqrt{g(T_{\text{coh}}^{(j+1)})}}{\sqrt{g(T_{\text{coh}}^{(j)})}}, \quad (4.14)$$

where we can recognize  $\gamma^{(j+1)}$  as a generalized version of the refinement factor  $\gamma$  introduced in Eq. (73) of [62] to account for the template bank refinement from a semicoherent stage to a fully-coherent one. To fully recover Eq. (4.11), we simply set  $\gamma^{(j+1)} = \mathcal{N}^*$  in every stage  $j$ .

According to this derivation, constructing a coherence ladder is equivalent to imposing a ratio of posterior volume shrinkage. For example, choosing  $\mathcal{N}^* \simeq 10^4$  is equivalent to imposing an overall volume shrinkage of  $\gamma \simeq 10^4$  (i.e. posterior volume is a ten-thousandth fraction of the prior volume) at each step of the ladder. As a result, the behaviour of an MCMC stage is dependent upon its capability to fulfill the required shrinkage rate.

### 4.3 Evaluating the hierarchical follow-up with a Bayes factor

A multi-stage MCMC follow-up analyzes CW outliers by converging towards parameter-space regions with a high posterior probability. After each stage, coherence time is increased, breaking up under-resolved regions into smaller ones and allowing the MCMC to further narrow down the parameters associated to the loudest outlier. We are interested in evaluating the significance of the loudest template resulting from the multi-stage follow-up by comparing its actual fully-coherent  $\mathcal{F}$ -statistic to the expected value predicted by a previous stage of the ladder.

We construct a new Bayes factor for this comparison, using the fully-coherent  $\mathcal{F}$ -statistic of the loudest candidate of the MCMC,  $2\tilde{\mathcal{F}}^*$ , in order to quantify the support for the presence or lack of a CW signal in the data. Following the definition in Eq. (4.2),

$$\ln \mathcal{B}_{S/N}^* = \ln \frac{\text{P}(2\tilde{\mathcal{F}}^*|\mathcal{H}_S)}{\text{P}(2\tilde{\mathcal{F}}^*|\mathcal{H}_N)}, \quad (4.15)$$

where the hypotheses  $\mathcal{H}_S$  and  $\mathcal{H}_N$  correspond to the presence or lack of a signal, respectively. As discussed in Sec. 4.3.1, the use of extreme value theory allows us to formulate  $\mathcal{H}_N$  such that it is not restricted to Gaussian noise, but includes any exponentially-bounded distributions with unbounded domain. The following subsections are devoted to deriving the probability distributions under each of these hypotheses.

#### 4.3.1 Noise hypothesis

The noise hypothesis  $\mathcal{H}_N$  ascribes the obtained value of  $2\tilde{\mathcal{F}}^*$  to pure noise. Under the presence of Gaussian noise, the coherent  $\mathcal{F}$ -statistic follows a chi-squared distribution with 4 degrees of freedom<sup>1</sup>,  $2\tilde{\mathcal{F}} \sim \chi_4^2$ . If we consider the resulting MCMC samples as a template bank  $\{\lambda\}$ , it is clear that  $2\tilde{\mathcal{F}}^* = \max_{\lambda \in \{\lambda\}} 2\tilde{\mathcal{F}}(\lambda)$  and the corresponding probability distribution is that of the maximum over a certain number of templates  $n$  [63]:

$$\text{P}(\max 2\tilde{\mathcal{F}}) = n \cdot \chi_4^2(\max 2\tilde{\mathcal{F}}) \cdot \left[ \int_0^{\max 2\tilde{\mathcal{F}}} d\xi \chi_4^2(\xi) \right]^{n-1}, \quad (4.16)$$

where  $\chi_4^2$  denotes the probability density function. The argument equally holds for the case of the *semicoherent*  $\mathcal{F}$ -statistic; in that case, however, the number of degrees of freedom of the chi-squared distribution would be  $4N_{\text{seg}}$ .

By construction, the *effective* number of templates in a CW template bank is different from the actual number of templates. This is because template banks are set up such that no parameter-space point is further than a certain mismatch  $\mathbf{m}$  from a template in the bank, implying a certain degree of correlation among neighbouring templates [56]. The problem of estimating the effective number of templates in a template bank has not found a definitive solution in the CW literature.

A common approach, see e.g. [64], is to evaluate the template bank on several realizations of Gaussian noise to numerically sample the probability distribution of the loudest outlier; the effective number of templates is then obtained by fitting  $n$  from Eq. (4.16) to the data. Another approach, firstly proposed in [65], splits the results of a wide parameter-space search into disjoint partitions such that they are equivalent to different realizations of a smaller search. The fraction of effective templates can be fitted using Eq. (4.16) to the loudest outlier per partition, obtaining  $n$  through extrapolation. Further developments on this method proposed a non-parametric ansatz to directly estimate the distribution of the loudest candidate of a search [66].

Here we will use a solution based on extreme value theory, which describes the three possible asymptotic distributions followed by the maximum of  $n$  independent trials according to the tail of their individual probability distribution. Short-scale correlated variables, such as the ones arising in the search for CW signals, are also covered by the theory [67]. The family of three distributions, usually referred to as the *generalized extreme value distribution*, is parametrized by a single parameter  $c \in \mathbb{R}$  (aside from the location and scale parameters), and encompasses every possible *max-stable* distribution: the

<sup>1</sup>We recall for the sake of consistency with the statistics literature that a chi-squared distribution with  $\nu$  degrees of freedom corresponds to a Gamma distribution with shape parameter  $k = \nu/2$  and scale parameter  $\theta = 2$ .




FIGURE 4.1: Distribution of the maximum  $2\mathcal{F}$  value of a template bank obtained from its evaluation at  $N_o = 600$  different off-sourced right ascensions, excluding  $90^\circ$  around the sky position of the outlier of interest. The template bank corresponds to MCMC samples from the fully-coherent stage follow-up of a simulated signal in Gaussian noise. The solid line represents the fit of a Gumbel distribution.

maximum value of a set of random variables following a generalized extreme value distribution follows itself a generalized extreme value distribution of the same class, albeit with different parameters. Each of the three possible distributions is related to  $c$  being positive, null or negative, and encloses a different set of probability distributions in its domain of attraction [68, 69, 70].

For our CW application, we focus on the case  $c = 0$ , also known as the *Gumbel distribution*

$$\text{Gumbel}(\xi; \mu, \sigma) = \frac{1}{\sigma} \exp \left[ - \left( \frac{\xi - \mu}{\sigma} \right) - e^{-\left( \frac{\xi - \mu}{\sigma} \right)} \right], \quad (4.17)$$

where  $\mu$  and  $\sigma$  are its location and scale parameters, respectively. The domain of attraction of this distribution comprises a variety of exponentially bounded distributions, including the chi-squared distribution. A similar procedure could be carried out for the other two families  $c \neq 0$ , including power-law and finite tails, if the behavior of the background noise required so. This argument is consistent with the empirical proposal of [71].

As noted in Appendix D of [53], the presence of correlated templates renders Eq. (4.16) unsuitable to describe the background noise distribution of CW searches. This is because the family of Gumbel distributions spanned by Eq. (4.16) as  $n \rightarrow \infty$  has a *fixed* scale parameter  $\sigma = 2$ . The inclusion of correlated templates makes the underlying distribution deviate from a chi-squared [72], but exponential tails still allow the distribution of the maxima to be described by a Gumbel distribution but with  $\sigma \neq 2$ . Further discussion on this topic is presented in Appendix 4.A.

As a result, we construct  $P(2\tilde{\mathcal{F}}^*|\mathcal{H}_N)$  by fitting *both* the location and scale parameters of a Gumbel distribution to the background distribution associated to  $2\tilde{\mathcal{F}}^*$

$$P(2\tilde{\mathcal{F}}^*|\mathcal{H}_N) = \frac{1}{\sigma_N} \exp \left[ - \left( \frac{2\tilde{\mathcal{F}}^* - \mu_N}{\sigma_N} \right) - e^{-\left( \frac{2\tilde{\mathcal{F}}^* - \mu_N}{\sigma_N} \right)} \right]. \quad (4.18)$$

This approach has the advantage of circumventing the computation of an effective number of templates by directly using the asymptotic distribution, the functional form of which is robust as long as the individual distribution tails fall off exponentially. The typical number of templates evaluated in an MCMC follow-up is consistent with a good convergence of the maximum distribution towards a Gumbel [73]. Further discussion on the suitable application of extreme value theory to evaluate the loudest outlier of a gravitational-wave search will be presented elsewhere [74].

To estimate the scale and location parameters of the background distribution  $\mu_N, \sigma_N$ , we apply the *off-sourcing* procedure, the effectiveness of which was studied in [75]. Off-sourcing consists in evaluating the  $\mathcal{F}$ -statistic on a template bank whose sky positions have been purposely *shifted* with respect to that of the outlier of interest. This blinds the detection statistic to the outlier under analysis while still sampling the same background distribution from the dataset. Incidentally, this takes into account

template-bank correlations induced by non-Gaussian noise components. These correlations do not arise due to different templates sampling the same spectrogram data (i. e. overlapping frequency-evolution tracks) [65, 25], but due to the presence of correlated spectrogram data spanning different iso-mismatch ellipsoids in the parameter space. The former kind is fundamental in the sense that it is independent of the background; the latter is entirely dependent upon the observed data: the wider the bandwidth of the disturbance, the lower the number of effective independent templates.

In our concrete application, we produce  $N_o = 600$  off-sourced template banks by randomly shifting the template's right ascension (azimuthal spherical angle), excluding a  $90^\circ$  region around the sky position of interest. The declination (polar spherical angle) is unchanged in order to maintain a constant level of sensitivity in terms of  $\mathcal{F}$ -statistic values. Fig. 4.1 shows an example of a background noise distribution obtained through this procedure.

The evaluation of off-sourced template banks represents the main contribution to the computing cost of the follow-up. The small number of outliers evaluated in this work allowed us to evaluate a set of off-sourced samples for each of them. For the case of a large-scale follow-up, however, one could benefit from the general properties of the Gumbel distribution to re-use a set of Gumbel parameters for different parameter-space regions, lowering the overall computing cost.

### 4.3.2 Signal hypothesis

The presence of a signal is characterized by its (squared) signal-to-noise ratio (SNR)  $\rho^2$ , which gauges the (squared) amplitude of a signal against that of the background noise [38, 42, 9, 47]. Exact expressions for  $\rho^2$ , which include amplitude-modulation effects due to the antenna pattern of the detectors, are available in [38, 52, 53]. The effect of this parameter on the probability distribution of the  $\mathcal{F}$ -statistic is to shift the chi-squared distribution towards a non-central chi-squared distribution,  $2\mathcal{F} \sim \chi_{4N_{\text{seg}}}^2(\rho^2)$ , were the fully-coherent case corresponds to  $N_{\text{seg}} = 1$ .

As previously discussed, we are interested in comparing the consistency of  $2\tilde{\mathcal{F}}^*$  to the values  $2\hat{\mathcal{F}}^*$  obtained in a previous stage of the ladder. Any semicoherent stage of the ladder can be used to construct a signal hypothesis; as discussed in more detail in Sec. 4.5.1, we select the second-to-last stage in order to benefit from the tighter constraints imposed by the signal model. For the remainder of this section we simplify our notation by removing the asterisks, assuming every  $\mathcal{F}$ -statistic value refers to that of the loudest candidate from the fully-coherent stage.

We construct  $P(2\tilde{\mathcal{F}}|\mathcal{H}_S)$  following the developments of [33]. The basic idea goes as follows: Assume a single-template search perfectly matching a signal is performed. The presence of a signal in the data, characterized by  $\rho^2$ , produces an  $\mathcal{F}$ -statistic value which depends only on  $\rho^2$  and the number of coherent segments  $N_{\text{seg}}$ . More specifically, obtaining a value of  $2\hat{\mathcal{F}}$  on  $N_{\text{seg}}$  segments automatically produces an estimate on  $\rho^2$ , which, in turn, yields an estimation of the expected  $2\tilde{\mathcal{F}}$  that will be retrieved after performing a fully-coherent search.

The exact flow of information from the semicoherent to the coherent statistic can be readily expressed by marginalizing over the unknown non-centrality parameter  $\rho^2$

$$P(2\tilde{\mathcal{F}}|\mathcal{H}_S) = \int_0^\infty d\rho^2 P(2\tilde{\mathcal{F}}|\rho^2, 2\hat{\mathcal{F}}, N_{\text{seg}}) P(\rho^2|2\hat{\mathcal{F}}, N_{\text{seg}}) \propto \int_0^\infty d\rho^2 P(2\tilde{\mathcal{F}}|\rho^2) P(2\hat{\mathcal{F}}|\rho^2, N_{\text{seg}}) P(\rho^2), \quad (4.19)$$

where constant factors with respect to  $2\hat{\mathcal{F}}$  and  $N_{\text{seg}}$  were omitted and the same data is being used to compute both statistics.<sup>2</sup> The choice of a prior distribution on  $\rho^2$  depends on the type of search carried out; for a wide parameter-space search such as the ones in which we are interested it is enough to consider an improper uniform prior.

In going to the second line in Eq. (4.19) we have assumed no dependency between  $2\hat{\mathcal{F}}$  and  $2\tilde{\mathcal{F}}$  in the sense of  $P(2\tilde{\mathcal{F}}|\rho^2, 2\hat{\mathcal{F}}, N_{\text{seg}}) = P(2\tilde{\mathcal{F}}|\rho^2)$ . This relation holds exactly if one computes each statistic on a different dataset, corresponding to the *fresh data mode* in [8]. On the other hand, if both statistics are evaluated on the same data, it represents a *conservative* choice in the sense of producing a wider distribution. This is because it neglects any correlations between  $2\hat{\mathcal{F}}$  and  $2\tilde{\mathcal{F}}$ . The lack of a simple way of quantifying correlations amongst said statistics in a general case justifies the safe approach of fresh data mode even though the same data is actually being used [33].

<sup>2</sup>This corresponds to  $\kappa = 1$  in the notation of [33].



Parameters in Gaussian noise	
$\mu_S$	$(3 - 10) \times 10^3$
$\mu_N$	10 - 20
$\sigma_S$	30 - 80
$\sigma_N$	2 - 3

TABLE 4.1: Typical location and scale parameters obtained from an injection campaign on Gaussian noise with an observing time of  $T_{\text{obs}} = 9$  months. Signal location and scale parameters were computed using the second-to-last stage of the coherence-time ladder. See Sec. 4.5 for further details.

The functional forms of the distributions in Eq. (4.19) have already been discussed in this subsection:

$$P(2\tilde{\mathcal{F}}|\rho^2) = \chi_4^2(2\tilde{\mathcal{F}}; \rho^2), \quad (4.20)$$

$$P(2\hat{\mathcal{F}}|\rho^2, N_{\text{seg}}) = \chi_{4N_{\text{seg}}}^2(2\hat{\mathcal{F}}; \rho^2). \quad (4.21)$$

It is useful to further simplify Eq. (4.19) to a closed analytical form. A proxy value for  $\rho^2$  can be obtained by simply subtracting the expected noise-only value of a chi-squared distribution with  $4N_{\text{seg}}$  degrees of freedom, namely  $\rho_0^2 = 2\hat{\mathcal{F}} - 4N_{\text{seg}}$ . Assuming  $\rho_0^2 \gg 1$ , chi-squared distributions can be replaced by Gaussian distributions [76, 77] and Eq. (4.19) can be further replaced by a Gaussian, the peak of which corresponds to  $\mu_S = \rho_0^2$ . We refer to [33] for further details on this derivation and simply quote the final result

$$P(2\tilde{\mathcal{F}}|2\hat{\mathcal{F}}, N_{\text{seg}}) = \text{Gauss}(2\tilde{\mathcal{F}}; \mu_S, \sigma_S), \quad (4.22)$$

where

$$\begin{aligned} \mu_S &= \rho_0^2, \\ \sigma_S^2 &= 8 \cdot (1 + N_{\text{seg}} + \rho_0^2). \end{aligned} \quad (4.23)$$

These expressions are useful to discuss the qualitative behavior of our newly proposed Bayes factor in different signal regimes. It will also be applicable in the analysis of software-injected signals in Sec. 4.5.1. However, due to the regime in which real-data outliers are typically found, we do not apply this Gaussian approximation to their analysis; instead, we numerically evaluate the full version of Eq. (4.19).

### 4.3.3 Bayes factor

We will now construct an overall Bayes factor to compare the two hypotheses supporting the presence or lack of a signal in a given stream of data. The distribution associated to the noise hypothesis, given in Eq. (4.18), is constructed by fitting the location and scale parameters of a Gumbel distribution to background data samples obtained through off-sourcing. The noise hypothesis can be defined in terms of said parameters, namely  $\mathcal{H}_N = \{\mu_N, \sigma_N\}$ , and the resulting distribution is

$$\ln P(2\tilde{\mathcal{F}}^*|\mathcal{H}_N) = - \left( \frac{2\tilde{\mathcal{F}}^* - \mu_N}{\sigma_N} + e^{-\left(\frac{2\tilde{\mathcal{F}}^* - \mu_N}{\sigma_N}\right)} + \ln \sigma_N \right). \quad (4.24)$$

The signal hypothesis compares the statistical behavior of the loudest candidate across different stages of the coherence-time ladder. We state the signal hypothesis as  $\mathcal{H}_S = \{\mu_S, \sigma_S\}$  and, to simplify the following discussion, we write everything in this section using the Gaussian approximation given in Eq. (4.22):

$$\ln P(2\tilde{\mathcal{F}}^*|\mathcal{H}_S) = -\frac{1}{2} \left[ \left( \frac{2\tilde{\mathcal{F}}^* - \mu_S}{\sigma_S} \right)^2 + \ln 2\pi\sigma_S \right]. \quad (4.25)$$

We note again that this approximated formula will *not* be applied to real-data candidates, as they are not located within the strong signal regime. Instead, we will then numerically evaluate Eq. (4.19).




FIGURE 4.2: Illustration of different regimes in which an outlier could be located. Shaded regions represent probability distributions associated to the indicated hypothesis. Dashed vertical lines refer to the enumerated labels in the text.

It is useful to introduce the following auxiliary variables

$$\xi_S = \frac{2\tilde{\mathcal{F}}^* - \mu_S}{\sigma_S}, \quad \xi_N = \frac{2\tilde{\mathcal{F}}^* - \mu_N}{\sigma_N}, \quad (4.26)$$

which measure the *discrepancy* of the retrieved  $2\tilde{\mathcal{F}}^*$  value with respect to the most probable values under the signal and noise hypothesis, respectively.

Combining Eqs. (4.24) and (4.25) we obtain an explicit expression for Eq. (4.15)

$$\ln \mathcal{B}_{S/N}^* = -\frac{1}{2}\xi_S^2 + \xi_N + e^{-\xi_N} + \ln \frac{\sigma_N}{\sqrt{2\pi}\sigma_S}. \quad (4.27)$$

Example values of the involved quantities for the use case later in this paper are summarized in Table 4.1. We proceed to analyze the general behavior of this new statistic under different conditions.

The operating point of wide parameter-space searches is generally such that outliers being followed up are significant enough so that  $\xi_N > 0$ , in the sense that a more sensitive method can be applied once the parameter-space region has been narrowed down. It is also reasonable to expect  $\mu_S > \mu_N$ , although this assumption may not be valid in case of very deep searches.

We distinguish three interesting regimes of behavior of Eq. (4.27), labeled in Fig. 4.2 using dashed vertical lines:

- a) The candidate is consistent with a noise fluctuation, returning  $\xi_N < \xi_S$ , hence  $\ln \mathcal{B}_{S/N}^* < 0$  and the signal hypothesis is disfavored.
- b) The candidate is consistent with the signal hypothesis  $\xi_S \sim 0$ ; hence, the dominant contribution to the Bayes factor is given by the discrepancy with respect to the noise hypothesis  $\ln \mathcal{B}_{S/N}^* \sim \xi_N$ . This is the expected behavior of a detection statistic: the favoring towards the signal hypothesis is directly proportional to the discrepancy with respect to background noise.
- c) The candidate is beyond the region expected by the signal hypothesis, meaning  $\ln \mathcal{B}_{S/N}^* \simeq -\frac{1}{2}\xi_S^2 + \xi_N$ . This novel behavior is due to the chosen signal hypothesis: As opposed to the  $\mathcal{F}$ -statistic's signal hypothesis, which results in a monotonic function of SNR, Eq. (4.19) establishes a particular region of interest centered at  $\xi_S = 0$ , penalizing deviations towards *both* sides of it.

A complementary description of Eq. (4.27) is shown in Fig. 4.3, where  $\ln \mathcal{B}_{S/N}^*$  is shown on the  $(\xi_N, \xi_S)$  plane. These two variables, which represent the discrepancy of  $2\tilde{\mathcal{F}}^*$  with respect to the noise and signal hypothesis, are related by

$$\xi_N = \frac{\sigma_S}{\sigma_N}\xi_S + \frac{\mu_S - \mu_N}{\sigma_N}, \quad (4.28)$$

meaning that once  $\{\mu_S, \sigma_S\}$  and  $\{\mu_N, \sigma_N\}$  are determined, the detection statistic is restricted to a straight line in  $(\xi_N, \xi_S)$ . This description also clarifies the behavior of  $\ln \mathcal{B}_{S/N}^*$  in case b) of Fig. 4.2. In said case,




FIGURE 4.3: Bayes factor in terms of the discrepancy of an outlier with respect to the noise and signal hypothesis as described in Eq. (4.27). Numerical values are computed using  $\sigma_N = 3$  and  $\sigma_S = 30$ , consistent with Table 4.1. This representation will be referred to as the  $(\xi_N, \xi_S)$  plane.

$\xi_S \sim 0$  and then  $\ln \mathcal{B}_{S/N}^* \sim (\mu_S - \mu_N)/\sigma_N$ , so that it is the *combined* action of a high SNR ( $\mu_S > \mu_N$ ) and a low discrepancy with respect to the expected value according to previous stages ( $\xi_S \sim 0$ ) what decides on the consistency of a CW candidate with respect to the signal or noise hypothesis.

A summary of the construction and practical computation of  $\ln \mathcal{B}_{S/N}^*$  is shown as a flowchart in Fig. 4.4.

## 4.4 Follow-up of outliers in LIGO O2 data

We now present the first application of a multi-stage MCMC-based hierarchical follow-up on real data by studying a set of 30 outliers obtained by different CW searches on Advanced LIGO O2 data. These are final-stage outliers resulting from the application of a complete search pipeline, including a set of vetoes depending upon the particularities of each search.

Section 4.4.1 briefly describes the main traits of the searches from which outliers are collected. The complete set of outliers to be followed up is reported in Table 4.2. The follow-up setup is described in Sec. 4.4.2.

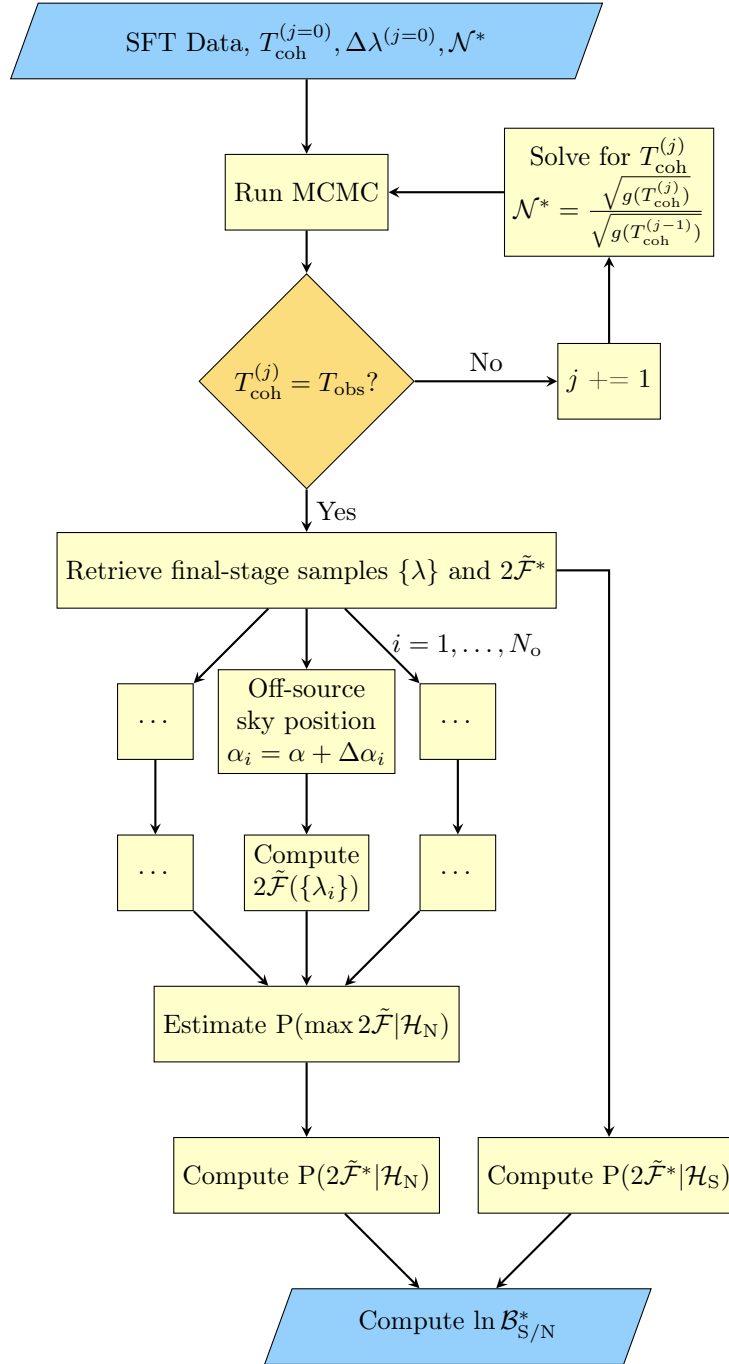
### 4.4.1 Continuous-wave search outliers from O2 data

#### All-sky Falcon search

The Falcon pipeline [13] is designed to survey wide parameter-space regions using a so-called *loosely coherent* approach [54, 78, 79], increasing its robustness against small deviations from the standard CW signal model [54].

We are interested in 18 outliers reported in two all-sky searches targeting two different frequency bands of the Advanced LIGO O2 dataset: mid frequency (500 – 1700 Hz) [80] and high frequency (1700 – 2000 Hz) [81]. These searches intended to unveil unknown low-ellipticity sources by analyzing a restricted set of spindown rates ( $|f_1| \lesssim 3 \cdot 10^{-12}$  Hz/s). These outliers are the result of a four-stage search using four different coherent times, namely 12, 24, 48 and 144 hours. After each stage, only those templates over a specified threshold were further followed up.

An additional low-frequency Falcon search was recently reported in [82]. As will be shown in Sec. 4.5.2 regarding low-frequency outliers from the other searches discussed below, the greater number and variety of instrumental artifacts in the low-frequency data somewhat hinder the effectiveness of this first incarnation of our follow-up method, as they are not directly addressed by the noise hypothesis. Therefore, we leave a re-analysis of the new low-frequency Falcon outliers for future work.

FIGURE 4.4: Flowchart illustrating the computation of  $\ln \mathcal{B}_{S/N}^*$  for a CW outlier.

### Directed Einstein@Home search

**Einstein@Home** is a large-scale computing framework based on the volunteer-computing platform BOINC [83] on which the Global Correlations Transform pipeline [84, 85, 86], intended to perform deep (very sensitive) searches across wide parameter-space regions, is deployed. This pipeline is flexible enough so as to be reconfigured into a directed pipeline, using astrophysical information obtained by electromagnetic means to restrict the sky positions to search on.

We are interested in the surviving outlier from a directed search for CWs from central compact objects in three supernova remnants [64]. Said outlier is associated to the central compact object known as 1 WGA J1713.4-949 [87] and located in SNR G347.3-05; for consistency with [64], we will simply refer to it as J1713.

This outlier is a sub-threshold candidate from an earlier **Einstein@Home** search on O1 data directed towards the same supernova remnants [29], which was then re-analyzed using O2 data. The statistical basis of the re-analysis was similar to the techniques explained in Sec. 4.3.2, comparing the significance of a candidate on different data streams with respect to the expected significance deduced from the initial analysis. As reported in [64], the outlier under analysis is inconsistent with Gaussian noise, but cannot be associated to the signal hypothesis either.

### Fomalhaut b Viterbi search

The **Viterbi** method spans a family of search pipelines which use a Hidden Markov model (HMM) to describe the frequency evolution of a CW signal [88, 71, 89, 90]. Such a signal model is able to incorporate stochastic contributions into the analysis (e.g. timing noise or spin-wandering due to the presence of an accreting companion [91]).

Ref. [92] reports on a **Viterbi** search for CWs directed at Fomalhaut b, an astrophysical object whose exact nature is still surrounded by debate [93, 94, 95, 96]. This search complements a previous one performed on Advanced LIGO O1 data using an  $\mathcal{F}$ -statistic search assuming the standard deterministic evolution of a CW [97].

The search setup assumes spindown to be the main contribution to the frequency evolution, considering timing noise as a sub-dominant component. This is done by imposing a *biased* random walk as a HMM, in the sense that evolution towards higher frequencies is forbidden<sup>3</sup>. The search was performed using  $T_{\text{coh}} = 5$  days and surviving candidates were sieved through a set of consistency vetoes. In the end, a single outlier was reported for further exploration.

### H.E.S.S. Viterbi search

Another implementation of the **Viterbi** pipeline, similar in scope and assumptions to that mentioned above, was used to perform a search on a set of ten pulsars observed by very high-energy  $\gamma$ -ray surveys in [98].

The search looks for CW emission at once, twice and  $4/3$  of the rotational frequency of the targeted pulsars in order to address several emission mechanisms [4]. After assessing the sub-dominant role of spin-wandering on frequency evolution, a biased random walk is implemented in a similar manner to [92], selecting the maximum  $T_{\text{coh}}$  allowed by the spindown rate of each pulsar so that the frequency evolution is within the range of the HMM.

After applying a set of consistency vetoes, twelve outliers are reported for further exploration; we only considered ten of them as independent follow-up targets since for two pairs of outliers, the corresponding prior parameter-space regions significantly overlap.

## 4.4.2 Follow-up setup

We demonstrate the general application of an MCMC-based multi-stage follow-up to a set of real-data outliers regardless of the pipeline producing them. To do so, outliers will be analyzed ignoring any information gathered from any of the vetoes or follow-up stages reported in their respective searches.

The second Advanced LIGO observing run [99, 100] comprises nine months of data taken by the two Advanced LIGO detectors H1 (Hanford) and L1 (Louisiana) [2]. The employed time segments are those

<sup>3</sup>This condition drastically reduces the space of possible frequency evolutions contemplated by the HMM model, easing the application of a model-based pipeline to follow up or estimate the exact parameters of any resulting candidates.

Outlier ID	Search	$f_0$ [Hz]	$f_1$ [Hz/s]	$f_2$ [Hz/s <sup>2</sup> ]	$\alpha$ [rad]	$\delta$ [rad]	$t_{\text{ref}}$ [GPS]	Ref.
Falcon 4	High-frequency Falcon	1891.756740	$-8.22 \cdot 10^{-12}$	—	2.986956	1.005798	1183375935	[81]
Falcon 5	High-frequency Falcon	1892.991060	$-1.08 \cdot 10^{-12}$	—	3.779161	-0.816273	1183375935	[81]
Falcon 15	Mid-frequency Falcon	900.218805	$-2.20 \cdot 10^{-12}$	—	2.084418	-0.102264	1183375935	[80]
Falcon 19	Mid-frequency Falcon	514.148927	$1.60 \cdot 10^{-12}$	—	2.170421	0.092501	1183375935	[80]
Falcon 23	Mid-frequency Falcon	1001.366228	$4.30 \cdot 10^{-12}$	—	1.355837	-0.770266	1183375935	[80]
Falcon 24	Mid-frequency Falcon	676.195421	$2.80 \cdot 10^{-12}$	—	3.847021	-0.101619	1183375935	[80]
Falcon 25	Mid-frequency Falcon	744.219166	$2.40 \cdot 10^{-12}$	—	3.344985	0.612566	1183375935	[80]
Falcon 29	Mid-frequency Falcon	512.490814	$1.20 \cdot 10^{-12}$	—	2.468975	-0.043050	1183375935	[80]
Falcon 31	Mid-frequency Falcon	983.151889	$2.20 \cdot 10^{-12}$	—	3.561119	0.017979	1183375935	[80]
Falcon 34	Mid-frequency Falcon	886.880087	$-1.60 \cdot 10^{-12}$	—	4.912788	-0.703498	1183375935	[80]
Falcon 35	Mid-frequency Falcon	988.373199	$1.20 \cdot 10^{-12}$	—	0.981835	0.778338	1183375935	[80]
Falcon 39	Mid-frequency Falcon	514.291681	$3.20 \cdot 10^{-12}$	—	0.569033	-0.128357	1183375935	[80]
Falcon 40	Mid-frequency Falcon	831.988473	$4.00 \cdot 10^{-13}$	—	4.917347	1.160537	1183375935	[80]
Falcon 41	Mid-frequency Falcon	873.524608	$4.00 \cdot 10^{-13}$	—	0.618991	-0.189450	1183375935	[80]
Falcon 42	Mid-frequency Falcon	895.421949	$3.60 \cdot 10^{-12}$	—	5.105590	0.249163	1183375935	[80]
Falcon 43	Mid-frequency Falcon	1224.745666	$-2.16 \cdot 10^{-12}$	—	1.715268	0.196184	1183375935	[80]
Falcon 45	Mid-frequency Falcon	698.728032	$-2.00 \cdot 10^{-13}$	—	4.557347	-0.724141	1183375935	[80]
Falcon 46	Mid-frequency Falcon	1095.557400	$-1.08 \cdot 10^{-12}$	—	4.354664	-0.260254	1183375935	[80]
J1713	Einstein@Home	368.801379	$-4.37 \cdot 10^{-9}$	$5.9 \cdot 10^{-19}$	4.509371	-0.695189	1131943508	[64]
Fomalhaut b	Fomalhaut b Viterbi	876.503400	$-1.00 \cdot 10^{-12}$	—	6.011130	0.517000	1167545066	[92]
J0534+2200	H.E.S.S. Viterbi	29.813738	$-3.77 \cdot 10^{-10}$	—	1.459675	0.384225	1164556817	[98]
J1420-6048	H.E.S.S. Viterbi	14.511294	$-1.70 \cdot 10^{-11}$	—	3.753057	-1.061240	1164556817	[98]
J1420-6048	H.E.S.S. Viterbi	19.515033	$-2.30 \cdot 10^{-11}$	—	3.753057	-1.061240	1164556817	[98]
J1420-6048	H.E.S.S. Viterbi	29.522611	$-3.50 \cdot 10^{-11}$	—	3.753057	-1.061240	1164556817	[98]
J1718-3825	H.E.S.S. Viterbi	17.503470	$-3.00 \cdot 10^{-12}$	—	4.530116	-0.670585	1164556817	[98]
J1831-0952	H.E.S.S. Viterbi	14.501823	$-1.00 \cdot 10^{-12}$	—	4.850147	-0.172213	1164556817	[98]
J1831-0952	H.E.S.S. Viterbi	15.401223	$-1.00 \cdot 10^{-12}$	—	4.850147	-0.172213	1164556817	[98]
J1831-0952	H.E.S.S. Viterbi	19.999146	$-2.00 \cdot 10^{-12}$	—	4.850147	-0.172213	1164556817	[98]
J1849-0001	H.E.S.S. Viterbi	26.308209	$-9.00 \cdot 10^{-12}$	—	4.850147	-0.000375	1164556817	[98]
J1849-0001	H.E.S.S. Viterbi	26.341209	$-9.00 \cdot 10^{-12}$	—	4.850147	-0.000375	1164556817	[98]

TABLE 4.2: CW search outliers of interest as reported by their original searches. H.E.S.S. Viterbi outliers will be further referred to by including their corresponding frequency.

Search	Estimated $\mathcal{D}^{95\%}$ [Hz <sup>-1/2</sup> ]
High-frequency Falcon	55 — 65
Mid-frequency Falcon	45 — 55
Directed Einstein@Home*	80 — 90 (75 — 85)
Fomalhaut b Viterbi	45 — 55
H.E.S.S. Viterbi	45 — 55

TABLE 4.3: Estimated ranges of 95% efficiency sensitivity depths achieved by each of the searches according to their reported results. The depth marked with an asterisk corresponds to a 90% efficiency instead. Values in parentheses refer to the sensitivity depth achieved by the original search producing the outlier [29].

Stage	0	1	2	3	4
$N_{\text{seg}}$	500	250	55	5	1
$T_{\text{coh}}$ [days]	0.5	1	5	55	270

TABLE 4.4: Coherence-time ladder constructed using  $\mathcal{N}^* = 10^4$  and including an initial stage of  $T_{\text{coh}} = 0.5$  days before imposing  $T_{\text{coh}} = 1$  days and applying the SuperSky metric. The results are independent of the parameter-space region at which the SuperSky metric was evaluated.

Hyperparameter	Value
Parallel chains	3
Walkers per chain	100
Burn-in & Production steps	250 + 250

TABLE 4.5: MCMC hyperparameter choices for each stage of the follow-up. The number of parallel chains equals the number of temperatures at which the likelihood is being sampled, following the recommendations in [58, 32].

with the “all” tag in [101]. The dataset was divided into segments with a duration of  $T_{\text{SFT}} = 1800$  s in which Fourier transforms were computed as explained in [20]. We take the observing time to be  $T_{\text{obs}} = 270$  days in order to convert the number of segments of a stage  $N_{\text{seg}}$  to a coherence time as  $T_{\text{coh}} = T_{\text{obs}}/N_{\text{seg}}$ .

Our follow-ups are conducted assuming a CW signal model with two spindown components. Since the second spindown component is only reported by the `Einstein@Home` search, we assume it to be compatible with a null value for the other outliers and apply a canonical uncertainty of  $\delta f_2 = 2 \cdot (T_{\text{coh}} \cdot T_{\text{obs}}^2)^{-1}$  [9]. As discussed in [102], this increases the robustness of a search method against unmodeled physics, such as neutron star glitches, due to an increase of the available parameter-space correlations.

Table 4.3 collects the approximated sensitivity depth achieved by each search according to their reported results. A comparison to the results in Figs. 8 and 9 of [32], which compute the detection efficiency of a four-stage MCMC follow-up starting at  $T_{\text{coh}} = 1$  day, places the outliers within the effective region of the follow-up procedure.

Most wide parameter-space searches currently operate at  $T_{\text{coh}} \sim \mathcal{O}(\text{hours})$ . As demonstrated in [30, 21], CW candidates with uncertainties at such short coherence times can be successfully recovered by an MCMC follow-up at  $T_{\text{coh}} = 0.5$  days.

We construct a hierarchical follow-up by imposing a first stage using  $T_{\text{coh}} = 0.5$  days followed by a second stage using  $T_{\text{coh}} = 1$  day. Further stages are constructed by means of Eq. (4.14) using  $\mathcal{N}^* = 10^4$ . The resulting coherence-time ladder, which is independent of the parameter-space region and the prior specification due to the locality of the analysis, is collected in Table 4.4. As per the previous discussion on the sensitivity of the considered searches, this ladder can be seamlessly applied to every one of the outliers under analysis.

Table 6.7 specifies the hyperparameter setup of every MCMC stage, following the setups employed in [30, 21]. As demonstrated in Sec. 4.5.1, this setup suffices to successfully follow up CW candidates within the probed sensitivity range.

The choice of initial priors is directly related to the outlier’s uncertainty returned by each of the analysis pipelines. Pipelines like `Falcon` or `Einstein@Home` return a well-determined parameter-space region in which the outlier was found. The `Viterbi` pipelines, on the other hand, return only the frequency-evolution track of each candidate, which can then be related to a certain parameter-space region if the stochastic contributions are sub-dominant. The scope of a search also affects the prior setup, as searches directed towards a particular sky position (such as the ones performed using `Viterbi`) allow us to place a narrower prior on the sky position of the outlier. It is recommended in [32] to choose a flat prior with fixed bounds containing the outlier’s parameters. Instead, we use a set of Gaussian priors centered at the outlier’s parameters with scale parameters corresponding to the uncertainty in each dimension. After each MCMC step, we re-center the priors on the median value of the resulting

posterior distribution, taking half of the (centered) 90% credible interval as the new scale parameter, and re-sample the initial state of the MCMC ensemble. This particular setup ensures a fresh start-up at each stage of the ladder, preventing spurious samples dissociated from the ensemble to pollute the final results. Moreover, the use of unbounded priors prevents the follow-up from missing the true parameters of an outlier due to the presence of parameter-space correlations [103, 47].

The uncertainty associated to Falcon outliers is specified in [80] as

$$\begin{aligned}\delta f_0 &= 5 \times 10^{-5} \text{ Hz} , \\ \delta f_1 &= 1 \times 10^{-12} \text{ Hz/s} , \\ \delta\theta &= 0.06 \text{ Hz}/f_0 \text{ rad} ,\end{aligned}\tag{4.29}$$

where  $\delta\theta$  refers to the sky position of an outlier projected onto the ecliptic plane. These uncertainties are conservatively lower than the canonical parameter-space resolution defined in [9] for a coherence time of  $T_{\text{coh}} = 0.5$  days, meaning their corresponding parameter-space size is within acceptable values to ensure an effective MCMC stage [21, 30]. The Einstein@Home search reports uncertainties corresponding to a coherence time of several months; since we start our follow-up at a lower coherence time, we used the same set of uncertainties as for the Falcon follow-up Eq. (4.29). Viterbi outliers were not reported as a parameter-space point, but as a frequency band on which a significant frequency-evolution track was found; since both searches were targeted at a particular sky position, we reduced the sky position uncertainty and increased the frequency uncertainty by the same factor in order to cover all possible frequencies at a similar parameter-space size.

## 4.5 Results

Before presenting results on the O2 outliers in Sec. 4.5.2, here we first describe an injection campaign in simulated Gaussian noise to demonstrate the efficacy of the follow-up procedure and calibrate a threshold on the newly introduced Bayes factor.

### 4.5.1 Injections in Gaussian noise

We characterize the behavior of  $\ln \mathcal{B}_{\text{S/N}}^*$  using three sets of 100 artificial signals at different signal strengths. These are injected into Gaussian noise data compatible with the O2 observing run characteristics, i.e. simulating data for both Advanced LIGO detectors and with a duration of  $T_{\text{obs}} = 9$  months, using `lalapps_Makefakedata_v5` [104]. The actual O2 data stream covers 60% of the duration of the run  $T_{\text{obs}}$  due to down time in the detectors (actual fractions are 65.3% and 61.8% for the H1 and L1, respectively) [99]. Since SNR scales as the square root of observing time, this would reduce the actual SNR of a signal to a fraction of 77%. For the simulated Gaussian noise, we set the average amplitude spectral density to a fiducial value of  $\sqrt{S_n} = 10^{-23} \text{ Hz}^{-1/2}$ . We injected the artificial signals at a fiducial frequency of 100 Hz, uniformly spread across the whole sky and log-uniformly distributed in spindown parameter  $f_1$  within  $[-10^{-8}, -10^{-11}] \text{ Hz/s}$ . The particular choice of a frequency band does not affect the results of this analysis, since its effects are automatically taken into account by parameter-space resolutions.

The CW amplitude  $h_0$  is fixed in terms of the sensitivity depth [16, 53]

$$\mathcal{D} = \frac{\sqrt{S_n}}{h_0} .\tag{4.30}$$

Additionally, we define an *effective* sensitivity depth by explicitly including the effects of the cosine of the inclination angle  $\iota$  [105]:

$$\mathcal{D}_{\text{Eff}} = \frac{\mathcal{D}}{\sqrt{\cos^4 \iota + 6 \cos^2 \iota + 1}} .\tag{4.31}$$

We selected three depth values, enumerated in Table 4.6, bracketing the estimated 95% efficiency depth of the analyzed pipelines. The rest of the amplitude parameters were randomly drawn from uniform distributions [106].



Depth [Hz <sup>-1/2</sup> ]	Efficiency (%)
40	97 ± 2
60	98 ± 1
80	96 ± 2
Overall	97 ± 1

TABLE 4.6: Detection efficiencies for each set of 100 injections. An injection was labeled as detected if the final-stage posterior probability contained the injection parameters in its support. Error bars correspond to binomial errors.

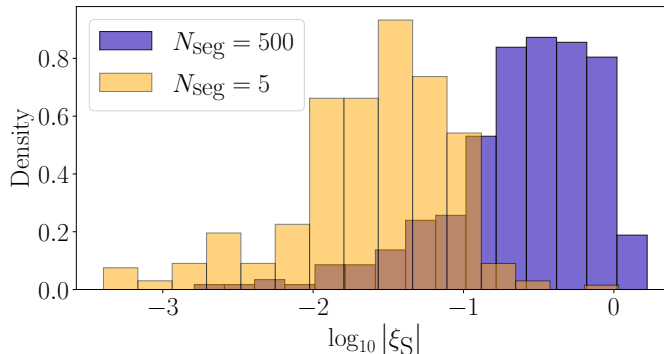


FIGURE 4.5: Distribution of  $|\xi_S|$  for the complete set of detected injections using different semicoherent stages, namely  $N_{\text{seg}} = 500$  and  $N_{\text{seg}} = 5$ , as the reference to compute  $\mu_S$  and  $\sigma_S$ .

We start by estimating the detection efficiency of the follow-up. To do so, we run the full hierarchical follow-up as specified in the previous section and count an injection as “detected” if the injection parameters are within the final-stage posterior probability support. This criterion ensures the CW signals are strong enough to guide the MCMC ensemble towards the relevant parameter space region, preventing a signal from being lost. Results are reported in Table 4.6. As expected from previous analyses in [32], we obtain a detection efficiency above 95% across the sensitivity range, meaning the follow-up is a suitable tool to further analyze the selected set of outliers.

The computation of  $\ln \mathcal{B}_{S/N}^*$  requires a particular semicoherent step from the ladder to be selected as the one from which the expected fully-coherent distribution will be propagated. As discussed in Sec. 4.3.2, using longer coherence times imposes a more restrictive signal model, reducing the number of outliers due to the presence of detector artifacts and increasing the significance of signal candidates (see e.g. Figure 6 of [32]). Fig. 4.5 shows the obtained distribution of signal-hypothesis discrepancies  $|\xi_S|$  for the complete set of detected injections with respect to two different stages. The use of a lower number of segments (i.e. a longer coherence time) yields a tighter consistency with respect to the expected distribution. We decide to carry out the analysis by taking the second-to-last stage of the ladder ( $N_{\text{seg}} = 5$ ) as the reference from which the expected fully-coherent  $\mathcal{F}$ -statistic distribution will be computed.

Figure 4.6 displays the distribution of injection results on the  $(\xi_N, \xi_S)$  plane, showing the discrepancy of an outlier with respect to the noise and signal hypotheses, respectively. The configuration is such that  $\xi_S \sim 0$  and  $\xi_N \gg \xi_S$ , corresponding to case b) in Sec. 4.3. This means that the computation of the signal contribution to  $\ln \mathcal{B}_{S/N}^*$  can be assumed to follow a Gaussian distribution and, correspondingly, Eq. (4.27) applies. Figure 4.7 shows the Bayes factor  $\ln \mathcal{B}_{S/N}^*$  computed by comparing the last two stages of the semicoherent ladder. The observed behavior  $\ln \mathcal{B}_{S/N}^* \propto \mathcal{D}_{\text{Eff}}^{-1}$  can be simply explained by noting that  $\mathcal{D}_{\text{Eff}}$  is inversely proportional to SNR by definition [53].

This injection campaign covers the sensitivity ranges reported in Table 4.3 for all searches except one. The `Einstein@Home` search differs in that it was built as a subthreshold search: the reported outlier was




FIGURE 4.6:  $(\xi_S, \xi_N)$  plane for the complete set of detected injections using  $N_{\text{seg}} = 5$  as the reference stage to compute  $\mu_S$  and  $\sigma_S$ . The horizontal axis represents the discrepancy with respect to the noise hypothesis, while the vertical axis represents the discrepancy with respect to the signal hypothesis.




FIGURE 4.7:  $\ln \mathcal{B}_{S/N}^*$  computed by applying the multi-stage MCMC follow up on the three sets of software injections. Reference values were computed with respect to the  $N_{\text{seg}} = 5$  stage and the Gaussian approximation was used to compute the signal contribution. Outliers marked by a star did not display an ensemble-level volume shrinkage, as explained in the text.




FIGURE 4.8: Estimation of 90% detection-probability threshold on  $\ln \mathcal{B}_{S/N}^*$  for different sensitivity depths beyond the injections shown in Fig. 4.7. Each dot represents an empirical estimate of the 90% detection-probability threshold using 100 simulated signals at a fixed depth value. Error bars correspond to the bootstrap standard deviation using 200 resamples of 50 samples each. The solid line and the associated envelopes represent a linear fit using `scipy.optimize.curve_fit` [107] with 1, 2, and 3 sigma uncertainties.

thoroughly scrutinized using a variety of tools, including a fully-coherent analysis on O2 data, following a similar scheme as the one presented in this work. In order to assess the follow-up capabilities of our proposed method, we perform a second injection campaign akin to the previous one, covering the deepest **Einstein@Home** search sensitivity range. Results are reported as 90% detection-probability thresholds in Fig. 4.8, following the approach proposed in [53]. Were any of the considered outliers due to a genuine CW signal, the corresponding  $\ln \mathcal{B}_{S/N}^*$  should lie within the shaded region or higher. Based on this argument, we set a safe decision threshold at  $\ln \mathcal{B}_{S/N}^* = 30$ , also accounting for the reduced SNR in the real dataset due to detector downtime.

Lastly, we comment on the behavior of the multi-stage MCMC itself in terms of the volume shrinkage rate introduced in Sec. 4.2.3. Figure 4.9 shows the behavior of the posterior volume of a successfully detected injection. The quantities  $V_{\text{prior}}^{(0)}$  and  $V_{\text{post}}$  represent approximations to the initial prior volume at the first stage of the ladder and the posterior volume after each of the MCMC stages. These quantities are computed by taking the product of parameter-wise central 90% credible intervals, since we are only interested in the overall scaling along the coherence-time ladder. The volume shrinkage shows a power-law behavior, the exponent of which (i.e. the slope in log-log scale) should be approximately given by  $\log_{10} \gamma^{(j+1)} \sim 4$  from Eq. (4.12). The same procedure is performed on the complete set of detected injections, collecting the power-law indices into a histogram in Fig. 4.10. The rate of volume shrinkage accumulates a prominent peak within the order of magnitude of the expected result.

Figure 4.10 also displays a small set of injections for which the MCMC ensemble did not produce a clear shrinkage of the (approximated) central 90% credible region, even though the true injection parameters are contained within said region. Their corresponding  $\ln \mathcal{B}_{S/N}^*$  values are marked using stars in Fig. 4.7, belonging to the weakest set of performed injections. This is a consequence of the parameter-space structure in the vicinity of a signal [47]: The effective (squared) SNR recovered by a template falls off as a linear function of the mismatch with respect to the true signal parameters. Strong injections, associated to higher SNR values, are able to sustain an  $\mathcal{F}$ -statistic above background throughout a wider parameter-space region than weak injections. Weak injections, as a result, require tighter priors to display a similar behavior to that of stronger injections. The fact that the ensemble is unable to focus into a particular parameter-space region, however, is still compatible with a good recovery of  $\ln \mathcal{B}_{S/N}^*$ , as for that it is only required to *sample* the region of interest during the production stage. This is in fact the principle upon which the application of a single-stage MCMC follow-up as a simple veto was based in [30, 21], and can be justified by interpreting the MCMC follow-up as being equivalent to a search starting from a *random* template bank at higher mismatches than traditionally suggested in CW searches [108, 109].




FIGURE 4.9: Posterior volume shrinkage of a successfully detected software injection. Parameter-space volumes are estimated by taking the product of parameter-wise 90% central credible regions as explained in the text. The vertical axis represents the posterior volume as a fraction of the initial prior volume. The slope of the log-log plot is an approximation to the inverse of the volume shrinkage  $1/\gamma^{(j+1)}$ , where the volume shrinkage  $\gamma^{(j+1)}$  was defined in Eq. (4.12).




FIGURE 4.10: Distribution of  $(\log_{10})$  posterior volume shrinkage rates of the detected injections. The prominent peak is within the order of magnitude of the expected value according to Eq. (4.12). The presence of *negative* shrinkage rates for a few injections is discussed in the text.

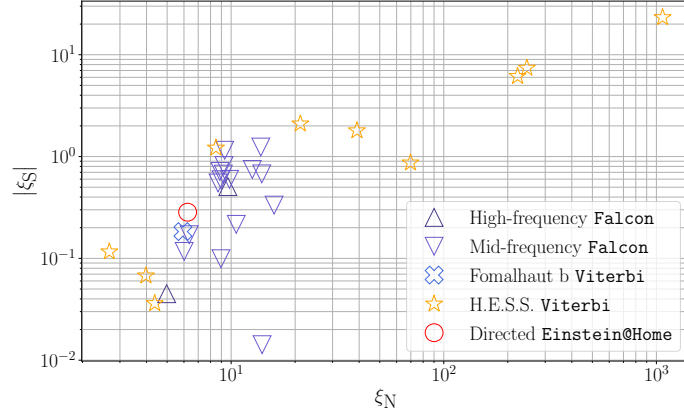


FIGURE 4.11:  $(\xi_N, \xi_S)$  plane associated to the outliers found in O2 data by the specified searches.

### 4.5.2 Follow-up of CW outliers from Advanced LIGO O2 data

We now report the results of the multi-stage MCMC-based follow-up on the set of outliers described in Sec. 4.4.1 in terms of the obtained  $(\xi_N, \xi_S)$  values and the corresponding  $\ln \mathcal{B}_{S/N}^*$ .

Figure 4.11 shows each pipeline’s outliers across the  $(\xi_N, \xi_S)$  plane, quantifying their discrepancy with respect to the noise and signal hypotheses. The bulk of outliers show discrepancies with respect to the signal hypothesis, quantified by  $|\xi_S|$ , an order of magnitude larger than those displayed by software injections in Fig. 4.6. Discrepancies with respect to background noise, quantified by  $\xi_N$ , are more than an order of magnitude lower. The retrieved values of  $\mu_N$  and  $\sigma_N$  are within the brackets obtained in Gaussian noise, suggesting these results are not because of an elevated background noise but rather a low SNR associated to the outliers. We note the presence of three *Viterbi* outliers at high values of  $\xi_N$ , namely J1831-0952@19.9991Hz, J1849-0001@26.3410Hz and J1831-0952@15.4012Hz; and another marginal pair stemming from the same pipeline in the middle ground, namely J1718-3825@17.5034Hz and J1831-0952@14.5018Hz.

We compute  $\ln \mathcal{B}_{S/N}^*$  by numerically integrating Eq. (4.19) due to the regime in which outliers are placed. Results are listed in Table 4.7 and displayed in Fig. 4.12. Five outliers score over the decision threshold  $\ln \mathcal{B}_{S/N}^* = 30$ , all of them related to the H.E.S.S. *Viterbi* pipeline.

The first set of outliers, J1831-0952@15.4012Hz, J1831-0952@19.9991Hz, and J1849-0001@26.3410Hz, is highlighted using circular markers in Fig. 4.12. The original search [98] ascribed them to instrumental artifacts in the L1 detector. We confirm that to be the case for the outlier J1831-0952@19.9991Hz: the loudest fully-coherent  $\mathcal{F}$ -statistic recovered by our follow-up is located at  $f_0 \simeq 20.0011$  Hz, crossing a well-known instrumental comb at both LIGO detectors [110]. For outlier J1849-0001@26.3410Hz, we note the presence of a hardware injection (a CW-like signal simulated by direct actuation of the interferometer mirrors, used to test calibration and analysis pipelines) at  $f_0 \simeq 26.3396$  Hz with an amplitude corresponding to  $\mathcal{D} \sim \mathcal{O}(1 \text{ Hz}^{-1/2})$  [111, 100]. Even though the spindown and sky positions are completely mismatched, such strong artificial signals are known to produce loud candidates across wide parameter space regions [112, 28, 113, 114, 30]. We are unable to relate J1831-0952@15.4012Hz to any of the listed narrow spectral artifacts in [110, 100].

A manual check of the segment-wise semicoherent  $\mathcal{F}$ -statistic values of J1831-0952@15.4012Hz reveals a rapid accumulation of  $\mathcal{F}$ -statistic as the frequency evolution crosses a narrow sub-band. This kind of behavior, shown in Fig. 4.13, is inconsistent with a CW signal and usually can be related to instrumental artifacts, but said identification becomes more difficult at low frequencies as they are populated by a wider variety of noise sources. As a result, outlier J1831-0952@15.4012Hz is also likely related to an instrumental artifact.

The second group of outliers is enclosed by diamonds in Fig. 4.12. Outlier J1831-0952@14.5018Hz’s loudest candidate is recovered at  $f_0 \simeq 14.4953$  Hz. This is consistent with a 1 Hz comb with an offset of




FIGURE 4.12:  $\ln \mathcal{B}_{S/N}^*$  values obtained after the hierarchical MCMC follow-up of O2 outliers. Relative outlier positions in this figure are consistent with Fig. 4.11. Outliers enclosed by circles and diamonds score a  $\ln \mathcal{B}_{S/N}^*$  value above 30. The outlier enclosed by a square returns a negative value of  $\ln \mathcal{B}_{S/N}^*$  and is displayed as white due to the logarithmic color scale.

0.5 Hz.<sup>4</sup> Another harmonic of the same comb can be related to outlier J1718-3825@17.5034Hz, whose loudest candidate is located at  $f_0 \simeq 17.5005$  Hz.

The remaining outliers from all searches return a  $\ln \mathcal{B}_{S/N}^*$  value below the decision threshold  $\ln \mathcal{B}_{S/N}^* = 30$ . For completeness, we list the parameters recovered by the final follow-up stage in Table 4.7. We highlight outlier J1713, initially found by the Einstein@Home search and scoring below our decision threshold. This result is consistent with the latest Einstein@Home search for J1713 reported in [115], covering up to 400 Hz in O2 data, not finding any significant outliers.

## 4.6 Conclusion

We have introduced the first complete framework to analyze outliers from arbitrary CW searches using a multi-stage MCMC-based follow-up. After demonstrating its general behavior on Gaussian noise, we applied it to a set of 30 outliers obtained by different CW search pipelines on O2 Advanced LIGO data [80, 81, 92, 98, 64].

The procedure constructs a Bayes factor comparing whether the behavior of the  $\mathcal{F}$ -statistic across different stages of the analysis is more consistent with the presence of a signal rather than with pure noise. The expected evolution of this detection statistic as the follow-up progresses can be derived from first principles. The noise contribution is described by applying extreme value theory to samples of background noise data. These samples can be obtained by sampling shifted sky positions with respect to the outliers, blinding the analysis from the presence of a signal.

The application of a multi-stage MCMC follow-up deemed 25 of the analyzed outliers as less consistent with a standard CW signal than with background noise. The remaining five outliers passed the specified threshold and were manually inspected. Four of them were successfully associated to known instrumental artifacts in the Advanced LIGO detectors. The fifth outlier displays a behavior inconsistent with a CW signal but consistent with an instrumental artifact; the exact instrumental cause, however, could not be identified.

Although the outliers were analyzed assuming a standard signal model corresponding to an isolated CW source, the framework presented here (and the PyFstat software used [32, 34, 59]) can be seamlessly applied to more general models, such as sources in binary systems [30, 21], sources producing glitches [35], and long gravitational-wave transient signals [36, 37].

<sup>4</sup>The spindown value reported by the original search is such that also positive values are covered by the initial prior volume, and indeed our followup recovered the loudest candidate at positive spindown.

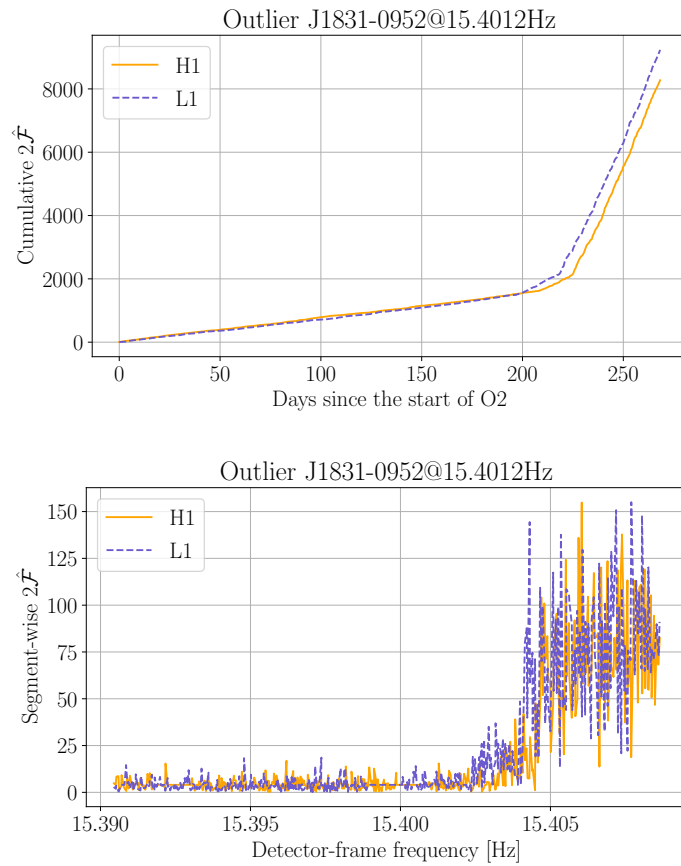


FIGURE 4.13: Segment-wise  $2\hat{\mathcal{F}}$  accumulation of the loudest template associated to the outlier J1831-0952@15.4012Hz throughout the observing run using 500 coherent segments. The solid orange line shows the results using the LIGO Hanford detector (H1) only, while the dashed blue line shows the results using the LIGO Livingston detector (L1). The upper panel shows the  $2\hat{\mathcal{F}}$  accumulation throughout the duration of the run. The lower panel shows the segment-wise  $2\hat{\mathcal{F}}$  values per frequency bin. Frequency values are computed evaluating the frequency-evolution template at the starting time of each segment.

Outlier ID	$f_0$ [Hz]	$f_1$ [Hz/s]	$f_2$ [Hz/s <sup>2</sup> ]	$\alpha$ [rad]	$\delta$ [rad]	$\ln \mathcal{B}_{S/N}^*$
Falcon 4	1891.756615	$-6 \cdot 10^{-12}$	$1 \cdot 10^{-20}$	2.987285	1.005941	2.10
Falcon 5	1892.991046	$-2 \cdot 10^{-12}$	$-1.8 \cdot 10^{-19}$	3.778973	-0.816265	6.46
Falcon 15	900.218764	$-1 \cdot 10^{-12}$	$7 \cdot 10^{-20}$	2.084412	-0.102330	10.49
Falcon 19	514.148984	$6 \cdot 10^{-12}$	$3.7 \cdot 10^{-19}$	2.170669	0.092978	9.04
Falcon 23	1001.366278	$2 \cdot 10^{-12}$	$4.9 \cdot 10^{-19}$	1.355553	-0.769952	5.88
Falcon 24	676.195493	$3 \cdot 10^{-12}$	$-2.6 \cdot 10^{-19}$	3.846438	-0.102138	5.80
Falcon 25	744.219196	$2 \cdot 10^{-12}$	$0.4 \cdot 10^{-20}$	3.344781	0.612270	7.42
Falcon 29	512.490782	$-8 \cdot 10^{-12}$	$-2 \cdot 10^{-20}$	2.468688	-0.041880	3.37
Falcon 31	983.151151	$-1 \cdot 10^{-12}$	$-7.0 \cdot 10^{-19}$	3.562362	0.018926	3.06
Falcon 34	886.880063	$-2 \cdot 10^{-12}$	$-2.1 \cdot 10^{-19}$	4.912748	-0.703663	10.85
Falcon 35	988.373241	$2 \cdot 10^{-12}$	$-1.3 \cdot 10^{-19}$	0.982043	0.778393	6.48
Falcon 39	514.291753	$3 \cdot 10^{-12}$	$-3.5 \cdot 10^{-19}$	0.569150	-0.128791	5.47
Falcon 40	831.988457	$-3 \cdot 10^{-12}$	$-1 \cdot 10^{-20}$	4.917884	1.160566	5.48
Falcon 41	873.524663	$3 \cdot 10^{-12}$	$4 \cdot 10^{-20}$	0.619107	-0.189295	5.77
Falcon 42	895.421995	$1 \cdot 10^{-12}$	$-1.9 \cdot 10^{-19}$	5.105728	0.249030	5.37
Falcon 43	1224.745693	$1 \cdot 10^{-12}$	$-1.2 \cdot 10^{-19}$	1.715372	0.196097	5.70
Falcon 45	698.728033	$1 \cdot 10^{-12}$	$1.6 \cdot 10^{-19}$	4.557448	-0.723930	12.69
Falcon 46	1095.557373	$-4 \cdot 10^{-12}$	$-5.7 \cdot 10^{-19}$	4.354405	-0.260292	9.60
J1713	368.801590	$-4.380 \cdot 10^{-9}$	$1.18 \cdot 10^{-18}$	4.511570	-0.694137	3.21
Fomalhaut b	876.517914	$-4.2979 \cdot 10^{-10}$	$-5.67 \cdot 10^{-18}$	6.011153	0.516952	2.96
J0534+2200	29.813469	$-2.3430 \cdot 10^{-10}$	$1.158 \cdot 10^{-17}$	1.461040	0.385286	-0.16
J1420-6048	14.511112	$-2.5 \cdot 10^{-11}$	$1.364 \cdot 10^{-17}$	3.750570	-1.061001	1.46
J1420-6048	19.512364	$-4.4 \cdot 10^{-11}$	$9.80 \cdot 10^{-18}$	3.753242	-1.061227	1.12
J1420-6048	29.526774	$3.9 \cdot 10^{-11}$	$5.30 \cdot 10^{-18}$	3.753011	-1.060915	4.46
J1718-3825	17.500500	$-4.0 \cdot 10^{-11}$	$7.58 \cdot 10^{-18}$	4.528719	-0.670563	<i>33.45</i>
J1831-0952	14.495361	$2.95 \cdot 10^{-10}$	$-3.89 \cdot 10^{-18}$	4.848071	-0.172356	<i>65.64</i>
J1831-0952	15.389002	$8.72 \cdot 10^{-10}$	$-2.980 \cdot 10^{-17}$	4.853052	-0.165697	<b>203.35</b>
J1831-0952	20.0016854	$-7.13 \cdot 10^{-10}$	$7.017 \cdot 10^{-17}$	4.859380	-0.179332	<b>633.204</b>
J1849-0001	26.3062476	$-6.4 \cdot 10^{-11}$	$1.101 \cdot 10^{-17}$	4.850380	-0.000331	14.71
J1849-0001	26.333433	$-7.1 \cdot 10^{-11}$	$2.297 \cdot 10^{-17}$	4.850122	0.003055	<b>192.71</b>

TABLE 4.7: Loudest template recovered by the multi-stage MCMC follow up for each of the analyzed outliers. Boldface and italic  $\ln \mathcal{B}_{S/N}^*$  values correspond to the two sets of outliers highlighted with circles and diamonds in Fig 4.12, respectively.



This represents the first application of a multi-stage MCMC-based follow-up to CW outliers from real data. The scalability of this development is such that it can be taken as a default follow-up strategy to outliers produced by virtually any CW search, as long as they can be related to a well-defined parameter space region. This allows for the general application of long-coherence follow-ups, massively reducing the complexity associated with the setup and calibration of ad hoc vetoes in CW searches.

## 4.A On the distribution of the maximum $\mathcal{F}$ -statistic and the effective number of templates

The validity of using an effective number of templates to fit Eq. (4.16) for the expected maximum  $\mathcal{F}$ -statistic from a search over a certain actual number of templates, in the presence of *non-independent* templates, has been discussed in the CW literature [52, 53]. We attempt to shed some light on the topic using extreme value theory. Concretely, we analyze the toy model posed in Appendix D of [53].

The basic point in [53] is that the presence of correlated templates not only changes the effective number of templates, but also the “functional form” of the resulting distribution, rendering Eq. (4.16) inaccurate. As an example, a toy model is constructed by generating a time series of zero-mean unit-variance Gaussian noise and computing the power of its Fourier transform. By choosing a suitable normalization, said power is the squared sum of two identical zero-mean Gaussian variables, following a chi-squared distribution with two degrees of freedom. This distribution can be properly fitted using Eq. (4.16), and the effective number of templates  $\mathcal{N}'$  is consistent with the number of frequency samples  $\mathcal{N} = N/2 - 1$ , where  $N$  is the number of elements from the original time series. Correlated templates are then introduced by over-resolving the Fourier transform applying zero-padding to the time series. The resulting distribution cannot be properly fitted using Eq. (4.16). The effective number of independent templates  $\mathcal{N}'$  is found to *increase* with the length of zero-padding, but it remains bounded by the actual number of power samples  $\mathcal{N}$ .

We provide an explanation for the two main issues raised in [53], namely what is the actual “functional form” of the target distribution and why the effective number of templates seems to increase as more correlated templates are included.

Let  $x_{n=1,\dots,N}$  be a zero-mean unit-variance Gaussian process. We define its Fourier transform as

$$\tilde{x}_k = \sum_{n=0}^{N-1} x_n e^{-2\pi i n \frac{k}{N}} \quad (4.32)$$

where  $k = 0, \dots, N - 1$ . Since  $x_n \in \mathbb{R}$ , the real and imaginary parts of Eq. (4.32) follow a zero-mean Gaussian distribution

$$\begin{aligned} \Re \tilde{x}_k &\sim \text{Gauss}(0, \sqrt{N/2}) \\ \Im \tilde{x}_k &\sim \text{Gauss}(0, \sqrt{N/2}) \end{aligned} \quad (4.33)$$

We then define power as

$$\tilde{\rho} = \left( \sqrt{\frac{2}{N}} \Re \tilde{x}_k \right)^2 + \left( \sqrt{\frac{2}{N}} \Im \tilde{x}_k \right)^2 \quad (4.34)$$

which, by definition, follows a chi-squared distribution with *two* degrees of freedom  $\tilde{\rho} \sim \chi_2^2$ . This same quantity is referred to as  $2\mathcal{F}_2$  in [53].

The case of a chi-squared distribution with two degrees of freedom is degenerate with an exponential distribution. For the sake of clarity, we re-express it as a gamma distribution with shape parameter  $k = 1$  and scale parameter  $\theta = 2$ , i.e.  $\tilde{\rho} \sim \Gamma(1, 2)$ . We note that chi-squared distributions correspond to the locus  $\theta = 2$  in the parameter space of Gamma distributions, with  $k$  equal to half the degrees of freedom; exponential distributions correspond to the locus  $k = 1$ , with  $\theta$  equal to the inverse of the rate parameter.

Let us now define  $x_n^p$  as the zero-padded time series containing  $Np$  elements, the last  $N(p - 1)$  of which are purposely zero. This padding re-scales the variance of the original distribution by a factor

$1/p$  and the resulting power can be expressed as

$$\tilde{\rho}_p = \left( \sqrt{\frac{2p}{N}} \Re \tilde{x}_k^p \right)^2 + \left( \sqrt{\frac{2p}{N}} \Im \tilde{x}_k^p \right)^2 = p\tilde{\rho}. \quad (4.35)$$

Then, by the properties of the Gamma function,  $\tilde{\rho}_p \sim \Gamma(1, 2p)$ , which is *not* a chi-squared distribution for  $p > 1$ , but an exponential distribution with rate parameter  $\lambda = (2p)^{-1}$ .

Finally, we discuss the asymptotics of the distribution followed by the maximum of a  $\Gamma$ -distributed random variable. As explained in Sec. 4.3.1, such light-tailed distributions fall under the domain of attraction of the Gumbel distribution, meaning

$$\max_{\mathcal{N}} \Gamma(k, \theta) \xrightarrow{\mathcal{N} \rightarrow \infty} \text{Gumbel}(\mu, \sigma), \quad (4.36)$$

where the location and scale parameters  $(\mu, \sigma)$  are given by [70]

$$\mu = \theta [\ln \mathcal{N} + (k-1) \ln \ln \mathcal{N} - \ln \Gamma(k)], \quad (4.37)$$

$$\sigma = \theta. \quad (4.38)$$

In particular, the case of  $\max_{\mathcal{N}} \tilde{\rho}_p$  results in

$$\mu_p(\mathcal{N}) = 2p \ln \mathcal{N}, \quad \sigma_p = 2p. \quad (4.39)$$

It is clear from Eq. (4.38) that the asymptotic distribution described by Eq. (4.16) is a Gumbel distribution with a scale parameter  $\sigma = 2$ . On the other hand, the asymptotic distribution followed by zero-padded Gaussian noise ( $p > 1$ ) follows a Gumbel distribution with a scale factor  $\sigma_p = 2p > 2$ . Since the scale parameter is *independent* of  $\mathcal{N}$ , Eq. (4.16) *fails* to describe the asymptotic distribution stemming from correlated templates. In other words, parameter-space correlations shift the distribution followed by the power statistic away from the locus of chi-squared distributions; since these correlations are generally contained in a certain characteristic length, the resulting light tails are still, however, within the Gumbel distribution's domain of attraction [67].

This result is consistent with the findings reported in Fig. 11 of [53], which we reproduce in Fig. 4.14. As the zero-padding increases,  $\sigma_p$  increases and the resulting distribution, which is well described by a Gumbel distribution, spreads beyond the fit provided by Eq. (4.16).

The location parameter  $\mu$ , on the other hand, does depend on the number of templates. Indeed, if one tries to compute the required effective number of templates  $\mathcal{N}'$  so that  $\mu_{p=1}(\mathcal{N}')$  coincides with  $\mu_p(\mathcal{N})$ ,

$$\mu_{p=1}(\mathcal{N}') = 2 \ln \mathcal{N}' = 2p \ln \mathcal{N} = \mu_p(\mathcal{N}), \quad (4.40)$$

the result is

$$\mathcal{N}' = \mathcal{N}^p, \quad (4.41)$$

which is a monotonic function of  $p$ . As a result, the effective number of templates *increases* with the zero-padding factor, again in agreement with [53]. We note, however, that this is just a consequence of the chosen Fourier normalization. If the normalized power was constructed using  $Np$  as a normalization (the actual number of samples) rather than  $N$  (the number of non-zero-padded samples), then Eq. (4.35) would be re-written as

$$\hat{\rho}_p = \frac{1}{p} \tilde{\rho}. \quad (4.42)$$

Consequently  $\hat{\rho}_p \sim \Gamma(1, 2/p)$  and the effects on the standard deviation would be exactly the opposite, as shown in Fig. 4.15. Indeed, in such a case the effective number of templates would be  $\mathcal{N}' = \mathcal{N}^{1/p}$ , which decreases as the zero-padding increases.

Our proposed solution to the problem of estimating the effective number of templates is then *not to do so*, as it depends strongly on the specific distribution followed by the noise, which is generally unknown in a real case. Instead, we propose to describe the background noise distribution by fitting an extreme value distribution to a set of samples (see e.g. Section 4.3.1). For light-tailed noise, the proper distribution is Gumbel; other distributions are available for noise falling off as a power law or presenting an upper cut-off.




FIGURE 4.14: Maximum Fourier power over  $N = 200$  samples of zero-mean unit-variance zero-padded Gaussian noise. In each panel, the stair-case line represents a histogram over  $10^6$  repeated trials of  $\max_{\mathcal{N}} \tilde{\rho}_p$ . The dashed line is the best fit of Eq. (4.16) on the effective number of templates  $\mathcal{N}'$ , and the solid line is the best fit of a Gumbel distribution on the location and scale parameters. Zero-padding is indicated by  $p$ , where  $p = 1$  represents no zero-padding, as explained in the text.

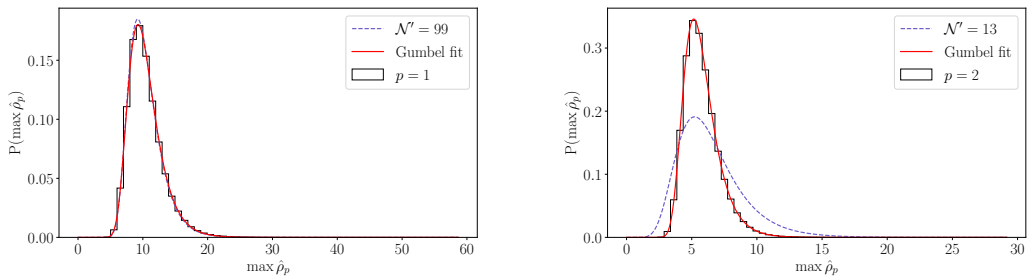


FIGURE 4.15: Equivalent figure to Fig. 4.14 using the alternative normalization of Fourier power  $\hat{\rho}_p$ . In this case, increasing the number of correlated templates *narrows* the resulting distribution with respect to Eq. (4.16).



# Bibliography

- [1] R. Abbott *et al.*, “GWTC-2: Compact Binary Coalescences Observed by LIGO and Virgo During the First Half of the Third Observing Run,” *Physical Review X*, vol. 11, p. 021053, 2021.
- [2] J. Aasi *et al.*, “Advanced LIGO,” *Classical and Quantum Gravity*, vol. 32, p. 074001, mar 2015.
- [3] F. Acernese *et al.*, “Advanced Virgo: a second-generation interferometric gravitational wave detector,” *Classical and Quantum Gravity*, vol. 32, p. 024001, dec 2014.
- [4] M. Sieniawska and M. Bejger, “Continuous gravitational waves from neutron stars: current status and prospects,” *Universe*, vol. 5, no. 11, p. 217, 2019.
- [5] S. J. Zhu, M. Baryakhtar, M. A. Papa, D. Tsuna, N. Kawanaka, and H.-B. Eggenstein, “Characterizing the continuous gravitational-wave signal from boson clouds around galactic isolated black holes,” *Physical Review D*, vol. 102, p. 063020, Sep 2020.
- [6] P. R. Brady, T. Creighton, C. Cutler, and B. F. Schutz, “Searching for periodic sources with LIGO,” *Physical Review D*, vol. 57, pp. 2101–2116, 1998.
- [7] P. R. Brady and T. Creighton, “Searching for periodic sources with LIGO II: Hierarchical searches,” *Physical Review D*, vol. 61, p. 082001, 2000.
- [8] C. Cutler, I. Gholami, and B. Krishnan, “Improved stack-slide searches for gravitational-wave pulsars,” *Physical Review D*, vol. 72, p. 042004, 2005.
- [9] B. Krishnan, A. M. Sintes, M. A. Papa, B. F. Schutz, S. Frasca, and C. Palomba, “Hough transform search for continuous gravitational waves,” *Physical Review D*, vol. 70, p. 082001, Oct 2004.
- [10] P. Astone, A. Colla, S. D’Antonio, S. Frasca, and C. Palomba, “Method for all-sky searches of continuous gravitational wave signals using the frequency-Hough transform,” *Physical Review D*, vol. 90, p. 042002, Aug. 2014.
- [11] K. Wette, S. Walsh, R. Prix, and M. A. Papa, “Implementing a semicoherent search for continuous gravitational waves using optimally-constructed template banks,” *Physical Review D*, vol. 97, no. 12, p. 123016, 2018.
- [12] M. Sieniawska, M. Bejger, P. Ciecieląg, and A. Królak, “Followup procedure in time-domain  $\mathcal{F}$ -statistic searches for continuous gravitational waves,” in *XXXVIII Polish Astronomical Society Meeting* (A. Rózańska, ed.), vol. 7, pp. 37–40, Aug. 2018.
- [13] V. Dergachev and M. A. Papa, “Sensitivity improvements in the search for periodic gravitational waves using O1 LIGO data,” *Physical Review Letters*, vol. 123, no. 10, p. 101101, 2019.
- [14] P. B. Covas and A. M. Sintes, “New method to search for continuous gravitational waves from unknown neutron stars in binary systems,” *Physical Review D*, vol. 99, no. 12, p. 124019, 2019.
- [15] L. Sancho de la Jordana and A. M. Sintes, “A  $\chi^2$  veto for continuous wave searches,” *Classical and Quantum Gravity*, vol. 25, p. 184014, 2008.
- [16] B. Behnke, M. A. Papa, and R. Prix, “Postprocessing methods used in the search for continuous gravitational-wave signals from the galactic center,” *Physical Review D*, vol. 91, p. 064007, Mar 2015.
- [17] P. Leaci, “Methods to filter out spurious disturbances in continuous-wave searches from gravitational-wave detectors,” *Physica Scripta*, vol. 90, no. 12, p. 125001, 2015.

- [18] S. J. Zhu, M. A. Papa, and S. Walsh, “New veto for continuous gravitational wave searches,” *Physical Review D*, vol. 96, no. 12, p. 124007, 2017.
- [19] F. Morawski, M. Bejger, and P. Ciecieląg, “Convolutional neural network classifier for the output of the time-domain  $\mathcal{F}$ -statistic all-sky search for continuous gravitational waves,” *Machine Learning: Science and Technology*, vol. 1, p. 025016, jun 2020.
- [20] B. P. Abbott *et al.*, “All-sky search for continuous gravitational waves from isolated neutron stars using Advanced LIGO O2 data,” *Physical Review D*, vol. 100, p. 024004, Jul 2019.
- [21] R. Abbott *et al.*, “All-sky search in early O3 LIGO data for continuous gravitational-wave signals from unknown neutron stars in binary systems,” *Physical Review D*, vol. 103, no. 6, p. 064017, 2021.
- [22] A. Singh, M. A. Papa, H.-B. Eggenstein, and S. Walsh, “Adaptive clustering procedure for continuous gravitational wave searches,” *Physical Review D*, vol. 96, no. 8, p. 082003, 2017.
- [23] B. Beheshtipour and M. A. Papa, “Deep learning for clustering of continuous gravitational wave candidates,” *Physical Review D*, vol. 101, no. 6, p. 064009, 2020.
- [24] B. Beheshtipour and M. A. Papa, “Deep learning for clustering of continuous gravitational wave candidates II: identification of low-SNR candidates,” *Physical Review D*, vol. 103, no. 6, p. 064027, 2021.
- [25] R. Tenorio, D. Keitel, and A. M. Sintes, “Time-frequency track distance for comparing continuous gravitational wave signals,” *Physical Review D*, vol. 103, no. 6, p. 064053, 2021.
- [26] M. Shaltev and R. Prix, “Fully coherent follow-up of continuous gravitational-wave candidates,” *Physical Review D*, vol. 87, no. 8, p. 084057, 2013.
- [27] M. Shaltev, P. Leaci, M. A. Papa, and R. Prix, “Fully coherent follow-up of continuous gravitational-wave candidates: an application to Einstein@Home results,” *Physical Review D*, vol. 89, no. 12, p. 124030, 2014.
- [28] M. A. Papa *et al.*, “Hierarchical follow-up of subthreshold candidates of an all-sky Einstein@Home search for continuous gravitational waves on LIGO sixth science run data,” *Physical Review D*, vol. 94, no. 12, p. 122006, 2016.
- [29] J. Ming *et al.*, “Results from an Einstein@Home search for continuous gravitational waves from Cassiopeia A, Vela Jr. and G347.3,” *Physical Review D*, vol. 100, no. 2, p. 024063, 2019.
- [30] P. B. Covas and A. M. Sintes, “First All-Sky Search for Continuous Gravitational-Wave Signals from Unknown Neutron Stars in Binary Systems Using Advanced LIGO Data,” *Physical Review Letters*, vol. 124, p. 191102, May 2020.
- [31] B. Steltner, M. A. Papa, H. B. Eggenstein, B. Allen, V. Dergachev, R. Prix, B. Machenschalk, S. Walsh, S. J. Zhu, and S. Kwang, “Einstein@Home All-sky Search for Continuous Gravitational Waves in LIGO O2 Public Data,” *The Astrophysical Journal*, vol. 909, no. 1, p. 79, 2021.
- [32] G. Ashton and R. Prix, “Hierarchical multistage MCMC follow-up of continuous gravitational wave candidates,” *Physical Review D*, vol. 97, no. 10, p. 103020, 2018.
- [33] Prix, Reinhard, “Coherent  $\mathcal{F}$ -statistic on semi-coherent candidate.” <https://dcc.ligo.org/LIGO-T1700236/public>, 2019.
- [34] D. Keitel, R. Tenorio, G. Ashton, and R. Prix, “PyFstat: a Python package for continuous gravitational-wave data analysis,” *Journal of Open Source Software*, vol. 6, no. 60, p. 3000, 2021.
- [35] G. Ashton, R. Prix, and D. I. Jones, “A semicoherent glitch-robust continuous-gravitational-wave search method,” *Physical Review D*, vol. 98, no. 6, p. 063011, 2018.
- [36] R. Prix, S. Giampanis, and C. Messenger, “Search method for long-duration gravitational-wave transients from neutron stars,” *Physical Review D*, vol. 84, p. 023007, 2011.

- [37] D. Keitel and G. Ashton, “Faster search for long gravitational-wave transients: GPU implementation of the transient  $\mathcal{F}$ -statistic,” *Classical and Quantum Gravity*, vol. 35, no. 20, p. 205003, 2018.
- [38] P. Jaranowski, A. Królak, and B. F. Schutz, “Data analysis of gravitational-wave signals from spinning neutron stars: The signal and its detection,” *Physical Review D*, vol. 58, p. 063001, Aug 1998.
- [39] E. T. Jaynes, *Probability Theory: The Logic of Science*. Cambridge University Press, 2003.
- [40] R. Prix and B. Krishnan, “Targeted search for continuous gravitational waves: Bayesian versus maximum-likelihood statistics,” *Classical and Quantum Gravity*, vol. 26, p. 204013, 2009.
- [41] J. T. Whelan, R. Prix, C. J. Cutler, and J. L. Willis, “New Coordinates for the Amplitude Parameter Space of Continuous Gravitational Waves,” *Classical and Quantum Gravity*, vol. 31, p. 065002, 2014.
- [42] C. Cutler and B. F. Schutz, “Generalized  $\mathcal{F}$ -statistic: Multiple detectors and multiple gravitational wave pulsars,” *Physical Review D*, vol. 72, p. 063006, Sep 2005.
- [43] P. Leaci and R. Prix, “Directed searches for continuous gravitational waves from binary systems: parameter-space metrics and optimal Scorpius X-1 sensitivity,” *Physical Review D*, vol. 91, no. 10, p. 102003, 2015.
- [44] D. Keitel, R. Prix, M. A. Papa, P. Leaci, and M. Siddiqi, “Search for continuous gravitational waves: Improving robustness versus instrumental artifacts,” *Physical Review D*, vol. 89, no. 6, p. 064023, 2014.
- [45] D. Keitel, “Robust semicoherent searches for continuous gravitational waves with noise and signal models including hours to days long transients,” *Physical Review D*, vol. 93, no. 8, p. 084024, 2016.
- [46] R. Prix and M. Shaltev, “Search for Continuous Gravitational Waves: Optimal StackSlide method at fixed computing cost,” *Physical Review D*, vol. 85, p. 084010, 2012.
- [47] R. Prix, “Search for continuous gravitational waves: Metric of the multi-detector  $\mathcal{F}$ -statistic,” *Physical Review D*, vol. 75, p. 023004, 2007. [Erratum: *Physical Review D* 75, 069901(E) (2007)].
- [48] K. Wette, “Empirically extending the range of validity of parameter-space metrics for all-sky searches for gravitational-wave pulsars,” *Physical Review D*, vol. 94, no. 12, p. 122002, 2016.
- [49] B. Allen, “Spherical ansatz for parameter-space metrics,” *Physical Review D*, vol. 100, no. 12, p. 124004, 2019.
- [50] B. Allen, “Optimal template banks,” *Physical Review D*, vol. 104, no. 4, p. 042005, 2021.
- [51] K. Wette, “Lattice template placement for coherent all-sky searches for gravitational-wave pulsars,” *Physical Review D*, vol. 90, no. 12, p. 122010, 2014.
- [52] K. Wette, “Estimating the sensitivity of wide-parameter-space searches for gravitational-wave pulsars,” *Physical Review D*, vol. 85, p. 042003, 2012.
- [53] C. Dreissigacker, R. Prix, and K. Wette, “Fast and Accurate Sensitivity Estimation for Continuous-Gravitational-Wave Searches,” *Physical Review D*, vol. 98, no. 8, p. 084058, 2018.
- [54] V. Dergachev, “On blind searches for noise dominated signals: a loosely coherent approach,” *Classical and Quantum Gravity*, vol. 27, p. 205017, 2010.
- [55] M. Shaltev, “Optimizing the StackSlide setup and data selection for continuous-gravitational-wave searches in realistic detector data,” *Physical Review D*, vol. 93, no. 4, p. 044058, 2016.
- [56] R. Prix, “Template-based searches for gravitational waves: Efficient lattice covering of flat parameter spaces,” *Classical and Quantum Gravity*, vol. 24, pp. S481–S490, 2007.

- [57] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, “emcee: The MCMC Hammer,” *Publications of the Astronomical Society of the Pacific*, vol. 125, p. 306, Mar. 2013.
- [58] W. D. Vousden, W. M. Farr, and I. Mandel, “Dynamic temperature selection for parallel tempering in Markov chain Monte Carlo simulations,” *Monthly Notices of the Royal Astronomical Society*, vol. 455, pp. 1919–1937, Jan. 2016.
- [59] G. Ashton, D. Keitel, R. Prix, and R. Tenorio, “Pyfstat/pyfstat: v1.11.3,” Feb. 2021. <https://doi.org/10.5281/zenodo.4542822>.
- [60] K. Wette and R. Prix, “Flat parameter-space metric for all-sky searches for gravitational-wave pulsars,” *Physical Review D*, vol. 88, no. 12, p. 123005, 2013.
- [61] K. Wette, “Parameter-space metric for all-sky semicoherent searches for gravitational-wave pulsars,” *Physical Review D*, vol. 92, no. 8, p. 082003, 2015.
- [62] H. J. Pletsch, “Parameter-space metric of semicoherent searches for continuous gravitational waves,” *Physical Review D*, vol. 82, p. 042002, 2010.
- [63] J. Abadie *et al.*, “First search for gravitational waves from the youngest known neutron star,” *The Astrophysical Journal*, vol. 722, pp. 1504–1513, 2010.
- [64] M. A. Papa, J. Ming, E. V. Gotthelf, B. Allen, R. Prix, V. Dergachev, H.-B. Eggenstein, A. Singh, and S. J. Zhu, “Search for Continuous Gravitational Waves from the Central Compact Objects in Supernova Remnants Cassiopeia A, Vela Jr., and G347.3–0.5,” *The Astrophysical Journal*, vol. 897, no. 1, p. 22, 2020.
- [65] K. W. Wette, *Gravitational waves from accreting neutron stars and Cassiopeia A*. PhD thesis, Australian Natl. U., Canberra, 2009.
- [66] K. Wette, L. Dunn, P. Clearwater, and A. Melatos, “Deep exploration for continuous gravitational waves at 171–172 Hz in LIGO second observing run data,” *Physical Review D*, vol. 103, no. 8, p. 083020, 2021.
- [67] M. Leadbetter, G. Lindgren, and H. Rootzen, *Extremes and Related Properties of Random Sequences and Processes*. Springer Series in Statistics, Springer New York, 1983.
- [68] J. Beirlant, Y. Goegebeur, J. Segers, J. Teugels, D. De Waal, and C. Ferro, *Statistics of Extremes: Theory and Applications*. Wiley Series in Probability and Statistics, Wiley, 2004.
- [69] L. de Haan and A. Ferreira, *Extreme Value Theory: An Introduction*. Springer Series in Operations Research and Financial Engineering, Springer New York, 2006.
- [70] P. Embrechts, C. Klüppelberg, and T. Mikosch, *Modelling Extremal Events: for Insurance and Finance*. Stochastic Modelling and Applied Probability, Springer Berlin Heidelberg, 2013.
- [71] S. Suvorova, P. Clearwater, A. Melatos, L. Sun, W. Moran, and R. J. Evans, “Hidden Markov model tracking of continuous gravitational waves from a binary neutron star with wandering spin. II. Binary orbital phase tracking,” *Physical Review D*, vol. 96, p. 102006, Nov. 2017.
- [72] A. Mathai and S. Provost, *Quadratic Forms in Random Variables*. Statistics: A Series of Textbooks and Monographs, Taylor & Francis, 1992.
- [73] A. Gasull, J. López-Salcedo, and F. Utzet, “Maxima of Gamma random variables and other Weibull-like distributions and the Lambert W function,” *TEST: An Official Journal of the Spanish Society of Statistics and Operations Research*, vol. 24, pp. 714–733, December 2015.
- [74] R. Tenorio, L. M. Modafferi, D. Keitel, and A. M. Sintes, “in preparation,” Tech. Rep. LIGO-P2100277, LIGO Scientific Collaboration, 2021.
- [75] M. Isi, S. Mastroianni, M. Pitkin, and O. J. Piccinni, “Establishing the significance of continuous gravitational-wave detections from known pulsars,” *Physical Review D*, vol. 102, no. 12, p. 123027, 2020.



- [76] R. Muirhead, *Aspects of Multivariate Statistical Theory*. Wiley Series in Probability and Statistics, Wiley, 2005.
- [77] D. Horgan and C. C. Murphy, “On the convergence of the chi square and noncentral chi square distributions to the normal distribution,” *IEEE Communications Letters*, vol. 17, no. 12, pp. 2233–2236, 2013.
- [78] V. Dergachev, “Loosely coherent searches for sets of well-modeled signals,” *Physical Review D*, vol. 85, p. 062003, 2012.
- [79] V. Dergachev, “Loosely coherent searches for medium scale coherence lengths.” arXiv:1807.02351 [astro-ph.IM].
- [80] V. Dergachev and M. A. Papa, “Results from the First All-Sky Search for Continuous Gravitational Waves from Small-Ellipticity Sources,” *Physical Review Letters*, vol. 125, no. 17, p. 171101, 2020.
- [81] V. Dergachev and M. A. Papa, “Results from high-frequency all-sky search for continuous gravitational waves from small-ellipticity sources,” *Physical Review D*, vol. 103, no. 6, p. 063019, 2021.
- [82] V. Dergachev and M. A. Papa, “Search for continuous gravitational waves from small-ellipticity sources at low frequencies,” *Phys. Rev. D*, vol. 104, no. 4, p. 043003, 2021.
- [83] D. P. Anderson, “BOINC: A Platform for Volunteer Computing,” *arXiv e-prints*, Mar. 2019.
- [84] H. J. Pletsch, “Parameter-space correlations of the optimal statistic for continuous gravitational-wave detection,” *Physical Review D*, vol. 78, p. 102005, 2008.
- [85] H. J. Pletsch and B. Allen, “Exploiting global correlations to detect continuous gravitational waves,” *Physical Review Letters*, vol. 103, p. 181102, 2009.
- [86] S. Walsh, K. Wette, M. A. Papa, and R. Prix, “Optimizing the choice of analysis method for all-sky searches for continuous gravitational waves with Einstein@Home,” *Physical Review D*, vol. 99, p. 082004, Apr. 2019.
- [87] E. Pfeiffermann and B. Aschenbach, “ROSAT observation of a new supernova remnant in the constellation Scorpius,” in *Roentgenstrahlung from the Universe* (H. U. Zimmermann, J. Trümper, and H. Yorke, eds.), pp. 267–268, Feb. 1996.
- [88] S. Suvorova, L. Sun, A. Melatos, W. Moran, and R. J. Evans, “Hidden Markov model tracking of continuous gravitational waves from a neutron star with wandering spin,” *Physical Review D*, vol. 93, no. 12, p. 123009, 2016.
- [89] L. Sun, A. Melatos, S. Suvorova, W. Moran, and R. J. Evans, “Hidden Markov model tracking of continuous gravitational waves from young supernova remnants,” *Physical Review D*, vol. 97, p. 043013, Feb. 2018.
- [90] J. Bayley, C. Messenger, and G. Woan, “Generalized application of the Viterbi algorithm to searches for continuous gravitational-wave signals,” *Physical Review D*, vol. 100, p. 023006, July 2019.
- [91] A. Mukherjee, C. Messenger, and K. Riles, “Accretion-induced spin-wandering effects on the neutron star in Scorpius X-1: Implications for continuous gravitational wave searches,” *Physical Review D*, vol. 97, no. 4, p. 043016, 2018.
- [92] D. Jones and L. Sun, “Search for continuous gravitational waves from Fomalhaut b in the second Advanced LIGO observing run with a hidden Markov model,” *Physical Review D*, vol. 103, no. 2, p. 023020, 2021.
- [93] P. Kalas, J. R. Graham, E. Chiang, M. P. Fitzgerald, M. Clampin, E. S. Kite, K. Stapelfeldt, C. Marois, and J. Krist, “Optical Images of an Exosolar Planet 25 Light Years from Earth,” *Science*, vol. 322, pp. 1345–1348, 2008.

- [94] T. Currie, J. Debes, T. J. Rodigas, A. Burrows, Y. Itoh, M. Fukagawa, S. J. Kenyon, M. Kuchner, and S. Matsumura, “Direct Imaging Confirmation and Characterization of a Dust-enshrouded Candidate Exoplanet Orbiting Fomalhaut,” *The Astrophysical Journal Letters*, vol. 760, p. L32, Dec. 2012.
- [95] R. Neuhäuser, M. M. Hohle, C. Ginski, J. G. Schmidt, V. V. Hambaryan, and T. O. B. Schmidt, “The companion candidate near Fomalhaut - a background neutron star?,” *Monthly Notices of the Royal Astronomical Society*, vol. 448, pp. 376–389, Mar. 2015.
- [96] A. Gaspar and G. Rieke, “New HST data and modeling reveal a massive planetesimal collision around Fomalhaut,” *Proceedings of the National Academy of Science*, vol. 117, pp. 9712–9722, Apr. 2020.
- [97] B. P. Abbott *et al.*, “Searches for Continuous Gravitational Waves from 15 Supernova Remnants and Fomalhaut b with Advanced LIGO,” *The Astrophysical Journal*, vol. 875, no. 2, p. 122, 2019.
- [98] D. Beniwal, P. Clearwater, L. Dunn, A. Melatos, and D. Ottaway, “Search for continuous gravitational waves from ten H.E.S.S. sources using a hidden Markov model,” *Physical Review D*, vol. 103, no. 8, p. 083009, 2021.
- [99] R. Abbott *et al.*, “Open data from the first and second observing runs of advanced ligo and advanced virgo,” *SoftwareX*, vol. 13, p. 100658, 2021.
- [100] LIGO Scientific Collaboration and Virgo Collaboration, “Gravitational Wave Open Science Center - Advanced LIGO O2 Data Release.” <https://www.gw-openscience.org>, 2019.
- [101] Goetz, Evan, “Segments used for creating standard SFTs in O2 data.” <https://dcc.ligo.org/LIGO-T1900085/public>.
- [102] G. Ashton, R. Prix, and D. I. Jones, “Statistical characterization of pulsar glitches and their potential impact on searches for continuous gravitational waves,” *Physical Review D*, vol. 96, no. 6, p. 063004, 2017.
- [103] R. Prix and Y. Itoh, “Global parameter-space correlations of coherent searches for continuous gravitational waves,” *Classical and Quantum Gravity*, vol. 22, pp. S1003–S1012, 2005.
- [104] LIGO Scientific Collaboration, “LIGO Algorithm Library - LALSuite.” free software (GPL), 2018.
- [105] J. T. Whelan, S. Sundaresan, Y. Zhang, and P. Peiris, “Model-Based Cross-Correlation Search for Gravitational Waves from Scorpius X-1,” *Physical Review D*, vol. 91, p. 102005, 2015.
- [106] S. Walsh, M. Pitkin, M. Oliver, S. D’Antonio, V. Dergachev, A. Królak, *et al.*, “Comparison of methods for the detection of gravitational waves from unknown neutron stars,” *Physical Review D*, vol. 94, no. 12, p. 124010, 2016.
- [107] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, *et al.*, “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python,” *Nature Methods*, vol. 17, pp. 261–272, 2020.
- [108] C. Messenger, R. Prix, and M. A. Papa, “Random template banks and relaxed lattice coverings,” *Physical Review D*, vol. 79, p. 104017, 2009.
- [109] B. Allen and A. A. Shoom, “Template banks based on  $\mathbb{Z}^n$  and  $A_n^*$  lattices.” arXiv:2102.11631 [astro-ph.IM].
- [110] P. B. Covas *et al.*, “Identification and mitigation of narrow spectral artifacts that degrade searches for persistent gravitational waves in the first two observing runs of Advanced LIGO,” *Physical Review D*, vol. 97, p. 082002, Apr. 2018.
- [111] C. Biwer, D. Barker, J. C. Batch, J. Betzwieser, R. P. Fisher, E. Goetz, *et al.*, “Validating gravitational-wave detections: The Advanced LIGO hardware injection system,” *Physical Review D*, vol. 95, no. 6, p. 062002, 2017.

- [112] J. Aasi *et al.*, “Directed search for continuous gravitational waves from the Galactic center,” *Physical Review D*, vol. 88, no. 10, p. 102002, 2013.
- [113] B. P. Abbott *et al.*, “First low-frequency Einstein@Home all-sky search for continuous gravitational waves in Advanced LIGO data,” *Physical Review D*, vol. 96, no. 12, p. 122004, 2017.
- [114] O. J. Piccinni, P. Astone, S. D’Antonio, S. Frasca, G. Intini, I. La Rosa, P. Leaci, S. Mastrogiovanni, A. Miller, and C. Palomba, “Directed search for continuous gravitational-wave signals from the Galactic Center in the Advanced LIGO second observing run,” *Physical Review D*, vol. 101, no. 8, p. 082004, 2020.
- [115] J. Ming, M. A. Papa, H.-B. Eggenstein, B. Machenschalk, B. Steltner, R. Prix, B. Allen, and O. Behnke, “Results from an Einstein@Home search for continuous gravitational waves from G347.3 at low frequencies in LIGO O2 data.” arXiv:2108.02808 [gr-qc].



## Chapter 5

# Empirically estimating the distribution of the loudest candidate from a gravitational-wave search

This chapter is an adaptation of the material presented in

*Empirically estimating the distribution of the loudest candidate from a gravitational-wave search*

Rodrigo Tenorio, Luana M. Modafferi David Keitel, Alicia M. Sintes

[Phys. Rev. D 105, 044029 \(2022\)](#) – [arXiv:2111.12032 \[gr-qc\]](#)

DOI: [10.1103/PhysRevD.105.044029](#)

### 5.1 Introduction

The search for gravitational-wave (GW) signals can be formulated as a multi-hypothesis test between a background-noise hypothesis and a set of signal hypotheses, each asserting the presence of a signal with a specific set of parameters [1]. Actual search implementations, however, usually split this process into three stages: a detection stage, which simply assesses the presence of a feature in the datastream unlikely to be caused by noise (null-hypothesis test); a validation stage, in which candidates are sieved through a set of vetoes to discard any instrumental causes; and a parameter-estimation stage, in which a proper Bayesian hypothesis test is carried out to infer the actual parameters of any detected signal. This division is motivated by the increasing computing cost of each stage [2, 3], as a simple null-hypothesis test (usually assuming Gaussian noise) is orders of magnitude more affordable than a single parameter-estimation stage.

The standard detection stage consists in performing a finite number of detection statistic evaluations over the parameter-space region of interest, usually using matched-filtering against a bank of waveform *templates* [4, 2, 5, 6, 7, 8, 9, 10, 11]. Loud templates, i.e. those scoring a high detection statistic, are deemed “signal candidates” and selected for the validation stage. The detection statistic can be usually interpreted as a Bayes factor, assessing the preference of the data for a particular signal hypothesis (represented by the template at hand) versus the background-noise hypothesis. Thus, the detection stage is a multi-hypothesis test in disguise in which parameter-space *marginalization* has been approximated to zeroth-order by *maximization* [12].

Loudest candidates from a template bank fall generally into one of two categories: The strongest excursions away from the background, such as an instrumental feature [13, 14] or a very clear GW signal (such as GW150914 [15]), are usually comparatively simple to deal with, as strong candidates tend to show characteristic signatures according to their cause. But weaker outliers that are in principle compatible with both a weak signal or an extreme event of the general noise background require a more careful analysis.

While much of the statistical framework used in this work is generally applicable, we mainly focus on the search for continuous gravitational-wave signals (CWs) [16], produced by long-standing quadrupolar deformations, such as in the case of non-axisymmetric spinning neutron stars (NS) [17]. From the point of view of the current generation of advanced detectors (Advanced LIGO [18], Advanced Virgo [19], and KAGRA [20]), they belong in the weak-signal regime, meaning they are expected to blend into the

background distribution. Characterizing the expected distribution of extreme background candidates is, thus, a simple approach to identify interesting outliers in a search and quantify their significance.

Pioneering work on describing the distribution of the loudest candidate from a CW search using the  $\mathcal{F}$ -statistic [21, 22] was presented in [23] and later extended in [24, 25]. Despite its wide applicability in the CW literature (see e.g. [26, 27, 28, 29, 30, 31, 32]), basic assumptions of the method make it insufficient for realistic template banks with a certain degree of correlation between neighbouring templates [33]. Latest developments on the subject used extreme value theory (EVT) to propose a suitable ansatz to circumvent the problems posed by template-bank correlations [34]; but the concrete method requires re-evaluating full template banks many times (similar in spirit to that in [24]) and is thus computationally unsuitable for wide parameter-space searches.

This work proposes **distromax**, a new method to describe the distribution of the loudest candidate stemming from a generic GW search. The method generalizes with respect to previous approaches presented in [23, 25] in two main aspects. First, the method is robust to typical degrees of template-bank correlations arising either due to the overlap of nearby templates or mild non-Gaussianities in the data. Second, the method is applicable to a wider class of detection statistics, including other  $\mathcal{F}$ -statistic-based detection statistics such as line-robust statistics [35, 36] or transient CW search statistics [37], as well as detection statistics from other search approaches. An implementation of **distromax** is publicly available as a homonymous Python package [38].

The paper is structured as follows: Section 5.2 introduces basic data-analysis tools for CW searches and discusses the origin of parameter-space correlations. Section 5.3 describes the quantitative effect of parameter-space correlations on the distribution of the loudest candidate, comparing standard approaches in the field to extreme value theory results. Section 5.4 introduces **distromax** to estimate the distribution of the loudest outlier of a search and discusses its basic phenomenology on synthetic data. In Section 5.5, we apply **distromax** to the results of a search on O2 Advanced LIGO data for (transient) CW signals. Appendix 5.A collects basic results in extreme value theory and provides further references for the interested reader. Appendix 5.B proposes a simple method to deal with narrow-band noise disturbances, common in realistic CW searches. The robustness of **distromax** to the presence of weak CW signals is discussed in Appendix 5.C.

## 5.2 Continuous wave searches

In this section, we revisit the basics of CW searches to frame our discussion of **distromax**. Section 5.2.1 reintroduces the  $\mathcal{F}$ -statistic and explicitly constructs its distribution under the noise hypothesis; Sec. 5.2.2 uses the explicit construction to discuss the two possible origins of parameter-space correlations affecting a template bank; Sec. 5.2.3 completes the analysis deriving the standard result for the distribution of the  $\mathcal{F}$ -statistic under the signal hypothesis.

The response of a ground-based GW detector to a passing CW or long-duration CW-like transient (tCW) is given by the linear combination of four linear filters [21, 37]

$$s(t; \mathcal{A}, \lambda, \mathcal{T}) = w(t; \mathcal{T}) \sum_{\mu=0}^3 \mathcal{A}_\mu h_\mu(t; \lambda), \quad (5.1)$$

where  $\mathcal{A}$  represents the source's amplitude parameters, namely GW amplitude  $h_0$ , inclination angle  $\iota$ , polarization angle  $\psi$ , and initial phase  $\phi_0$ , which can be combined into the so-called JKS decomposition  $\{\mathcal{A}_\mu, \mu = 0, 1, 2, 3\}$ ; and  $\lambda$  describes the phase-evolution parameters, namely the GW frequency and spindown  $\{f_0, f_1, f_2, \dots\}$ , the sky position  $\vec{n}$ , and possibly binary orbital parameters if the source orbits a companion. The time-dependent quadratures  $h_\mu(t; \lambda)$  encompass the detector's antenna pattern effects on the signal. The window function  $w(t; \mathcal{T})$  is a time-dependent amplitude modulation parametrized by the transient parameters  $\mathcal{T}$  to account for tCW signals [37]. The standard CW signal model is recovered for  $w(t; \mathcal{T}) = 1 \forall t$ .

Given a datastream  $x$ , the detection problem consists in deciding between the background noise hypothesis  $\mathcal{H}_N$ , under which the data stream contains only Gaussian noise  $x = n$ , and the signal hypothesis  $\mathcal{H}_S$ , according to which there is a (t)CW signal with a defined set of parameters  $x = n + s(\lambda, \mathcal{A}, \mathcal{T})$ .

Further hypotheses accounting for different non-Gaussian populations, such as narrow instrumental artifacts in the data [13], can be also included in the analysis [35, 36], although the usual approach is to apply post-processing veto strategies targeting specific types of disturbances [39, 40, 28, 41].

### 5.2.1 $\mathcal{F}$ -statistic under the noise hypothesis

A basic tool to conduct CW searches is the  $\mathcal{F}$ -statistic, first introduced in [21, 22] as a maximum-likelihood estimator with respect to amplitude parameters  $\mathcal{A}$ , and later re-introduced in a Bayesian context [12, 42, 43, 44]. The basic idea is to exploit the linear dependency of Eq. (5.1) on  $\mathcal{A}$  to analytically marginalize the matched-filtering likelihood using a suitable set of priors. The result can be readily expressed as a quadratic form [45]

$$2\mathcal{F}(\lambda) = \sum_{\mu,\nu=0}^3 x_{\mu}(\lambda) \mathcal{M}_{\mu\nu}^{-1}(\lambda) x_{\nu}(\lambda), \quad (5.2)$$

where  $x_{\mu}$  are the projections of the data stream  $x$  onto the four quadrature functions

$$x_{\mu}(\lambda) = \langle h_{\mu}(\lambda), x \rangle \quad (5.3)$$

and  $\mathcal{M}^{-1}(\lambda)$  is the inverse Gram matrix associated to the four quadrature functions

$$\mathcal{M}_{\mu\nu}(\lambda) = \langle h_{\mu}(\lambda), h_{\nu}(\lambda) \rangle. \quad (5.4)$$

The functional scalar product [46]

$$\langle x, y \rangle = 4 \Re \int_0^{\infty} df \frac{x(f) y^*(f)}{S_n(f)} \quad (5.5)$$

accounts for the presence of correlated noise in the data stream through the single-sided power spectral density (PSD)  $S_n$ . Current implementations of Eq. (5.5) make use of the so called  $\mathcal{F}$ -statistic *atoms* [45], evaluated over individual Short Fourier Transforms (SFTs) of the data. These could be simply described as a set of complex-valued spectrograms (from now on *atomic spectrograms*) containing both phase and amplitude information, whose proper combination results in an efficient computation of Eq. (5.2).

Under the noise hypothesis  $\mathcal{H}_N$ , the data stream is composed of zero-mean Gaussian noise and Eq. (5.3) implies the four projections  $\{n_{\mu}(\lambda)\}$  are drawn from a 4-dimensional Gaussian distribution with covariance matrix  $\mathcal{M}(\lambda)$ . Hence,

$$\{n_{\mu}(\lambda)\} \sim \text{Gauss}(0, \mathcal{M}(\lambda)), \quad (5.6)$$

and  $n_{\mu}(\lambda)$  values can be constructed as a linear combination of four zero-mean unit-variance Gaussian random variables

$$n_{\mu}(\lambda) = \sum_{\nu=0}^3 L_{\mu\nu}(\lambda) g_{\nu}[\lambda], \quad (5.7)$$

where  $g_{\nu}[\lambda] \sim \text{Gauss}(0, 1)$  and  $L$  is a  $4 \times 4$  matrix such that  $LL^T = \mathcal{M}$  (e.g. Cholesky decomposition). Here the square brackets indicate that Gaussian numbers are to be drawn independently for each template  $\lambda$ , but their distribution does not depend on  $\lambda$ ; as opposed to round brackets, which represent deterministic relations.

Introducing these results into Eq. (5.2),

$$2\mathcal{F}(\lambda) = \sum_{\mu=0}^3 g_{\mu}[\lambda]^2, \quad (5.8)$$

we obtain  $2\mathcal{F}(\lambda)$  as the Euclidean norm of a 4-dimensional Gaussian vector. Consequently, the probability distribution associated to  $2\mathcal{F}$  under the noise hypothesis  $\mathcal{H}_N$  for a fixed template  $\lambda$  is given by a

$\chi^2$  distribution with four degrees of freedom

$$p(2\mathcal{F}|\mathcal{H}_N) = \chi_4^2(2\mathcal{F}) . \quad (5.9)$$

### 5.2.2 Template-bank correlations

The statistical properties of the right-hand side of Eq. (5.14) are independent of the specific phase-evolution template  $\lambda$  under consideration. This suggests that evaluating  $2\mathcal{F}$  over a template bank using a single noise realization could, under suitable conditions, be equivalent to evaluating  $2\mathcal{F}$  for a single template over an ensemble of noise realizations.

Gaussian vectors  $\{g_\nu[\lambda]\}$  are constructed from a noise stream as follows:

$$g_\nu[\lambda] = \sum_{\mu=0}^3 L_{\nu\mu}^{-1}(\lambda) \langle h_\mu(\lambda), n \rangle . \quad (5.10)$$

The noise stream is projected onto four different deterministic functions of  $\lambda$ ,  $\{h_\mu(\lambda)\}$ , and combined using a set of weights  $L_{\nu\mu}^{-1}(\lambda)$ , also dependent on  $\lambda$ . Such a projection is a weighted average of the atomic spectrogram bins visited by the frequency-evolution track associated to  $\lambda$ . Since the atomic spectrograms are constructed using finite time and frequency resolutions, the number of independent Gaussian vectors constructible out of them is equivalent to the number of templates with *non-overlapping* frequency tracks over the spectrograms (i.e. crossing different spectrogram bins). This result was stated in a simpler fashion in [23] by arguing that the typical number of bins in a narrow-banded atomic spectrogram is orders of magnitude smaller than the number of templates in a typical CW search crossing said spectrogram.

As discussed in [47], the average dissimilarity in frequency-evolution tracks of nearby parameter-space points is related to the fractional loss in detection statistic, usually referred to as *mismatch* [48]

$$m = 1 - \frac{2\mathcal{F}(\lambda + \Delta\lambda)}{2\mathcal{F}(\lambda)} \simeq \sum_{i,j} \Delta\lambda_i \Delta\lambda_j g_{ij} + \mathcal{O}(\Delta\lambda^3) , \quad (5.11)$$

where  $\Delta\lambda_i$  represents an offset in an arbitrary parameter-space dimension and  $g_{ij}$  is the parameter-space metric [48, 49, 50, 51, 52, 53]. In the context of a grid-based CW search, the parameter-space metric can be employed to set up a template bank at a pre-specified maximum mismatch value [7, 54, 55]: the higher the mismatch, the coarser the template bank. An ensemble of templates with non-overlapping frequency-evolution tracks, then, corresponds to a coarse-enough template bank in the sense of large parameter-space mismatch.

In a real search, template banks tend to be set up using a moderate mismatch (e. g.  $m \sim 0.2$ ) in order to produce dense-enough parameter-space coverings [7]. A first kind of template-bank correlation arises, then, as a result of the template-bank construction strategy. Latest developments on the subject [56, 53, 10, 55], however, suggest higher mismatch values ( $m \sim 1$ ) could actually be compatible with a successful CW search, potentially suppressing the effect of these correlations.

A second kind of template-bank correlations, briefly discussed in [34], arises due to non-Gaussianities in the data (e.g. narrow instrumental features [13] or transient “pizza-slice” disturbances [57, 36, 29]). In this case, it is not a matter of re-using the same data on different templates; rather, a region of a priori independent spectrogram bins gets correlated due to the presence of a strong disturbance. As a result, non-overlapping templates crossing said correlated spectrogram region become correlated as well.

This same formalism applies to the search for tCWs, as the standard strategy in such cases is either to maximize or marginalize out any dependency on the transient parameters [37], obtaining in the end a detection statistic over an equivalent template bank to that of CW searches. Discussion on specific tCW detection statistics is postponed to Sec. 5.5.

The presence of correlations in a template bank, thus, is a generic property of (t)CW searches, and their effects on any newly proposed method should be properly understood before attempting to interpret results on a real setup.



### 5.2.3 $\mathcal{F}$ -statistic under the signal hypothesis

We conclude this summary of standard CW search methods by considering the distribution of the  $\mathcal{F}$ -statistic when there is a signal in the data. To derive its probability distribution under the signal hypothesis  $\mathcal{H}_S$ , we simply apply  $n = x - s(\lambda, \mathcal{A})$  and repeat the same reasoning up to Eq. (5.7), obtaining

$$x_\mu(\lambda) = \sum_{\nu=0}^3 L_{\mu\nu}(\lambda) (g_\nu[\lambda] + m_\nu(\lambda, \mathcal{A})) , \quad (5.12)$$

where

$$m_\nu(\lambda, \mathcal{A}) = \sum_{\kappa=0}^3 L_{\nu\kappa}^{-1}(\lambda) s_\kappa(\lambda, \mathcal{A}) \quad (5.13)$$

and, consequently,  $g_\nu[\lambda] + m_\nu(\lambda, \mathcal{A})$  is a Gaussian random number with mean  $m_\nu$  and unit variance.

Introducing these results in Eq. (5.2),

$$2\mathcal{F}(\lambda) = \sum_{\nu=0}^3 (g_\nu[\lambda] + m_\nu(\lambda, \mathcal{A}))^2 , \quad (5.14)$$

the  $2\mathcal{F}$  under the signal hypothesis corresponds to the norm of a four-dimensional uncorrelated Gaussian vector with identity covariance matrix and mean vector equal to  $\{m_\nu(\lambda, \mathcal{A})\}$ . The probability distribution is, as a result, a *non-central* chi-squared distribution with four degrees of freedom

$$p(2\mathcal{F}|\rho^2\mathcal{H}_S) = \chi_4^2(2\mathcal{F}; \rho^2) , \quad (5.15)$$

where the non-centrality parameter  $\rho^2$  is defined as

$$\rho^2 = \sum_{\nu=0}^3 m_\nu^2 = \sum_{\mu, \nu=0}^3 \mathcal{A}_\mu \mathcal{M}_{\mu\nu} \mathcal{A}_\nu = \langle s, s \rangle . \quad (5.16)$$

This quantity is referred to as the (squared) signal-to-noise ratio (SNR) in the literature. Concretely,  $\rho^2$  is the maximum attainable SNR corresponding to the case where signal parameters are perfectly matched by a phase-evolution template [45].

## 5.3 Loudest candidates, the “effective number of templates”, and extreme value theory

The standard problem of estimating the distribution of the loudest candidate in a search is posed as follows: Let  $\xi = \{\xi_i, i = 1, \dots, \mathcal{N}\}$  be a set of detection statistic values obtained by evaluating a template bank with  $\mathcal{N}$  templates in a noise-only data stream. Let  $f$  be the probability distribution of such a detection statistic under the noise hypothesis. Describe the probability distribution of the loudest candidate  $\max_{i=1, \dots, \mathcal{N}} \xi_i$ .

We can easily construct said distribution using the joint cumulative density function (CDF) of the entire template bank

$$P(\max_i \xi_i \leq \xi^* | \mathcal{N}) = P(\xi_1 \leq \xi^* \text{ and } \dots \text{ and } \xi_{\mathcal{N}} \leq \xi^*) . \quad (5.17)$$

For the case of an uncorrelated template bank, each template is independent and the joint CDF factors into the product of individual CDFs:

$$P(\max_i \xi_i \leq \xi^* | \mathcal{N}) = \prod_{i=1}^{\mathcal{N}} P(\xi_i \leq \xi^*) = \left[ \int^{\xi^*} d\xi f(\xi) \right]^{\mathcal{N}} . \quad (5.18)$$

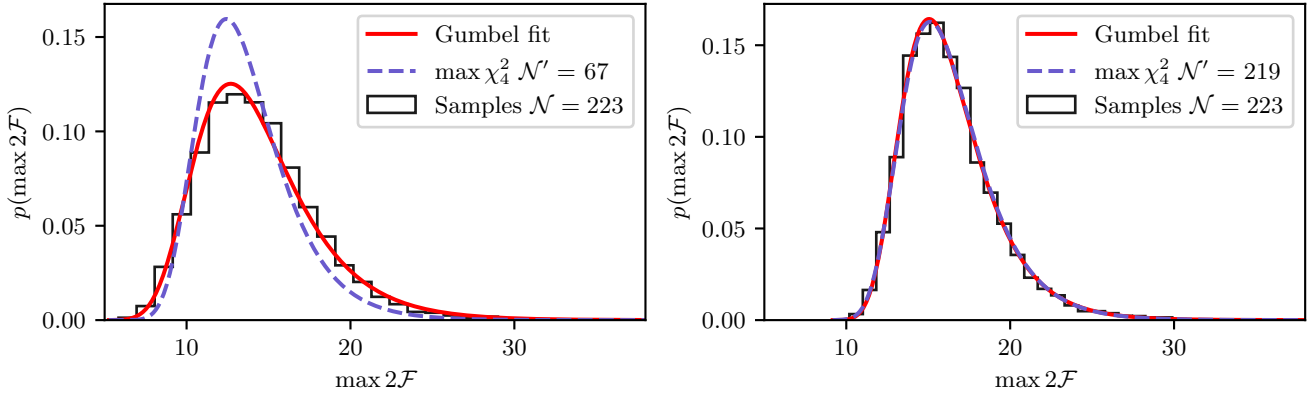


FIGURE 5.1: Distribution of the loudest  $2\mathcal{F}$  values produced by the evaluation of a template bank on a Gaussian-noise data stream lasting for 7 days. The template bank was set up using the `gridType=8` option of `ComputeFstatistic_v2` [59] with mismatch  $m = 0.2$ ,  $f_0 = 49.5$  Hz and  $f_1 = -10$  nHz/s covering bands of  $\Delta f_0 = 0.22$  Hz and  $\Delta f_1 = 45$  pHz/s. The sky position was fixed to a fiducial value  $(\alpha, \delta) = (0, 0)$  in equatorial coordinates. Loudest values were obtained by selecting the loudest  $2\mathcal{F}$  over different segmentations of the template bank. The left panel corresponds to selecting the loudest value within every 5 mHz subband. The right panel corresponds to shuffling the results and taking the loudest values from batches of the same size as the subbands. The stepped line is the histogram of the data; the dashed line is the best fit value  $\mathcal{N}'$  for  $\mathcal{N}$  in Eq. (5.19); and the solid line is the best fit of a Gumbel distribution.

Consequently, the probability density function associated to  $\xi^* = \max_{i=1, \dots, \mathcal{N}} \xi_i$  is simply

$$p(\xi^* | \mathcal{N}) = \mathcal{N} f(\xi^*) \left[ \int^{\xi^*} d\xi f(\xi) \right]^{\mathcal{N}-1}. \quad (5.19)$$

Template-bank correlations imply that we sample fewer independent combinations of the data than with an uncorrelated bank of the same  $\mathcal{N}$ . In other words, they reduce the “trials factor” of a search, diminishing the expected detection statistic of the loudest candidate in a similar fashion to evaluating a smaller template bank. Given a fixed false-alarm probability, neglecting template-bank correlations and naively using Eq. (5.19) would overestimate the corresponding threshold, potentially leading to missing interesting candidates.

Extensive analyses in [23, 58] concluded the effect of template-bank correlations on Eq. (5.19) could be reproduced to an acceptable level by adjusting  $\mathcal{N}$  to the “effective number of templates” in the template bank at hand. Although in some cases an empirical estimate was possible [26], most applications obtained an effective number  $\mathcal{N}'$  via numerical fits to search results [28, 27, 30, 31, 32, 24]. Further studies on this topic [33], however, exposed a systematic discrepancy between the family of distributions spanned by Eq. (5.19) and the actual distributions obtained due to template-bank correlations.

An example of this discrepancy is illustrated in Fig. 5.1. We evaluated a template bank containing  $2.23 \times 10^6$  CW templates over frequency and spindown parameters  $(f_0, f_1)$  with a realistic mismatch of  $m = 0.2$  on 7 days of simulated Gaussian noise. We grouped the resulting  $2\mathcal{F}$ -statistic values into batches containing 223 templates each, from which the loudest  $2\mathcal{F}$ -statistic value was retrieved. The effect of parameter-space correlations was tested by either grouping templates within contiguous 5 mHz frequency bands or pooling an equivalent number of templates after randomly shuffling the results.

Shuffling the results before retrieving the loudest value tends to break any contribution from parameter-space correlations, as nearby templates are likely to end up in different batches. The resulting distribution can be properly fitted assuming an uncorrelated template bank. The apparent mismatch between the obtained effective number of templates  $\mathcal{N}' = 218$  and the actual number of independent templates  $\mathcal{N} = 223$  is consistent with the basic claim in [26] about the robustness of Eq. (5.19) with respect to small changes in  $\mathcal{N}$ .

Grouping contiguous frequency bins, on the other hand, produces a distribution out of the scope of

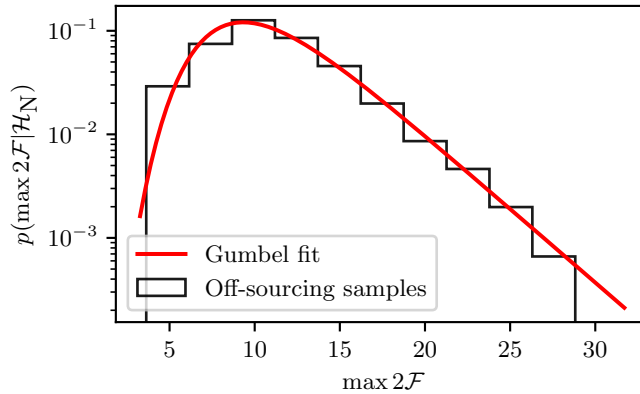


FIGURE 5.2: Distribution of the loudest  $2\mathcal{F}$  of a template bank obtained using 600 off-sourcing evaluations. The template bank corresponds to MCMC samples from a fully-coherent follow-up of a simulated signal in Gaussian noise, in a similar manner to Fig. 1 in [34]. Each histogram entry corresponds to the loudest  $2\mathcal{F}$  retrieved from the template bank evaluated at a different right ascension. The solid line represents the fit of a Gumbel distribution using `scipy.stats.gumbel_r.fit` [77].

Eq. (5.19). This was understood in [34] using extreme value theory (EVT).<sup>1</sup> In the limit of  $\mathcal{N} \rightarrow \infty$ , Eq. (5.19) converges to a max-stable distribution [62, 63, 64, 65, 66], whose functional form is determined by the behaviour of the tail of the distribution  $f$  of the detection statistic. We are primarily interested in the cases when  $f$  is a  $\chi^2$ ,  $\Gamma$  or Gaussian distribution, for all of which  $p(\xi^*|\mathcal{N})$  converges to a Gumbel distribution

$$\text{Gumbel}(x; \mu, \sigma) = \frac{1}{\sigma} \exp \left[ - \left( \frac{x - \mu}{\sigma} \right) - e^{-\left( \frac{x - \mu}{\sigma} \right)} \right] \quad (5.20)$$

where  $\mu$  and  $\sigma$  refer to the location and scale parameters, respectively. Analytical expressions for  $\mu(\mathcal{N})$  and  $\sigma(\mathcal{N})$  for different distributions  $f$  are widely available in the literature [66, 67, 68]. As discussed in Sec. 5.2, the individual  $2\mathcal{F}$  follow a  $\chi_4^2$  distribution on Gaussian noise; the scale parameter of the associated Gumbel distribution is, consequently, fixed to  $\sigma = 2$  *regardless of* the value of  $\mathcal{N}$  [68]. Naively fitting Eq. (5.19) corresponds then to simply adjusting the location of the Gumbel distribution’s peak, as clearly seen in the top panel of Fig. 5.1. The apparent mismatch is resolved if one instead tries to fit *both* the location and scale parameters of the Gumbel distribution to the data.

This solution can be directly applied to computationally cheap searches, such as narrow-band searches [69, 70], directed searches using the Viterbi method [71, 72, 73], or the follow-up of particular outliers [34], using the “off-sourcing” method [74]. The basic idea is that evaluating the same template bank while shifting the sky position away from the outlier will sample a subset of templates uncorrelated to the outlier but with a consistent noise background. Each off-sourced template bank is thus equivalent to a different noise realization. To describe the distribution of the loudest outlier, then, it suffices to evaluate  $10^2 - 10^3$  off-sourced template banks retrieving the loudest outlier of each. A Gumbel distribution can then be fitted to the resulting distribution [34].

Figure 5.2 exemplifies this procedure using a template bank constructed by MCMC sampling as implemented in `PyFstat` [75, 76]. The template bank, containing  $2.5 \times 10^5$  highly correlated templates across frequency, spindown and sky positions, was shifted to 600 different sky positions excluding a  $90^\circ$  wedge around the outlier’s position. The resulting distribution is well described by a Gumbel distribution, with parameters fitted using a standard maximum-likelihood estimation.

Extreme value theory thus allows directly tackling the actual problem posed at the start of this section, namely estimating the distribution of the loudest candidate under the noise hypothesis. The “effective number of independent templates” does not play any major role, as the parameters being fitted are the location and scale of a well-described probability distribution.

<sup>1</sup>We acknowledge previous attempts to apply EVT to the search for CWs [60, 61]. Ref. [34] is the first work presenting a *practical* application of an EVT result improving over previous methods.

## 5.4 How to estimate the distribution of the loudest outlier: an empirical approach

Estimating the loudest candidate’s distribution typically entails fitting an ansatz to a set of samples generated using a numerical procedure. As briefly demonstrated in Sec. 5.3, EVT provides sensible ansätze for this purpose; generating samples, however, quickly becomes a burden for wide parameter-space searches, as template banks are orders of magnitude larger. In such cases, the distribution of the loudest outlier can be estimated using the search results themselves as a proxy for background samples [23, 24, 25].

In this section, we combine the EVT ansatz described in Sec. 5.3 [Eq. (5.20)] with the proposal from [25]. Our new generalized method, **distromax**, covers any sort of detection statistic whose noise-hypothesis distribution falls into one of the three possible max-stable domains of attraction, i.e. not only the standard  $\mathcal{F}$ -statistic, but also “line-robust” statistics [35, 36], generalizations of the  $\mathcal{F}$ -statistic to look for tCW [37]. Other detection statistics used in the CW literature [16], such as Hough number-count [78, 79, 80, 81, 82], cross-correlation [83], or power-based statistics [84, 85, 86], could potentially benefit from **distromax** as well.

### 5.4.1 Basic formulation

We are interested in describing  $p(\xi^*|\mathcal{H}_N)$  solely using the available detection statistic samples from the search  $\xi$ , that is, without any further evaluation of the template bank (e.g. off-sourcing). Following the argument in Sec. 5.3, the evaluation of a detection statistic over a generic template bank can be interpreted as equivalent to the evaluation of said detection statistic over different realizations of noise with a certain (and unknown) degree of correlation. If correlations were negligible, a direct application of Eq. (5.19) would give us the desired answer.

The key realization of [23, 25] is that the loudest outlier from a template bank  $\xi^*$  can be obtained in two steps: estimate the distribution of the loudest candidate of a *smaller* template bank, then extrapolate such distribution to account for the template bank reduction. Dividing the initial template bank into smaller subsets makes multiple loudest candidates available to properly fit a distribution.

For the first step, one splits the dataset  $\xi$ , containing  $\mathcal{N}$  (possibly correlated) values, into  $B$  batches, each of them with  $n = \mathcal{N}/B$  elements. This partition can be done such that each batch contains a similar subset of the overall population so that the per-batch maxima (*batchmax* samples)  $\{\xi_b^*, b = 1, \dots, B\}$  are independent draws from the same unknown distribution:  $\xi_b^* \sim p_n$ . If we choose a sufficiently high number of batches  $B$ , the *batchmax* distribution  $p_n$  can be obtained by fitting a suitable ansatz to the data.

As the second step, the overall loudest value  $\xi^*$  is then

$$\xi^* = \max_{b=1, \dots, B} \xi_b^*, \quad (5.21)$$

which corresponds to the loudest of  $B$   $p_n$ -distributed random variables. This operation was already described in Eq. (5.19), which here we recast as an operator in the space of probability distributions for later convenience: given a probability distribution  $f$ ,  $\text{MaxProp}_B f$  corresponds to the distribution of the loudest candidate over a set of  $B$  independent samples of  $f$ :

$$\text{MaxProp}_B f(x) = Bf(x) \left[ \int_0^x dx' f(x') \right]^{B-1}. \quad (5.22)$$

The distribution of the overall loudest value  $\xi^*$  is then simply<sup>2</sup>

$$p(\xi^*|\mathcal{H}_N) = \text{MaxProp}_B p_n(\xi^*). \quad (5.23)$$

The initial proposal in [25] described the batchmax distribution  $p_n$  using a Gaussian Kernel Density Estimation (KDE) over the set of batchmax samples  $\{\xi_b^*\}$ . Eq. (5.23) was then implemented as a numerical integration. The final decision threshold was based on the support of the resulting distribution.

<sup>2</sup>Note that [23], inserting an empirical histogram as the  $f$  in Eq. (5.22) (to account for a bias due to implementation details in the  $\mathcal{F}$ -statistic, see also Sec. 5.5), constitutes an earlier application of this principle.

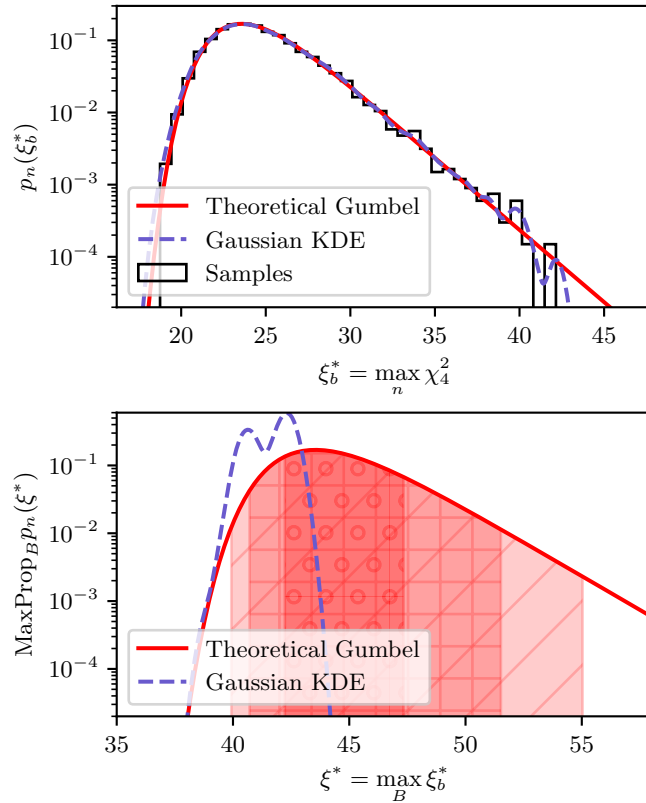


FIGURE 5.3: Upper panel: KDE fit to a set of  $B = 10^4$  samples drawn from the theoretical Gumbel distribution [68] of the loudest sample out of an ensemble of  $n = 10^4$   $\chi_4^2$  random variables. The stepped line corresponds to the histogram of samples. The blue dashed line corresponds to the Gaussian KDE. The red solid line corresponds to the theoretical distribution. Lower panel: Application of the numerical  $\text{MaxProp}_B$  operator with  $B = 10^4$  to the KDE computed from the upper panel (blue dashed line). We compare the result to the theoretical distribution of the maximum sample over  $\mathcal{N} = n \times B = 10^8$   $\chi_4^2$  samples (red solid line). Shaded regions correspond to the 68%, 95%, and 99% probability intervals. KDE bandwidths are estimated using the default method (“scott”) implemented in [77].

We find that the use of Gaussian KDEs introduces inaccuracies into the estimation of  $p(\xi^*|\mathcal{H}_N)$ . The reason is twofold. First, KDEs are prone to overfitting histogram artifacts which arise due to finite sample sizes. This is illustrated in the upper panel of Fig. 5.3. As a result, the propagated distribution in this case displays an unintended bimodality, as shown in the lower panel of Fig. 5.3. Second, for the detection statistics we consider here, the batchmax distribution falls off exponentially (see Appendix 5.A), at a much slower pace than a Gaussian tail. This tends to cause  $\text{MaxProp}_B$  to underestimate the variance of the resulting distribution, as shown in the lower panel of Fig. 5.3.

Our main innovation with `distromax` is to propose a cogent ansatz to circumvent the non-parametric description of the batchmax distribution. Our specific proposal, a max-stable distribution, corresponds to the asymptotic behaviour of the batchmax distribution in the limit of  $n \rightarrow \infty$ . The max-stable property also simplifies the  $\text{MaxProp}_B$  operator into a simple algebraic operation.

### 5.4.2 Introducing `distromax`

Batchmax samples in Eq. (5.21) are constructed so that they correspond to independent and identically distributed random variables from a certain underlying distribution  $p_n$ . In the case of a data stream free of loud disturbances, this can be simply achieved by randomly shuffling the results of a search before grouping them into batches. (A discussion of the effects of shuffling data with loud disturbances is deferred to Secs. 5.4.5 and 5.4.6 and Appendix 5.B.) The batchmax distribution then corresponds

to that of the loudest candidate over  $n$  templates, which, as discussed in Sec. 5.3, tends to a Gumbel distribution as  $n \rightarrow \infty$ . Hence, we propose the following ansatz for the batchmax distribution:

$$p_n(x) = \text{Gumbel}(x; \mu_n, \sigma_n), \quad (5.24)$$

where  $\mu_n, \sigma_n$  are obtained by direct fit to the batchmax samples. This choice is similar to that of [87], which directly fitted an exponential tail (upper tail of a Gumbel distribution) to the batchmax distribution.

EVT distributions, such as Gumbel, are max-stable distributions: the distribution of the loudest outlier from a set of EVT distributions is itself an EVT distribution of the same kind, albeit with different parameter values. As a result, the MaxProp operator can be re-expressed in a closed form in terms of the location and scale parameters of the distribution. Concretely, it is straightforward to show that

$$\text{MaxProp}_B \text{Gumbel}(x; \mu_n, \sigma_n) = \text{Gumbel}(x; \mu_*, \sigma_*) \quad (5.25)$$

where

$$\mu_* = \mu_n + \sigma_n \ln B, \quad (5.26)$$

$$\sigma_* = \sigma_n. \quad (5.27)$$

Thus, the target distribution is readily obtainable through a simple algebraic calculation after performing a fit to the batch-max samples:

$$p(\xi^* | \mathcal{H}_N) = \text{Gumbel}(\xi^*; \mu_*, \sigma_*). \quad (5.28)$$

Summarizing, `distromax` exploits the max-stability of the Gumbel distribution to estimate the distribution of the loudest candidate of a search,  $p(\xi^* | \mathcal{H}_N)$ . To do so, search results are shuffled into  $B$  disjoint batches from which the loudest candidates are retrieved. These  $B$  batchmax candidates, by construction, can be interpreted as draws from an EVT distribution  $p_n$  whose parameters can be estimated using a standard maximum-likelihood fit such as `scipy.rv_continuous.fit` [77].

In broad terms, the batch size  $n$  determines how close the batchmax distribution is to an EVT one, whereas the number of batches  $B$  determines how sharply the parameters of  $p_n$  can be determined. Real searches usually contain a fixed number of templates  $\mathcal{N} = nB$ , meaning a trade-off is required: On the one hand, choosing a large  $n$  (hence a small  $B$ ) produces a small number of samples, each well consistent with an EVT distribution, but increases the variance of the  $p_n$  fit. On the other hand, a large  $B$  (hence a small  $n$ ) produces a big number of samples drawn from a distribution which has not fully converged to an EVT distribution, meaning the estimated parameters may be biased with respect to the actual distribution. We devote the following subsections to further discuss the role played by each of these parameters.

### 5.4.3 The MaxProp operator

We now characterize the phenomenology of the MaxProp operator on a Gumbel distribution, which corresponds to the asymptotic distribution followed by standard (t)CW detection statistics.

Let us consider a template bank with  $\mathcal{N} = 10^6$  templates. Given a batch size  $n$ , we model a batchmax distribution as a Gumbel distribution with  $\sigma_n = 2$  and  $\mu_n = \sigma_n \ln n$ . This is equivalent to considering a detection statistic following a  $\chi_4^2$  distribution which has already converged to its corresponding EVT distribution, preventing finite sample-size effects from polluting the analysis. The ground truth distribution of the loudest candidate from the template bank then corresponds to the propagation of said Gumbel distribution over  $B = 10^6/n$  batches, i.e. a Gumbel distribution with  $\sigma = \sigma_n$  and  $\mu = \sigma \ln \mathcal{N}$ .

Batch sizes  $B \in [1, 10^6]$  are analyzed by drawing  $B$  batchmax samples from the aforementioned batchmax distribution with  $n = 10^6/B$ ;  $\mu_n$  and  $\sigma_n$  are fitted using `scipy.stats.gumbel_r.fit` and propagated using Eq. (5.25). The resulting CDF is compared against the ground truth CDF, shown in Figure 5.4. The relative error in the estimated location and scale parameters is shown in Fig. 5.5.

As previously anticipated, a low number of batches  $B \lesssim 10^3$  results in a greater dispersion of the estimated parameters. As the number of batches reach the  $10^3 \lesssim B \lesssim 10^4$  range, batchmax histograms become more robust and relative parameter deviations achieve sub-percent levels.




FIGURE 5.4: CDF comparison between the ground truth Gumbel distribution and the MaxProp Gumbel distributions using different numbers of batches  $B$ . Each line corresponds to a propagated Gumbel distribution obtained by drawing  $B$  samples from the batchmax distribution (as discussed in the main text), fitting a Gumbel distribution, and applying the MaxProp operator. This procedure is repeated 500 times for each of the 48 selected values of  $B$  between 10 and  $10^6$ .

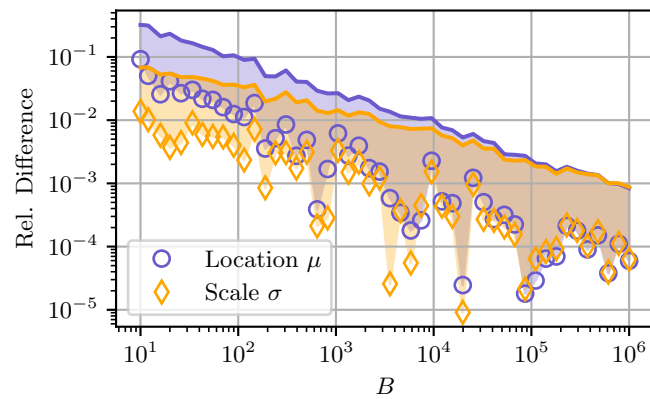


FIGURE 5.5: Relative difference in Gumbel parameters obtained by the application of MaxProp using different numbers of batches  $B$ . Each marker shows the average relative deviation over 500 realizations, corresponding to the distributions shown in Fig. 5.4. The upper envelope represents sample standard deviation. Lower envelopes are omitted due to the logarithmic scale.




FIGURE 5.6: Batchmax distribution of an ensemble of  $n$   $\chi_4^2$  random variables for different choices of  $n$ . Shaded regions represent histograms of  $10^4$  batchmax values obtained by drawing  $10^4 \times n$  values and retrieving the loudest out of each group of  $n$ . Each solid line shows a direct fit of a Gumbel distribution to the data using `scipy.stats.rv_continuous.fit` [77]. Each dashed line represents the corresponding Gumbel distribution using Eqs. (32) and (33) of [68], implemented in the `distromax` package [38].

#### 5.4.4 Characterizing the batchmax distribution

On the other hand, to test convergence of batchmax distributions, we take as an example the case of a  $\chi^2$  distribution with 4 degrees of freedom ( $\Gamma$  distribution with shape parameter  $k = 2$  and scale parameter  $\theta = 2$ ), deferring to Appendix 5.A other generic distributions and further references. For a  $\chi_4^2$  distribution, the limit  $n \rightarrow \infty$  corresponds to a Gumbel distribution [66], as shown in Fig. 5.6.

Given a set of random variables following a specific distribution  $f$ , the convergence of the loudest draw towards an EVT distribution is driven by the behaviour of  $f$ 's tail (*tail-equivalence* [66]). More specifically, the role of the batch size  $n$  is related to how likely it is to draw a sample within the tail of the distribution: the higher the number of samples  $n$ , the more likely it is to retrieve a value from the upper tail, hence the lower the dependency on other details of the distribution's shape. Using a low batch size causes batchmax samples to be dominated by the bulk instead of the tail, keeping the resulting distribution from properly converging to an EVT distribution.

$\chi^2$  random variables are non-negative, as they are the sum of the squares of standard Gaussian random variables. Gumbel distributions, on the other hand, present a double exponential decay in their lower tail [Eq. (5.20)]. Consequently, as shown in the case  $n = 1$  in Fig. 5.6, batchmax distributions with low  $n$  tend not to follow a Gumbel distribution. As the number of samples  $n$  increases, the effects due to the distribution's bulk become milder and we find a better agreement to the expected distribution.

#### 5.4.5 Parameter estimation accuracy and comparison to previous approaches

Finally, we present a more realistic set of results evaluating the  $\mathcal{F}$ -statistic over an actual template bank on 7 days of Gaussian noise using `ComputeFstatistic_v2` [59]. The template bank is constructed using the `gridType=8` option with maximum mismatch  $m = 0.2$ , for a fixed sky position  $(0, 0)$  in equatorial coordinates around  $f_0 = 50$  Hz and  $f_1 = -10^{-8}$  Hz/s, containing  $\mathcal{N} \simeq 8 \times 10^6$  templates. A ground truth distribution is numerically constructed by evaluating this template bank on 900 realizations of Gaussian noise and retrieving the loudest  $2\mathcal{F}$  value from each. `distromax` is then applied to the individual realizations in order to test its accuracy.

We compare two different batching approaches: batching contiguous frequency bins and shuffling the results into random batches. To produce comparable results, the shuffled batches contain the same number of templates  $n$  as each contiguous batch. The motivation behind these two approaches is related to the potential presence of correlated outliers in real detector data: instrumental artifacts tend to affect relatively well-localized frequency bands [13]. Frequency-wise batching could thus prevent very loud outliers from polluting a high number of batches and overestimating the expected loudest outlier. Not shuffling the template bank, however, could require an increase in the batch size to obtain proper






FIGURE 5.7: Comparison of methods to estimate the distribution of the loudest  $\mathcal{F}$ -statistic outlier from a template bank. The data stream and template bank are constructed as explained in the main text. In this figure, the  $\mathcal{N} \simeq 8 \times 10^6$  templates are grouped together by joining 25 consecutive frequency bins or shuffled into  $B = 7236$  batches with a batch size of  $n = 1105$ . Red circles and orange triangles represent **distromax** results with and without shuffling, while blue crosses and light blue plus signs represent the results obtained using a Gaussian KDE again with and without shuffling. Solid lines show the mean and standard deviation of the ground truth distribution. Dotted lines represent one, two and three standard deviations with respect to the ground truth mean.

convergence to a Gumbel distribution, but then the reduced number of batches would imply an increase in the variance of the estimate.

Results are shown in Figs. 5.7, 5.8, and 5.9 in terms of the estimated mean and standard deviation from each method against those of the ground-truth distribution. We also compare to the original proposal of [25] by using a Gaussian KDE to approximate  $p_n$  in Eq. (5.24).

We start by discussing the performance of the Gaussian KDE. The first significant feature is the lack in precision of the estimated parameters, which are over-dispersed regardless of the choice of  $B$  and  $n$ . We also note that the bulk of these results tend to underestimate both the location and scale parameters of the Gumbel distribution with respect to the ground truth. This is related to the shape of the kernel function being used, as previously discussed in Sec. 5.4.1: the tails of a Gaussian distribution fall off more rapidly than those of a Gumbel distribution, yielding a lower mean and standard deviation. The insensitivity of these results to the choice of  $B$  and  $n$  and to shuffling suggests that this particular KDE-based ansatz does not return a reliable estimate of the batchmax distribution.

**distromax** results, on the other hand, return a more consistent picture. Sample variance increases as the number of batches goes down both with and without shuffling. Complementarily, the bias in the estimated parameters reduces as  $n$  increases ( $B$  decreases), although this effect is only significant for the method without shuffling.

We observe a significant bias reduction by shuffling. Randomly shuffling samples results in more homogeneous batches with weaker inner correlations; as a result, batchmax samples are closer to the expected Gumbel distribution, improving the accuracy of the recovered parameters. Also, as previously anticipated, using bigger batches in the non-shuffling case does improve accuracy, although with a significant increment in variance due to the correspondingly lower number of batches.

These features are consistent with the basic phenomenology discussed in Secs. 5.4.3 and 5.4.4. The different regimes in which **distromax** operates depending on  $\mathcal{N}$ ,  $n$ , and  $B$  are summarized in Fig. 5.10.




FIGURE 5.8: Comparison of methods to estimate the distribution of the loudest  $\mathcal{F}$ -statistic outlier from a template bank using the same dataset as in Fig. 5.7. In this figure, the  $\mathcal{N} \simeq 8 \times 10^6$  templates are grouped together by joining 100 consecutive frequency bins or shuffled into  $B = 1801$  batches with a batch size of  $n = 4419$ .




FIGURE 5.9: Comparison of methods to estimate the distribution of the loudest  $\mathcal{F}$ -statistic outlier from a template bank using the same dataset as in Fig. 5.7. In this figure, the  $\mathcal{N} \simeq 8 \times 10^6$  templates are grouped together by joining 3000 consecutive frequency bins or shuffled into  $B = 60$  batches with a batch size of  $n = 132540$ .




FIGURE 5.10: Summary of regimes in which `distromax` operates. Shaded regions represent combinations of  $n$  and  $B$  for which the `distromax` results suffer from low precision (high variance) due to a low number of batches (Sec. 5.4.3) or low accuracy (high bias) due to a low batch size (Sec. 5.4.4).

These values are extracted from the general behaviour of `distromax` observed throughout the tests performed in this section. As a general working principle, `distromax` requires at least  $\mathcal{N} \simeq 10^6$  in order to return a cogent answer; for smaller template banks, on the other hand, off-sourcing as described in Sec. 5.3 requires little computational effort.

### 5.4.6 Discussion

The results above on simulated data demonstrate the overall performance of `distromax` under Gaussian-noise conditions. The concrete output of `distromax` is a simple estimation on the probability of the loudest candidate of a search under the noise hypothesis  $p(\xi^*|\mathcal{H}_N)$ . However, the specific statement to be drawn from  $p(\xi^*|\mathcal{H}_N)$ , such as a threshold choice, is entirely dependent upon the scope of the analysis at hand. For the sake of completeness, we briefly review the assumptions on which the `distromax` method relies, as well as possible consequences of violating them.

First, the distribution of the detection statistic at hand must belong to the domain of attraction of the Gumbel distribution; roughly, this means its probability distribution should be unbounded and decay at a slower rate than a power-law [66]. Nonetheless, as discussed in Appendix 5.A, this method could be easily adapted to detection statistics within a different domain of attraction; in such cases, Eqs. (5.26) and (5.27) would have to be adapted to the corresponding EVT distribution.

Second, the data at hand must be free of strong disturbances. As discussed in Sec. 5.2.2, loud disturbances in a data stream typically translate into parameter-space regions returning enhanced detection statistics with respect to a non-disturbed data stream. The width of the affected region will depend on the characteristics of the disturbance and the template bank, but in general there can be an extended set of templates with correlated response to the disturbance. Attempting to construct a batchmax distribution by shuffling the samples into different batches would result in a distribution shifted towards the right-hand side of the expected Gumbel distribution. Not using shuffling would suppress the effect of disturbances if the resulting associated population of templates was well localized in frequency; the resulting distribution, however, would be less accurate and in particular could still overestimate the Gumbel parameters, as discussed in Figs. 5.7, 5.8, and 5.9. The robustness of `distromax` results to the effect of mild disturbances on the data can also be tested by generating several sets of batchmax samples and comparing the location and scale parameters of the corresponding Gumbel distributions. The wider the distribution over shuffling realizations, the bigger the effect of noise disturbances.

Narrow spectral features (“lines”), in particular, are common noise disturbances affecting (t)CW searches, with excess power typically concentrated within a few frequency bins [13]. CW signals themselves are another typical example of well-localized “disturbances”: should a (strong) CW signal be present in a datastream, a blind application of `distromax` could result in overestimation of the loudest candidate’s distribution, potentially flagging the CW signal itself as a background-noise fluctuation.

The sensitivity of current interferometric detectors, however, makes CW searches to operate in the weak-signal regime (see Appendix 5.C). As a result, CW signals are unlikely to actually affect the estimations provided by `distromax`.

Due to their crucial role in CW searches, we discuss in Appendix 5.B a simple proposal to reduce the effect of narrow-band disturbances so that `distromax` results can still provide a cogent answer. Whenever possible, we recommend the application of informed veto strategies against instrumental artifacts (see [16] and references therein) *before* attempting to process the results using `distromax`. The method discussed in Appendix 5.B is just a complementary algorithm to prevent a specific type of strong disturbances from invalidating an analysis. The characterization and improvement of this or similar algorithms to deal with more generic disturbances is left for future work.

Third, in principle `distromax` assumes the complete set of detection statistic values from the full template bank is available. Several wide parameter-space searches, however, use *toplists* [78, 88, 89], meaning they only keep a small fraction of detection statistic values corresponding to the louder templates. The basic requirement is to use a toplist such that the tail of the distribution is properly represented. Falling short (i. e. not reaching the bulk of the distribution) could result in inaccurate fits to batchmax samples. Incidentally, the results discussed in Appendix 5.B clarify the suitability of `distromax` to these searches.

Fourth, and closely related to the second point, template banks must not be too strongly correlated. EVT ensures `distromax` is robust to a certain degree of template-bank correlations; more specifically, answers provided by `distromax` will be cogent as long as the dominant contribution to batchmax samples comes from the tails of the involved distributions. As discussed in e.g. [48], template bank setups using a small mismatch return highly correlated samples in the vicinity of local parameter-space maxima, adding additional features to the results distribution (see e. g. Fig. 2 in [33]). This may have similar effects to the presence of loud disturbances, as a significant fraction of the resulting batchmax samples would come from samples around a few local maxima and not be representative of the tail of the background distribution. Consequently, the corresponding batchmax distribution will not be fully converged to an EVT distribution and the final Gumbel parameter estimation will be affected. No simple amendment, other than using a higher mismatch [53, 10], is currently available to obtain robust results with `distromax` in this situation.

This phenomenon was observed in [90, 91], where `distromax` was applied to process the result of both a search for CW signals and a search for long-duration transient GWs from glitching pulsars. The CW search used a mismatch of  $m \simeq 0.02$  combined with a toplist, which reduced the effect of such a dense parameter-space converging on the batchmax distribution. Preliminary studies for the long-duration transient GW search using  $m = 0.02$  *without a toplist*, on the other hand, revealed a poor performance of `distromax`; in that case, the solution was to increase the mismatch to  $m = 0.2$ , using a similar setup as in [32].

As we will discuss during Sec. 5.5, `distromax` is suitable to be applied in real-data searches with typical mismatch setups. Moreover, further detection statistics beyond  $2\mathcal{F}$ , such as “line-robust” statistics [35, 36] or tCW search statistics [37], can be processed using the same method. This has the effect of improving the quality of `distromax` results, as these statistics are designed to diminish the effect of noise disturbances on (t)CW searches, providing a cleaner set of batchmax samples.

## 5.5 Application to O2 data

As a demonstration, we apply `distromax` to data of the two advanced LIGO detectors [18] from the O2 observing run [92] in Hanford and Livingston. We study the statistics of results obtained in [32] for narrow-band searches targeting the Vela and Crab pulsars, which experienced glitches on 12th December 2016 and on 27th March 2017 respectively. The template bank for each target was a grid in  $\lambda = (f_0, f_1)$  of size  $N_\lambda \approx 1.15 \times 10^7$ . This template bank was not constructed with a fixed mismatch but an estimate of  $m \approx 0.2$  was given in [32]. Here we consider the results using various  $\mathcal{F}$ -statistic based detection statistics for both CWs and tCWs. For both cases we choose a batch size  $n = 1000$ , and the number of batches as  $B = N_\lambda/n$ . Since this is a narrow-band search, the low number of templates places this particular application of `distromax` at the border of the suitable regime described in Fig. 5.10; the obtained results, however, are not negatively affected by this.




FIGURE 5.11:  $2\mathcal{F}$  values obtained in [32] analysing LIGO O2 data after glitches in the Vela (left panel) and Crab (right panel) pulsars. In each panel, the gray dotted line is the histogram of the samples obtained by the search. The blue dotted line is the expected  $\chi_4^2$  distribution for independent samples. This is propagated using the total number of templates  $N_\lambda \approx 1.15 \times 10^7$  as  $B$  in Eq. (5.22), yielding an estimated distribution for the maximum of  $2\mathcal{F}$  shown by the dashed-dotted blue line. On the other hand using `distromax` we plot the batchmax histogram (using  $n = 1000$ , black solid line) and we fit it with a Gumbel distribution (dashed red line). The propagated distribution is obtained by applying Eq. (5.25) with  $B = N_\lambda/n$  (solid red line).

### 5.5.1 CW detection statistics

We first consider results for CWs of duration four months corresponding to the maximum observation time in [32]. We begin with the standard  $2\mathcal{F}$  as its distribution is well known. As we see from Fig. 5.11, the histogram of the full  $2\mathcal{F}$  results matches well with a standard  $\chi_4^2$  distribution. One can also fit Eq. (5.19) treating  $\mathcal{N}$  as a free parameter, obtaining  $\mathcal{N}' \approx 1$ . This is equivalent to considering each  $2\mathcal{F}$  sample as the trivial maximum of a single draw from a  $\chi_4^2$  distribution. This  $\chi_4^2$  distribution can then be propagated using the total number of templates  $N_\lambda$  as  $B$  in Eq. (5.22), yielding an estimated distribution for the maximum of  $2\mathcal{F}$ , which assumes that the template bank correlations are negligible.<sup>3</sup> We then compare the resulting distribution with the one obtained by the `distromax` method in Fig. 5.11. The two resulting distributions for the maximum agree well.

We also apply the `distromax` method on a different statistic for the CW search, namely the line-robust statistic  $B_{S/GL}$  [35]. This is a Bayes factor derived from the likelihood ratio between the signal hypothesis and the combined hypothesis noise hypothesis of Gaussian noise and lines. The lines are modelled based on the assumption that they look exactly like a signal, but are present in only one detector. The results are shown in Fig. 5.12. Since the underlying distribution of  $B_{S/GL}$  is unknown, one cannot do the equivalent of fitting Eq. (5.19). Nevertheless, we can still apply the `distromax` method, for which the only constraint is that said distribution falls off faster than a power-law (for the case here discussed involving a Gumbel distribution). The details of the exact distribution are not needed. Indeed, the batchmax distribution is well-fitted by a Gumbel distribution. In previous studies using  $B_{S/GL}$  on real data [94, 24, 95, 31, 30, 96, 29, 97], it was used only to improve the robustness of the search against disturbances by using it as the toplist ranking statistic, but the final significance statements were made returning to  $2\mathcal{F}$  because no closed-form distribution was known for  $B_{S/GL}$ . Now with `distromax` we can directly estimate thresholds from the samples allowing for end-to-end analysis using  $B_{S/GL}$ .

### 5.5.2 tCW detection statistics

We now investigate the case of tCWs, which was the main focus of [32]. As briefly mentioned in Sec. 5.2, such signals can be modelled as CWs modulated by a window function dependent on the transient parameters  $\mathcal{T}$ , namely the start time of the transient signal  $t_0$  and its duration  $\tau$ . For this analysis a

<sup>3</sup>As discussed in detail in Sec. 8.7.1 of [23], the result of any fit of Eq. (5.19) to  $\mathcal{F}$ -statistic-based detection statistic samples cannot be directly interpreted as an “effective number of templates” even in the absence of template-bank correlations. This is due to a small upwards implementation bias in the  $\mathcal{F}$ -statistic computation [45, 93]. For the  $\mathcal{N}' \approx 1$  fit to the  $\mathcal{F}$ -statistic samples, this seems to approximately cancel with the effect of template bank correlations.




FIGURE 5.12: Values of the  $\log_{10} B_{S/GL}$  statistic obtained in the same analysis [32] as in Fig. 5.11. Vela results are on the left panel, Crab results are on the right panel. No closed-form expression for the distribution of  $B_{S/GL}$  is known, so only the search samples (dotted gray line), the batchmax histogram (black solid line), the Gumbel fit to it (dashed red line), and the propagated final distribution (solid red line) from `distromax` are shown.

rectangular window function was used. As a detection statistic for tCWs, the  $\mathcal{F}$ -statistic at fixed  $\lambda$  can be maximized over transient parameters [37], thus obtaining  $2\mathcal{F}_{\max} = \max_{\mathcal{T}} 2\mathcal{F}$ . (We use this notation instead of simply  $\max 2\mathcal{F}$  to avoid confusion with the maximum CW detection statistic  $2\mathcal{F}$  over a full template bank.) We expect local correlations to have a more severe impact when using this statistic because it can pick up short-duration non-Gaussianities or simple fluctuations that the CW  $\mathcal{F}$ -statistic would not be susceptible to.

When the `distromax` package shuffles the dataset in the batchmax stage, several batches can be contaminated by the same noise fluctuation, and therefore the batchmax distribution reflects this contamination. The result is a more ragged distribution with peak-like features, as one can see in Fig. 5.13.

To estimate the distribution of the loudest candidate for this dataset, [32] made several simplifying assumptions. While there is no known distribution for  $2\mathcal{F}_{\max}$  that could be directly inserted into Eq. (5.19), its value at each template  $\lambda$  is the maximum of  $2\mathcal{F}$  values over the transient parameters  $\mathcal{T}$ , which individually follow a  $\chi_4^2$  distribution. However, there is a high degree of correlations in the transient parameter space. Hence, fitting Eq. (5.19) to the  $2\mathcal{F}_{\max}$  samples, the result is an “effective number of transient templates”  $\mathcal{N}' \approx 55$  (compared to a nominal number of  $N_{\mathcal{T}} \approx 2 \times 10^6$  transient templates at each  $\lambda$ ).<sup>4</sup> It was then assumed that  $N_{\lambda} \times \mathcal{N}'$  could be interpreted as an “effective number of templates” over the full, non-maximized, parameter space  $(\lambda, \mathcal{T})$ . Consequently, the distribution of the overall loudest was obtained by propagating a  $\chi_4^2$  distribution using  $B = N_{\lambda} \times \mathcal{N}'$  in Eq. (5.22).

The fits to the  $2\mathcal{F}_{\max}$  sample histograms approximately catch the peak of the distribution, but fail to correctly recover the overall shape. On the other hand, with `distromax`, the Gumbel fits to the batchmax samples are noticeably better aligned to both the peaks and the tails of the batchmax histograms than the  $\mathcal{N}'$  based fits are to the full samples histograms. The propagated distributions from both methods still overlap, but their differences are larger than in the CW case.

Again, we apply `distromax` also on an alternative detection statistic for the tCW search, namely the  $B_{tS/G}$  statistic, also derived in [37]. This statistic does not deal with the transient parameters by maximizing over the  $\mathcal{T}$  space, but rather marginalizes over it using a uniform prior. This results in less contamination from disturbances and noise fluctuations. Despite its known better detection efficiency [37], one reason why this detection statistic has not been used in [32] is that its distribution is not analytically known, hence no simple fit of Eq. (5.19) could be done. With the `distromax` method [38], this is no longer a problem. The results are shown in Fig. 5.14. Using the same data sets as before, the batchmax histograms are much smoother and better fit by a Gumbel distribution than their  $2\mathcal{F}_{\max}$  counterparts. This indicates that  $B_{tS/G}$  is a more robust detection statistic than  $2\mathcal{F}_{\max}$  on real data. As for  $B_{S/GL}$  for CWs, with `distromax` it can now also be used as an end-to-end detection statistic.

<sup>4</sup>Here the difference between the fitted “effective” and nominal number of templates is much larger than for the CW  $\mathcal{F}$ -statistic and hence the previously discussed bias is small enough to be ignored.




FIGURE 5.13: Values of  $2\mathcal{F}_{\max}$  obtained in the tCW analysis of [32] on LIGO O2 data after glitches in the Vela (left panel) and Crab (right panel) pulsars. The gray dotted line is the histogram of the samples obtained by the search. There is no known distribution for  $2\mathcal{F}_{\max}$ , but it is a maximum over  $\mathcal{F}$ -statistics over the transient parameters which individually follow a  $\chi_4^2$  distribution. Due to the high degree of correlations in the transient parameter space, in [32] the “effective number of transient templates” was obtained by fitting Eq. (5.19) to the  $2\mathcal{F}_{\max}$  samples, obtaining  $\mathcal{N}' \approx 55$  (dotted blue line). The distribution of the overall loudest was then obtained by propagating the  $\chi_4^2$  distribution using  $B = N_\lambda \times \mathcal{N}'$  in Eq. (5.22) (dash-dotted blue line). This corresponds to treating each sample as a batch with a single element. On the other hand using **distromax** we plot the batchmax histogram ( $n = 1000$ , black solid line) and we fit it with a Gumbel distribution (dashed red line). The propagated distribution is obtained by applying Eq. (5.25) with  $B = N_\lambda/n$  (solid red line).

## 5.6 Conclusion

We have introduced **distromax**, a new method to estimate the distribution of the loudest candidate in a gravitational-wave search. This method culminates a series of developments in the continuous gravitational-wave literature aimed at re-cycling wide parameter-space search results into a proxy distribution for the expectation over different background noise realizations. An implementation of the method is freely available as a Python package [38].

Our specific proposal uses max-stable distributions from extreme value theory to provide a generic approach, applicable to any detection statistic displaying a light-tailed distribution under the noise hypothesis (that is, unbounded and decaying faster than a power-law). This is in contrast with previous approaches based on the  $\mathcal{F}$ -statistic, whose very specific assumptions prevented a successful generalization.

Although we have focused on the case of detection statistics with light-tailed distributions, as that is the standard encountered in CW searches, extensions to other kinds of distributions are possible by using a different family of max-stable distributions.

We have demonstrated the general applicability of **distromax** using both synthetic Gaussian-noise data and the results of a real search on Advanced LIGO O2 data for (transient) continuous gravitational-wave signals from the Vela and Crab pulsars. Results show a significant improvement with respect to previous estimation methods due to the robustness of **distromax** to realistic template-bank correlations.

Additionally, the possibility of using further detection statistics suppressing the effect of lines ( $B_{S/GL}$ ,  $B_{S/GLtL}$ ) or transient instrumental artifacts ( $B_{tS/G}$ ) presents two further advantages for (transient) continuous gravitational-wave searches: first, **distromax** allows us to process the results directly in terms of these more informative statistics; second, the built-in suppression of instrumental features in these statistics itself improves the convergence of batchmax samples to a max-stable distribution, improving the quality of the results provided by **distromax**. This last point also makes plots of the batchmax distribution a useful tool to diagnose the data quality of a specific frequency band using its deviation with respect to the expected max-stable distribution.




FIGURE 5.14: Values of  $\log_{10} B_{\text{ts/G}}$  obtained in the same [32] tCW analysis as in Fig. 5.13. Vela results are on the left panel, Crab results are on the right panel. Here there is no direct fit of Eq. (5.19) because no closed-form expression for the distribution of  $B_{\text{ts/G}}$  is known, and it cannot be easily related to the original  $\mathcal{F}$ -statistic, so only the search samples (dotted gray line), the batchmax histogram (black solid line), the Gumbel fit to it (dashed red line), and the propagated final distribution (solid red line) from `distromax` are shown.

## 5.A Basic results of extreme value theory

Let us consider a set of  $n$  independent and identically distributed random variables  $\{x_1, \dots, x_n\}$  each following a probability distribution  $f$ . These variables can be identified with a detection statistic evaluated on a set of parameter-space templates, with  $f$  corresponding to the detection statistic's distribution under the noise hypothesis. We are interested in describing the probability distribution of the highest detection statistic value (usually referred to as the largest order statistic [62])  $x^* = \max\{x_1, \dots, x_n\}$ , denoted as  $f_n^*$ , in order to evaluate the significance of outliers resulting from a CW search.

A first ansatz for  $f_n^*$  can be constructed by considering the probability of drawing  $x_i = x^*$  for a single  $i \in [1, n]$  and  $x_i < x^*$  for the remaining  $n - 1$  values, taking into account all possible sortings:

$$f_n^*(x) = n f(x) \left[ \int^x dx' f(x') \right]^{n-1}. \quad (5.29)$$

(An alternative derivation of this results is presented in Sec. 5.3.) This approach is sufficient if the probability distribution  $f$  is well understood and different random variables  $x_i$  are independent from one another so that the joint distribution factors into the product of individual distributions. In the case of searches on real data, however, parameter-space correlations cause  $f_n^*$  to deviate from Eq. (5.29) [58, 33, 34].

Extreme Value Theory (EVT) provides asymptotic closed forms for Eq. (5.29) in the limit of  $n \rightarrow \infty$

$$f_n^* \xrightarrow{n \rightarrow \infty} \text{GEV}(\gamma), \quad (5.30)$$

where GEV refers to the generalized extreme value distribution and  $\gamma \in \mathbb{R}$  is referred to as the extreme value index [62, 63, 65, 66]. According to the specific properties of the random variables at hand, GEV distributions can be shifted and rescaled by location and scale parameters,  $\mu_n$  and  $\sigma_n$ . The general dependency on the number of random variables being drawn  $n$  is due to the increased chances of drawing an extreme value as the number of independent trials increases. This is usually referred to as the *trials factor*.

In a practical case, assuming  $n$  so that the convergence is suitable for the application at hand, Eq. (5.31) can be recast into a closed form

$$f_n^*(x) = \text{GEV}(\mu_n + x \sigma_n; \gamma). \quad (5.31)$$

Typical prescriptions for  $\mu_n$  and  $\sigma_n$  [67, 68] tend to be valid for  $n \gtrsim 10^4$ . The GEV distribution has the specific property of being max-stable, meaning that the distribution of the maximum sample out of  $n$  draws from a GEV distribution is again a GEV of the same kind (same  $\gamma$ ).






FIGURE 5.15: CDF comparison of the loudest sample out of  $n$  draws from a uniform (upper panel), exponential (middle panel), and Cauchy (lower panel) distribution to their corresponding generalized extreme value distribution. Different line colors represent different numbers of draws  $n$  over which the maximization was performed.

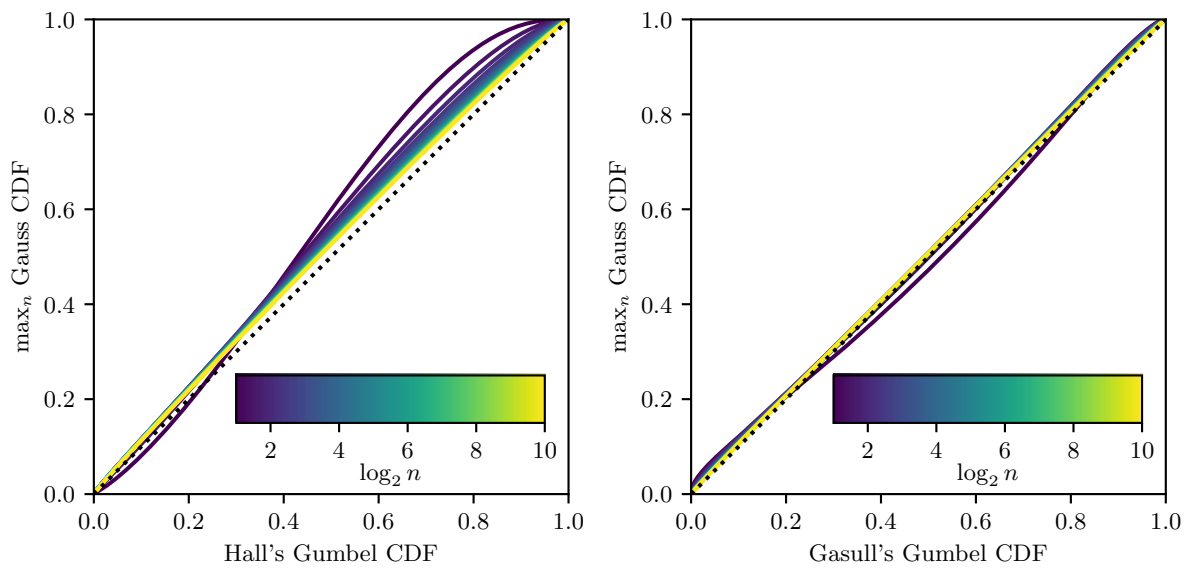


FIGURE 5.16: Comparison of the speed of convergence of the distribution of the loudest sample out of  $n$  draws from a standard Gaussian distribution. The upper panel shows the classical result, derived in [98]. The lower panel shows an improvement later presented in [67], which achieves a lower level of discrepancy than the previous one at the same number of draws. Different line colors represent different numbers of draws  $n$  over which the maximization was performed.

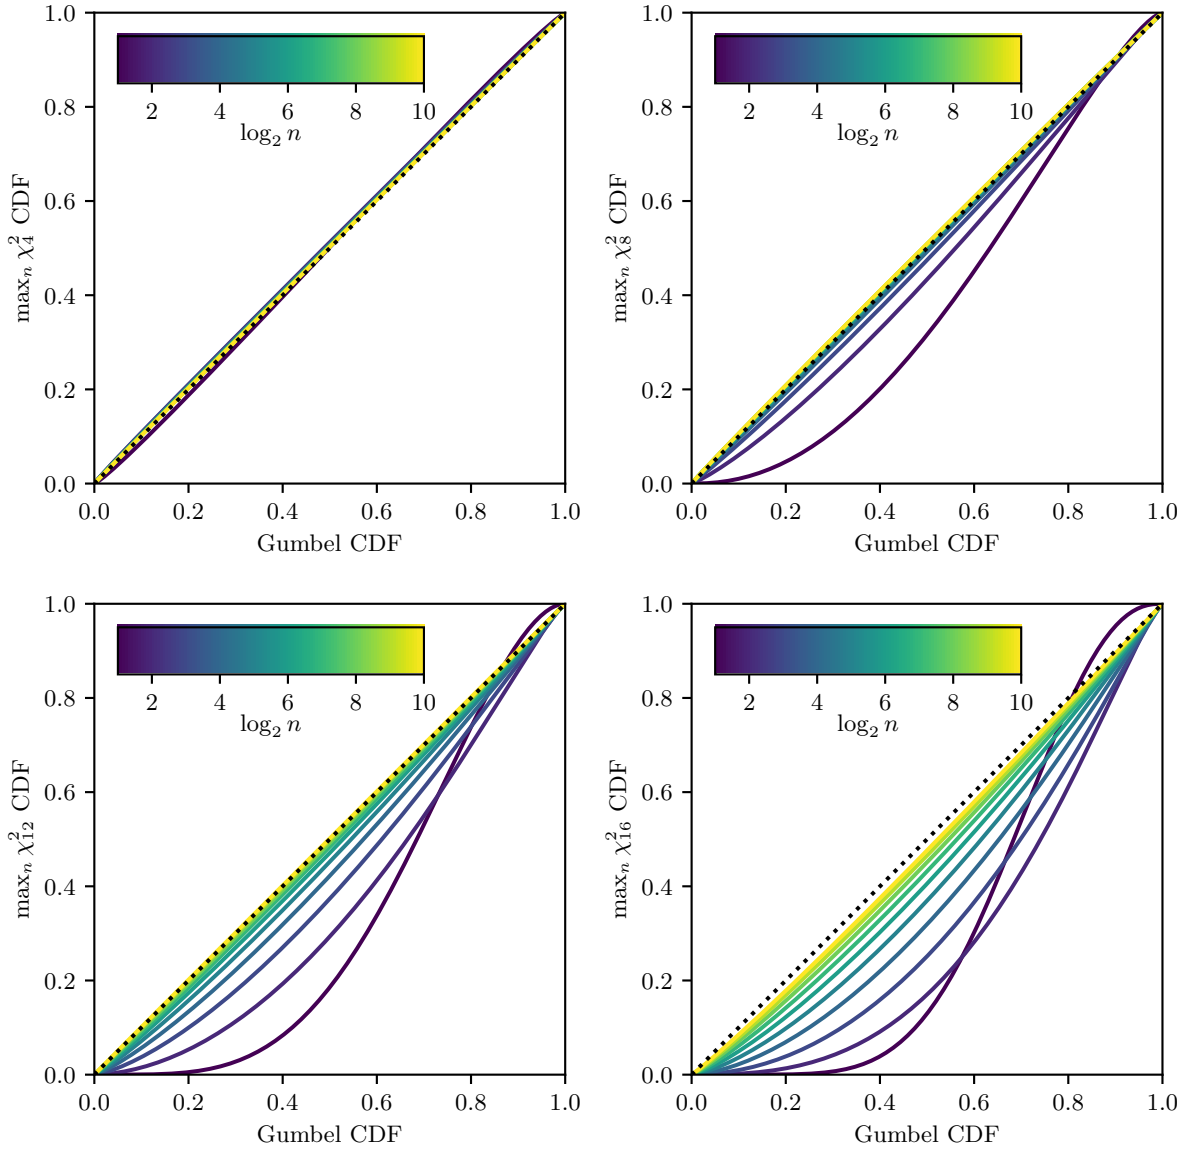


FIGURE 5.17: CDF comparison of the maximum sample out of  $n$  draws from a  $\chi_k^2$ -distributed random variable with different number of degrees of freedom  $k$  to their asymptotic Gumbel distribution. Different line colors represent different numbers of draws  $n$  over which the maximization was performed.

The value of  $\gamma$  depends on the right-hand tail behaviour of  $f$ , and determines the functional form of  $f^*$  out of three possibilities. We follow the definitions given in [66] (where  $\gamma$  is referred to as  $\xi$ ): finite tails with power-law behaviour correspond to  $\gamma < 0$  (Weibull distribution), light tails correspond to  $\gamma = 0$  (Gumbel distribution), and power-law tails correspond to  $\gamma > 0$  (Fréchet distribution)<sup>5</sup>. Fig. 5.15 illustrates the convergence towards each of these families using paradigmatic probability distributions, namely a uniform distribution in  $[0, 1]$ , a standard exponential distribution, and a standard Cauchy distribution.

We focus our attention on the location and scale parameters, as they are relevant in terms of convergence speed. EVT imposes very loose conditions on them, so the choice of  $\mu_n$  and  $\sigma_n$  as functions of  $n$  is not unique for a given distribution, and the main difference across different choices is the speed with which the resulting distribution will approach the GEV one. We illustrate this using a Gaussian distribution, which is in the domain of attraction of the Gumbel distribution and is famous for being quite slow to converge. Figure 5.16 compares the prescription of location and scale parameters originally proposed by Hall [98] to the improvement proposed by Gasull [67].

However, in this paper we are mainly interested in  $\chi_k^2$  distributions, where  $k \in \mathbb{N}$  denotes the degrees of freedom of the distribution, as CW statistics are quite frequently constructed as the norm of a Gaussian vector and hence follow  $\chi_k^2$  distributions. A significant improvement over the classical literature was presented in [68], where closed expressions for  $\mu$  and  $\sigma$  for a generic  $\Gamma$  distribution were obtained. Figure 5.17 shows the convergence of different  $\chi_k^2$  distributions towards a Gumbel distribution.

We provide an implementation of the corresponding expressions discussed in [68] within the `distromax` Python package [38]: Although the `distromax` method itself does not use any of these results (since  $\mu$  and  $\sigma$  are estimated from the data), they can still be used to produce theoretical estimates.

## 5.B Addressing disturbed data

The intended output of `distromax` is an empirical estimation of the distribution of the loudest candidate produced by noise-only data in a CW search  $p(\xi^*|\mathcal{H}_N)$ . To do so, the basic assumption is that the output of a search  $\xi$  corresponds mostly to samples of the detection statistic from a single well-behaved distribution. In practice, this generally means Gaussian noise, plus only a small number of samples coming from another population such as a non-Gaussianity in the data or a CW signal. If the number of such samples is negligible compared to the number of batches used in `distromax`, so would be their effect in the batchmax distribution.

CW searches in real data, however, are populated by various kinds of noise disturbances. Concretely, a prominent type are narrow-band instrumental features (“lines”), which tend to concentrate their effect within a few frequency bins, but especially for higher-dimensional searches (several spin-down terms and/or all-sky searches) can still affect a large number of templates. Probability theory provides the right tools to deal with this situation. Specifically, as discussed in Chapter 21 of [1], one should describe the results of a search  $\xi$  as a mixture of two populations, namely a population of samples belonging to the background and another one belonging to the noise disturbance. Further populations describing additional effects, such as the presence of a CW signal, can also be included in the analysis. The distribution of the loudest candidate produced by the background,  $p(\xi^*|\mathcal{H}_N)$ , would then be obtained by marginalizing out all but the background component of the mixture. Part of the idea of using multiple candidate populations was implemented in [99]. In this Appendix, however, we concentrate on the typical case of loud noise disturbances polluting a small number of frequency bins, for which the ad hoc approach of excising or *notching* the disturbed frequency band returns a similar result to a proper Bayesian analysis. A full treatment of the mixture model problem is left for future work.

As previously discussed in Sec. 5.4, this approach is conceived to deal with a specific set of common noise disturbances so that `distromax` can be applied on a larger range of real-data results. However, users are encouraged to understand and curate their search results using standard CW vetoes (see [16] and references therein) before falling back to this specific notching algorithm.

Leveraging thresholding algorithms from the image-processing literature [100] we propose a simple algorithm capable of notching frequency bands containing prominent disturbances. Since we focus

<sup>5</sup>The `scipy` Python package [77] implements these three distributions under the `stats` module, although it uses a different sign criterion for the extreme value index, therein referred to as  $c$ . Setting  $c = \gamma$ , the Gumbel distribution is `gumbel_r`, the Fréchet distribution is `invweibull`, and the Weibull distribution is `weibull_max`.




FIGURE 5.18: Simulated samples corresponding to  $10^6$  templates in a square grid across  $(f_0, f_1)$ .  $2\mathcal{F}$  background samples correspond to draws from a  $\chi_4^2$  distribution.  $2\mathcal{F}$  outlier samples are drawn from a non-central  $\chi_4^2$  distribution with non-centrality parameter  $\rho^2 = 25$ . This figure shows the third notching iteration. Orange diamonds correspond to the loudest outlier per frequency bin  $\xi^*(f_0)$ . Cyan dots mark samples notched in a previous iteration. Red crosses and the corresponding vertical lines denote the frequency bins being notched in the present iteration. The solid horizontal line corresponds to the threshold computed on the maximum samples using `skimage.filters.threshold_minimum`.

on noise disturbances within well-localized frequency bands, we attempt to flag their corresponding candidates focusing on the loudest detection statistic in each frequency bin  $f_0$

$$\xi^*(f_0) = \max_{\lambda} \xi(f_0, \lambda), \quad (5.32)$$

where  $\lambda$  contains any other relevant parameter-space dimension. The resulting envelope can be thought of as a one-dimensional gray-scale image in which we are interested to discern the background from an object (the polluted band); the distinction is made by properly selecting a gray-value (detection-statistic) threshold such that object pixels (polluted-band samples) lie above it, leaving nothing but background below.

We illustrate the effects of our notching algorithm using a synthetic template bank containing a narrow-band disturbance. The template bank contains  $1000 \times 1000$  templates spanning the  $(f_0, f_1)$  parameter space over the  $[100, 100.1]$  Hz frequency band. The corresponding  $2\mathcal{F}$  is drawn from a  $\chi_4^2$  distribution for each template in the bank. We refer to this  $\chi_4^2$ -drawn set of samples as the ground truth. An outlier is introduced by replacing samples in the  $[100.04, 100.05]$  Hz sub-band with an equal amount of samples drawn from a non-central  $\chi_4^2$  distribution with non-centrality parameter  $\rho^2 = 25$ . The sample projection over the  $f_0$  subspace is shown in Fig. 5.18.

We tested different thresholding techniques, including standard approaches such as the Otsu threshold [101], the minimum cross-entropy threshold [102, 103], and the minimum method threshold [104], using the implementations available in the `skimage` package [105]. We find the minimum method threshold `skimage.filters.threshold_minimum` performs best in our specific study, noting that the implementation of the notching procedure in [38] allows for a flexible selection of thresholding strategies.

Once an appropriate threshold  $\xi_T$  has been established, we proceed to *notch* any frequency bin containing at least one sample above threshold. Specifically, we remove *all* the samples from the frequency bins  $f_0$  where  $\xi^*(f_0) > \xi_T$ . This step can be applied multiple times in order to take care of multiple lines in a band with very different amplitudes, “shoulders” of broad lines, or features such as spectral leakage. The specific implementation provided in [38] implements a simple stopping criterion: notching iterations stop whenever the threshold  $\xi_T$  falls below a pre-specified quantile of  $\xi^*(f_0)$ . The default value, which performs well for our specific example, stops whenever  $\xi_T$  is lower than the top 20% values of  $\xi^*(f_0)$ .

The result of notching, as opposed to simply removing samples over threshold, is illustrated in Fig. 5.19, where the distribution being notched is shown as a blue histogram. Simply removing samples above threshold would be equivalent to *cutting* the tail of the histogram while leaving the bulk untouched. While such an approach would be relatively harmless in the case of a disturbance strong enough to be




FIGURE 5.19: Effects of notching disturbed frequency bins on the background distribution. Each histogram shows a set of  $2\mathcal{F}$  samples akin to that in Fig. 5.18. The black dashed histogram represents ground truth samples in which no disturbance has been included. The blue histogram represents the ground truth samples plus a loud narrow-band disturbance. The red histogram represents the result of notching the blue histogram as discussed in the main text and the caption of Fig. 5.18.




FIGURE 5.20: Distribution of batchmax samples. Each histogram contains the maxima of 5000 batches generated by randomly shuffling samples from Fig. 5.19. Numbers in the legend indicate the location and scale parameters of a Gumbel distribution fitted to each histogram.

cleanly separated from the background, it is rendered ineffective in the case of a relatively mild disturbance, as the polluted-band samples tend to overpopulate the tail of the distribution itself. Completely notching the band, on the other hand (red histogram), properly deals with the overpopulation of outliers and returns a distribution consistent with an undisturbed background.

After notching disturbed bands, we can simply apply the shuffling and batching procedure described in Sec. 5.4.2. The resulting batchmax distributions, including the unnotched and ground truth distributions, are shown in Fig. 5.20. Location and scale parameters of the best fitting Gumbel distributions are shown in the legend and compared to the ground truth distribution in Fig. 5.21.

In this example, the estimated parameters using notching show only 1% relative difference with respect to ground truth parameters, as opposed to the strong bias suffered by the unnotched estimates. Moreover, we remark the robustness of the method to a mild overnotching of non-polluted frequency bins: the convergence to a Gumbel distribution is mainly related to the properties of the “bulk of the tail”; trimming the most extreme events from the background distribution does not affect significantly the fitting of a Gumbel distribution. This is clearly seen in Fig. 5.19, where the notched distribution differs from the ground truth by a few samples. These samples correspond to the background samples over threshold in Fig. 5.18, which belong to the tail of the non-disturbed distribution. As briefly commented in Sec. 5.4.6, this result justifies the extent up to which `distromax` is applicable to toplist-based searches:




FIGURE 5.21: CDF comparison of the batchmax distributions obtained with (red) and without (blue) the application of notching to the ground truth CDF. Each line compares the corresponding CDF of the distribution shown in Fig. 5.20 to the CDF of the ground truth (black histogram). The dashed gray line represents equality.

so long as the top-list reaches the bulk of the distribution, `distromax` should be capable of returning a cogent answer.

The application of more notching iterations than strictly required, however, could result in an underestimation of the Gumbel parameters due to the removal of too many samples in the tail of the distribution. This is particularly important for the scale parameter: underestimations around 5% are often obtained across several realizations of the example setup discussed here when using the notching procedure with the stopping criterion as described above. For a typical batch size of  $B \simeq 10^3 - 10^4$  ( $\ln B \simeq 10$ ) and fiducial values of  $\mu_* = 50$  and  $\sigma_n = 2$  [Eq. (5.26)] (similar to the values encountered in a CW search using the  $2\mathcal{F}$  statistic), a 5% underestimation in  $\sigma_n$  implies about 20% of underestimation in  $\mu_*$ . The main consequence of this is a shift of the resulting  $p(\xi^*|\mathcal{H}_N)$  towards lower values, potentially resulting in an increased number of candidates scoring over the specified threshold.

## 5.C Robustness against injections

As discussed in Sec. 5.4.6, (t)CW signals themselves can be considered as disturbances when trying to estimate a background distribution. We therefore need to test that the `distromax` method is robust to the presence of (t)CW signals and will not be biased upwards, which would lead to picking too high a threshold and missing those signal-related candidates. Such a bias is not entirely avoidable, but should ideally only appear for signals stronger than expected in any practical real-data search situation.

To test this, we re-use the upper limits injections in O2 data from the same analysis [32] as in Sec. 5.5. Simulated signals of increasing amplitude  $h_0$  were added to the original short Fourier Transform (SFT) data (“injections”). For each chosen  $h_0$ , there are 50 data sets with parameters  $\{f_0, f_1, t_0\}$  uniformly distributed over their respective search ranges (see [32]) and the remaining amplitude parameters  $\{\cos \iota, \psi, \phi_0\}$  randomized over their natural ranges. `ComputeFStatistic_v2` [59] was used to reanalyze a small range around the injection point and we combine these results with the original detection statistic samples for the rest of the search band. We then apply `distromax` as in Sec. 5.5, checking the resulting propagated distribution for the maximum of each detection statistic. We also test the notching procedure introduced in Appendix 5.B by gradually increasing the `num_iterations` option,

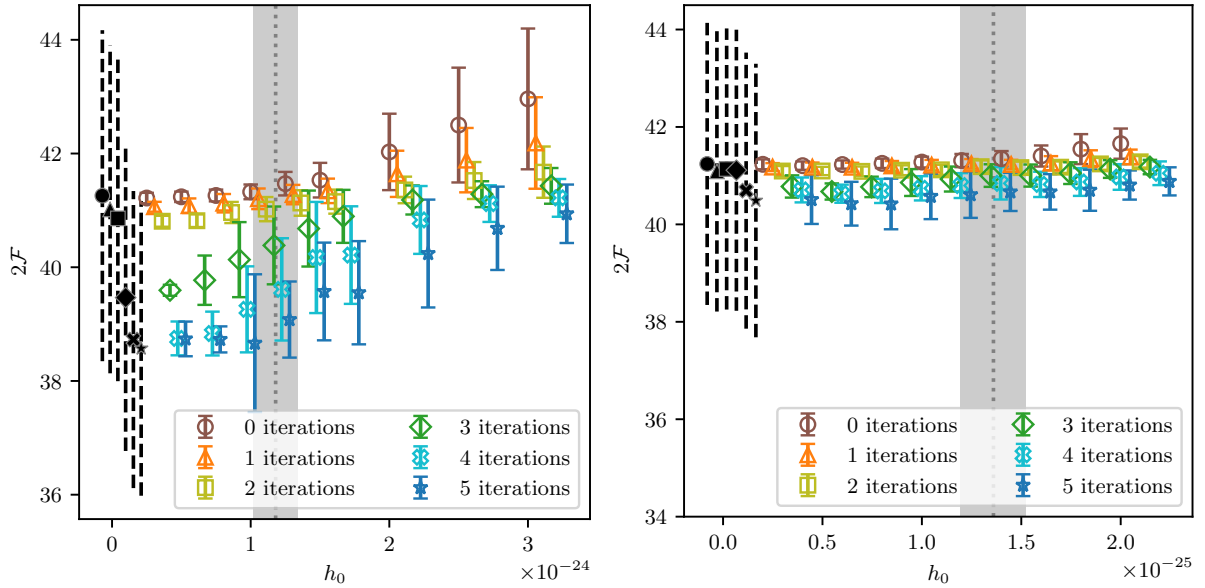


FIGURE 5.22: Results of testing `distromax` on  $2\mathcal{F}$  values from simulated signal injections for the Vela (left) and Crab (right) O2 search parameter spaces, matching the setup from [32], plotted as a function of injected amplitude  $h_0$ . The black data points corresponding to  $h_0 = 0$  are the means of the estimated Gumbel distribution on the original data without injected signals and their dashed vertical lines correspond to the standard deviation of the same distribution. The different markers correspond to different choices of notching iterations. For each injected amplitude, the colored data points with different markers are the average distribution means over 50 injected signals. Their error bars show the standard deviations of these 50 means. A small horizontal shift between the data points belonging to the same  $h_0$  has been inserted for readability. The gray highlight areas are the 90% upper limits taken from [32] for Vela and Crab, considering their uncertainty.

from 0 to 5. The means and standard deviations of the propagated distributions for the four statistics discussed in Sec. 5.5 are shown in Fig. 5.22, Fig. 5.23, Fig. 5.24 and Fig. 5.25.

As one increases the amplitude  $h_0$  of the injected signal, an increasing number of templates will produce elevated values of the detection statistic. In Sec. 5.4.5, we found that shuffling of batches in the `batchmax` step is generally preferred. However, with a strong signal present that affects multiple templates, the shuffled `batchmax` distribution will inevitably become contaminated, leading to an overestimation of the final distribution parameters. Such a trend is indeed visible in the results for all detection statistics, more clearly for the Vela analyses (left-side panels). However, the effect is small compared to the actual increase of the detection statistic at the templates with injections. For the highest amplitudes tested, the detection statistic can reach values  $\sim 20$  times above the expected loudest background sample. Furthermore, the upwards shift in the estimated distribution is generally mild, with the mean shifting by less than one standard deviation of the original injection-free mean, at least as long as  $h_0$  does not reach significantly above the 90% upper limits set in [32].

In addition, the notching feature can be useful in limiting the rise of the estimated means in the presence of signals with large  $h_0$ , by treating the templates with elevated detection statistic as disturbances for the purpose of background estimation and removing them before applying the `batchmax` procedure. Generally 1–2 iterations of notching have little influence on the estimated Gumbel distribution mean of the original data without injections while helping to reduce the rise of the estimated distribution mean with injection  $h_0$ . With more iterations of notching, results become more robust towards strong injections, while for the original data the means in some cases are estimated lower. This would always be conservative in the sense that one would retain more candidates for follow-up even when “over-notching” clean data, however it would lead to additional human and computing effort to follow up candidates that are clearly noise fluctuations.

Hence, as already discussed in Appendix 5.B, we recommend using the notching feature only if

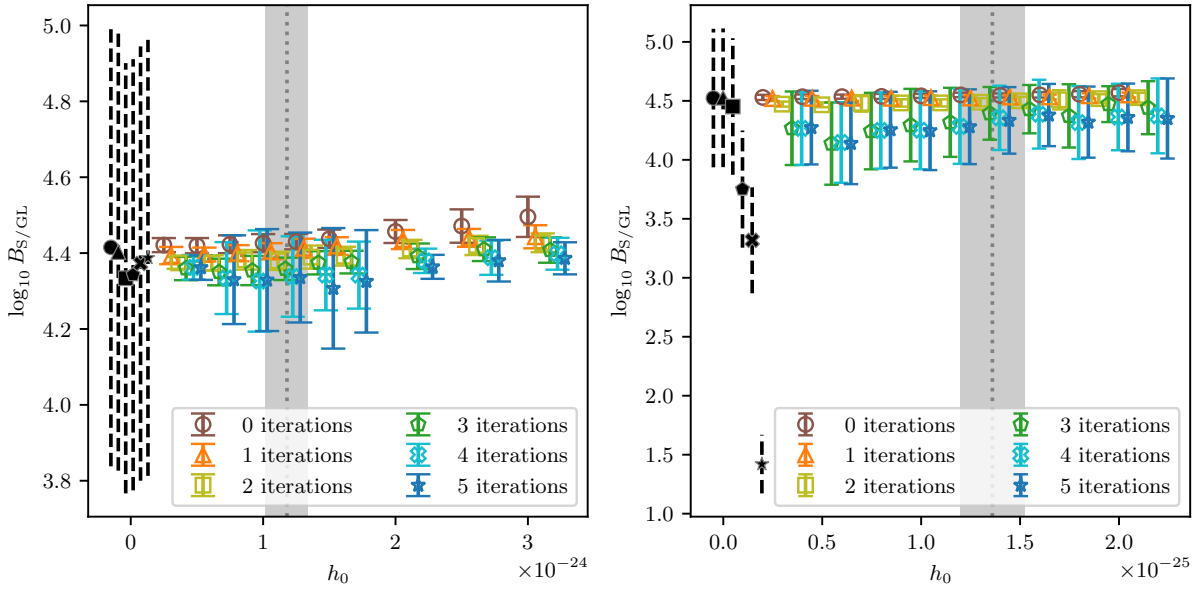


FIGURE 5.23: Results of testing **distromax** on  $\log_{10} B_{S/GL}$  values from simulated signal injections for the Vela (left) and Crab (right) O2 search parameter spaces, with all details as in Fig. 5.22.

required. If data is clean – i.e. no unusual features in the batchmax histograms – notching is not necessary, and more likely leads to underestimated distribution parameters which would correspond to an overly conservative threshold choice. If, on the other hand, data exhibits strong and numerous spikes in the statistic, the plain **distromax** method may lead to overestimated distribution parameters, and notching can be a useful tool in such situations. Regarding the potential presence of (t)CW signals in the data, **distromax** seems robust to these, with or even without notching, for the typical target signals of current (t)CW searches (not standing out far above the noise background); and if in doubt, notching can still help to provide more conservative thresholds.



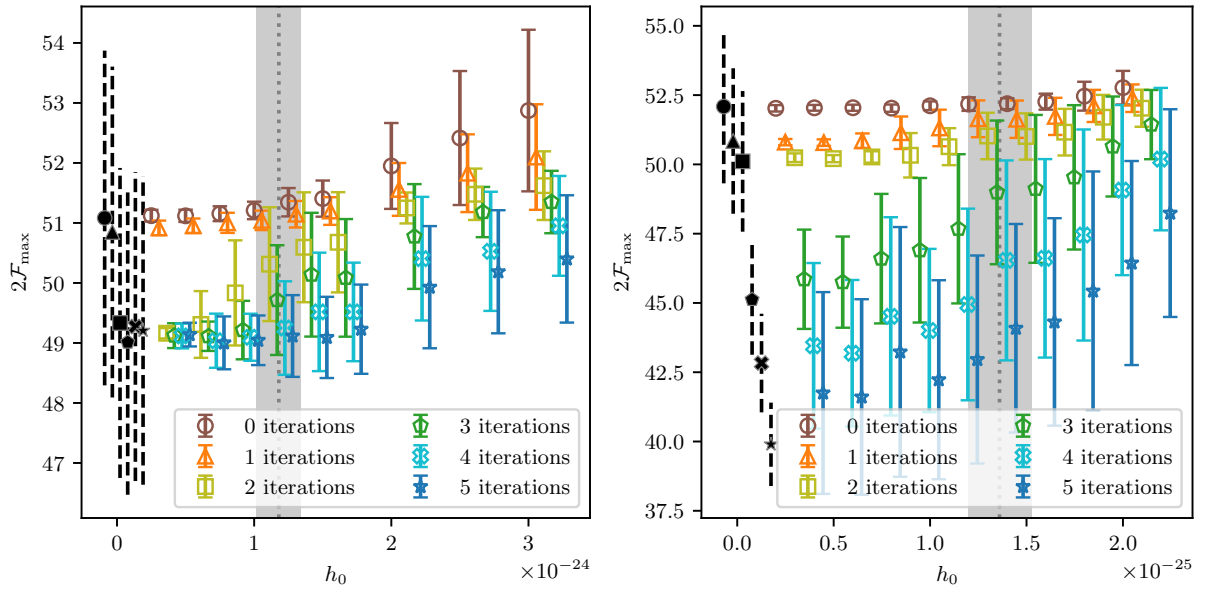


FIGURE 5.24: Results of testing **distromax** on  $2\mathcal{F}_{\max}$  values from simulated signal injections for the Vela (left) and Crab (right) O2 search parameter spaces, with all details as in Fig. 5.22.

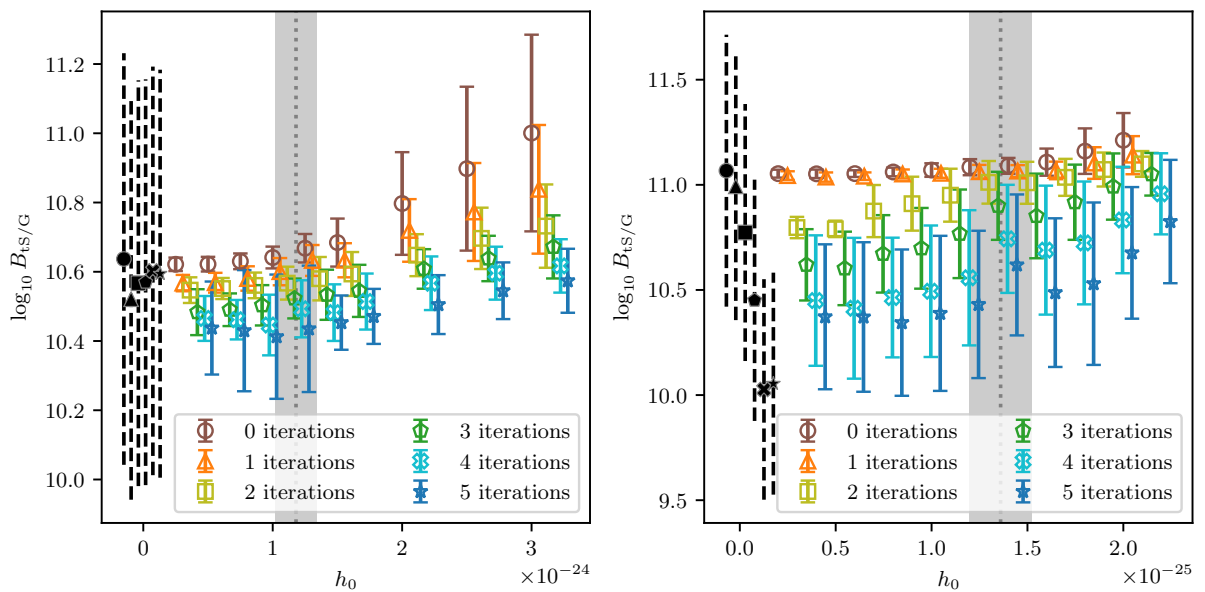


FIGURE 5.25: Results of testing **distromax** on  $\log_{10} B_{\text{ts}/G}$  values from simulated signal injections for the Vela (left) and Crab (right) O2 search parameter spaces, with all details as in Fig. 5.22.



# Bibliography

- [1] E. T. Jaynes, *Probability Theory: The Logic of Science*. Cambridge University Press, 2003.
- [2] B. J. Owen, “Search templates for gravitational waves from inspiraling binaries: Choice of template spacing,” *Phys. Rev. D*, vol. 53, pp. 6749–6761, 1996.
- [3] P. R. Brady and T. Creighton, “Searching for periodic sources with LIGO. 2. Hierarchical searches,” *Phys. Rev. D*, vol. 61, p. 082001, 2000.
- [4] B. S. Sathyaprakash and S. V. Dhurandhar, “Choice of filters for the detection of gravitational waves from coalescing binaries,” *Phys. Rev. D*, vol. 44, pp. 3819–3834, 1991.
- [5] B. J. Owen and B. S. Sathyaprakash, “Matched filtering of gravitational waves from inspiraling compact binaries: Computational cost and template placement,” *Phys. Rev. D*, vol. 60, p. 022002, 1999.
- [6] B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, and J. D. E. Creighton, “FINDCHIRP: An Algorithm for detection of gravitational waves from inspiraling compact binaries,” *Phys. Rev. D*, vol. 85, p. 122006, 2012.
- [7] R. Prix, “Template-based searches for gravitational waves: Efficient lattice covering of flat parameter spaces,” *Class. Quant. Grav.*, vol. 24, pp. S481–S490, 2007.
- [8] I. W. Harry, B. Allen, and B. S. Sathyaprakash, “A Stochastic template placement algorithm for gravitational wave data analysis,” *Phys. Rev. D*, vol. 80, p. 104014, 2009.
- [9] C. Messenger, R. Prix, and M. A. Papa, “Random template banks and relaxed lattice coverings,” *Phys. Rev. D*, vol. 79, p. 104017, 2009.
- [10] B. Allen, “Optimal template banks,” *Phys. Rev. D*, vol. 104, no. 4, p. 042005, 2021.
- [11] K. J. Wagner, J. T. Whelan, J. K. Wofford, and K. Wette, “Template Lattices for a Cross-Correlation Search for Gravitational Waves from Scorpius X-1.” arXiv:2106.16142 [gr-qc].
- [12] R. Prix and B. Krishnan, “Targeted search for continuous gravitational waves: Bayesian versus maximum-likelihood statistics,” *Class. Quant. Grav.*, vol. 26, p. 204013, 2009.
- [13] P. B. Covas *et al.*, “Identification and mitigation of narrow spectral artifacts that degrade searches for persistent gravitational waves in the first two observing runs of Advanced LIGO,” *Phys. Rev. D*, vol. 97, p. 082002, Apr. 2018.
- [14] D. Davis, L. V. White, and P. R. Saulson, “Utilizing aLIGO Glitch Classifications to Validate Gravitational-Wave Candidates,” *Class. Quant. Grav.*, vol. 37, no. 14, p. 145001, 2020.
- [15] B. P. Abbott *et al.*, “Observation of Gravitational Waves from a Binary Black Hole Merger,” *Phys. Rev. Lett.*, vol. 116, no. 6, p. 061102, 2016.
- [16] R. Tenorio, D. Keitel, and A. M. Sintes, “Search Methods for Continuous Gravitational-Wave Signals from Unknown Sources in the Advanced-Detector Era,” *Universe*, vol. 7, no. 12, p. 474, 2021.
- [17] M. Sieniawska and M. Bejger, “Continuous gravitational waves from neutron stars: current status and prospects,” *Universe*, vol. 5, no. 11, p. 217, 2019.
- [18] J. Aasi *et al.*, “Advanced LIGO,” *Class. Quant. Grav.*, vol. 32, p. 074001, mar 2015.

- [19] F. Acernese *et al.*, “Advanced Virgo: a second-generation interferometric gravitational wave detector,” *Class. Quant. Grav.*, vol. 32, p. 024001, dec 2014.
- [20] T. Akutsu *et al.*, “KAGRA: 2.5 Generation Interferometric Gravitational Wave Detector,” *Nat. Astron.*, vol. 3, no. 1, pp. 35–40, 2019.
- [21] P. Jaranowski, A. Królak, and B. F. Schutz, “Data analysis of gravitational-wave signals from spinning neutron stars: The signal and its detection,” *Phys. Rev. D*, vol. 58, p. 063001, Aug 1998.
- [22] C. Cutler and B. F. Schutz, “Generalized  $\mathcal{F}$ -statistic: Multiple detectors and multiple gravitational wave pulsars,” *Phys. Rev. D*, vol. 72, p. 063006, Sep 2005.
- [23] K. W. Wette, *Gravitational waves from accreting neutron stars and Cassiopeia A*. PhD thesis, Australian Natl. U., Canberra, 2009.
- [24] M. A. Papa, J. Ming, E. V. Gotthelf, B. Allen, R. Prix, V. Dergachev, H.-B. Eggenstein, A. Singh, and S. J. Zhu, “Search for Continuous Gravitational Waves from the Central Compact Objects in Supernova Remnants Cassiopeia A, Vela Jr., and G347.3–0.5,” *Astrophys. J.*, vol. 897, no. 1, p. 22, 2020.
- [25] K. Wette, L. Dunn, P. Clearwater, and A. Melatos, “Deep exploration for continuous gravitational waves at 171–172 Hz in LIGO second observing run data,” *Phys. Rev. D*, vol. 103, no. 8, p. 083020, 2021.
- [26] J. Abadie *et al.*, “First search for gravitational waves from the youngest known neutron star,” *Astrophys. J.*, vol. 722, pp. 1504–1513, 2010.
- [27] J. Aasi *et al.*, “Directed search for continuous gravitational waves from the Galactic center,” *Phys. Rev. D*, vol. 88, no. 10, p. 102002, 2013.
- [28] B. Behnke, M. A. Papa, and R. Prix, “Postprocessing methods used in the search for continuous gravitational-wave signals from the galactic center,” *Phys. Rev. D*, vol. 91, p. 064007, Mar 2015.
- [29] B. P. Abbott *et al.*, “Results of the deepest all-sky survey for continuous gravitational waves on LIGO S6 data running on the Einstein@Home volunteer distributed computing project,” *Phys. Rev. D*, vol. 94, no. 10, p. 102002, 2016.
- [30] S. J. Zhu, M. A. Papa, H.-B. Eggenstein, R. Prix, K. Wette, B. Allen, O. Bock, D. Keitel, B. Krishnan, B. Machenschalk, M. Shaltev, and X. Siemens, “An Einstein@home search for continuous gravitational waves from Cassiopeia A,” *Phys. Rev. D*, vol. 94, p. 082008, 2016.
- [31] B. P. Abbott *et al.*, “First low-frequency Einstein@Home all-sky search for continuous gravitational waves in Advanced LIGO data,” *Phys. Rev. D*, vol. 96, no. 12, p. 122004, 2017.
- [32] D. Keitel, G. Woan, M. Pitkin, C. Schumacher, B. Pearlstone, K. Riles, A. G. Lyne, J. Palfreyman, B. Stappers, and P. Weltevrede, “First search for long-duration transient gravitational waves after glitches in the Vela and Crab pulsars,” *Phys. Rev. D*, vol. 100, p. 064058, Sept. 2019.
- [33] C. Dreissigacker, R. Prix, and K. Wette, “Fast and Accurate Sensitivity Estimation for Continuous-Gravitational-Wave Searches,” *Phys. Rev. D*, vol. 98, no. 8, p. 084058, 2018.
- [34] R. Tenorio, D. Keitel, and A. M. Sintes, “Application of a hierarchical MCMC follow-up to Advanced LIGO continuous gravitational-wave candidates,” *Phys. Rev. D*, vol. 104, no. 8, p. 084012, 2021.
- [35] D. Keitel, R. Prix, M. A. Papa, P. Leaci, and M. Siddiqi, “Search for continuous gravitational waves: Improving robustness versus instrumental artifacts,” *Phys. Rev. D*, vol. 89, no. 6, p. 064023, 2014.
- [36] D. Keitel, “Robust semicoherent searches for continuous gravitational waves with noise and signal models including hours to days long transients,” *Phys. Rev. D*, vol. 93, no. 8, p. 084024, 2016.
- [37] R. Prix, S. Giampanis, and C. Messenger, “Search method for long-duration gravitational-wave transients from neutron stars,” *Phys. Rev. D*, vol. 84, p. 023007, 2011.

- [38] R. Tenorio, L. M. Modafferi, D. Keitel, and A. M. Sintes, “`distromax`: empirically estimating the distribution of the loudest candidate from a gravitational-wave search.” <https://github.com/Rodrigo-Tenorio/distromax>, 2021.
- [39] L. Sancho de la Jordana and A. M. Sintes, “A  $\chi^2$  veto for continuous wave searches,” *Class. Quant. Grav.*, vol. 25, p. 184014, 2008.
- [40] P. Leaci, “Methods to filter out spurious disturbances in continuous-wave searches from gravitational-wave detectors,” *Phys. Scr.*, vol. 90, no. 12, p. 125001, 2015.
- [41] S. J. Zhu, M. A. Papa, and S. Walsh, “New veto for continuous gravitational wave searches,” *Phys. Rev. D*, vol. 96, no. 12, p. 124007, 2017.
- [42] J. T. Whelan, R. Prix, C. J. Cutler, and J. L. Willis, “New Coordinates for the Amplitude Parameter Space of Continuous Gravitational Waves,” *Class. Quant. Grav.*, vol. 31, p. 065002, 2014.
- [43] S. Dhurandhar, B. Krishnan, and J. L. Willis, “Marginalizing the likelihood function for modeled gravitational wave searches.” arXiv:1707.08163 [gr-qc].
- [44] J. J. Bero and J. T. Whelan, “An Analytic Approximation to the Bayesian Detection Statistic for Continuous Gravitational Waves,” *Class. Quant. Grav.*, vol. 36, no. 1, p. 015013, 2019. [Erratum: *Class.Quant.Grav.* 36, 049601 (2019)].
- [45] Prix, Reinhard, “The F-statistic and its implementation in ComputeFStatistic\_v2.” <https://dcc.ligo.org/LIGO-T0900149/public>, 2018.
- [46] L. S. Finn, “Detection, measurement and gravitational radiation,” *Phys. Rev. D*, vol. 46, pp. 5236–5249, 1992.
- [47] R. Tenorio, D. Keitel, and A. M. Sintes, “Time-frequency track distance for comparing continuous gravitational wave signals,” *Phys. Rev. D*, vol. 103, no. 6, p. 064053, 2021.
- [48] R. Prix, “Search for continuous gravitational waves: Metric of the multi-detector F-statistic,” *Phys. Rev. D*, vol. 75, p. 023004, 2007. [Erratum: *Phys. Rev. D* 75, 069901(E) (2007)].
- [49] H. J. Pletsch, “Parameter-space metric of semicoherent searches for continuous gravitational waves,” *Phys. Rev. D*, vol. 82, p. 042002, 2010.
- [50] K. Wette and R. Prix, “Flat parameter-space metric for all-sky searches for gravitational-wave pulsars,” *Phys. Rev. D*, vol. 88, no. 12, p. 123005, 2013.
- [51] K. Wette, “Parameter-space metric for all-sky semicoherent searches for gravitational-wave pulsars,” *Phys. Rev. D*, vol. 92, no. 8, p. 082003, 2015.
- [52] P. Leaci and R. Prix, “Directed searches for continuous gravitational waves from binary systems: parameter-space metrics and optimal Scorpius X-1 sensitivity,” *Phys. Rev. D*, vol. 91, no. 10, p. 102003, 2015.
- [53] B. Allen, “Spherical ansatz for parameter-space metrics,” *Phys. Rev. D*, vol. 100, no. 12, p. 124004, 2019.
- [54] K. Wette, “Lattice template placement for coherent all-sky searches for gravitational-wave pulsars,” *Phys. Rev. D*, vol. 90, no. 12, p. 122010, 2014.
- [55] B. Allen and A. A. Shoom, “Template banks based on Zn and An\* lattices,” *Phys. Rev. D*, vol. 104, no. 12, p. 122007, 2021.
- [56] K. Wette, “Empirically extending the range of validity of parameter-space metrics for all-sky searches for gravitational-wave pulsars,” *Phys. Rev. D*, vol. 94, no. 12, p. 122002, 2016.
- [57] O. Piccinni, “Mitigation of transient disturbances in wide parameter space searches for continuous gravitational wave signals,” Master’s thesis, Università di Roma La Sapienza, 2014.

- [58] K. Wette, “Estimating the sensitivity of wide-parameter-space searches for gravitational-wave pulsars,” *Phys. Rev. D*, vol. 85, p. 042003, 2012.
- [59] LIGO Scientific Collaboration, “LIGO Algorithm Library - LALSuite.” free software (GPL), 2018.
- [60] C. Rover, C. Messenger, and R. Prix, “Bayesian versus frequentist upper limits,” in *PHYSTAT 2011*, (Geneva), pp. 158–163, CERN, 2011.
- [61] S. Suvorova, P. Clearwater, A. Melatos, L. Sun, W. Moran, and R. J. Evans, “Hidden Markov model tracking of continuous gravitational waves from a binary neutron star with wandering spin. II. Binary orbital phase tracking,” *Phys. Rev. D*, vol. 96, p. 102006, Nov. 2017.
- [62] M. Leadbetter, G. Lindgren, and H. Rootzen, *Extremes and Related Properties of Random Sequences and Processes*. Springer Series in Statistics, Springer New York, 1983.
- [63] S. Coles, *An introduction to statistical modeling of extreme values*. Springer Series in Statistics, London: Springer-Verlag, 2001.
- [64] J. Beirlant, Y. Goegebeur, J. Segers, J. Teugels, D. De Waal, and C. Ferro, *Statistics of Extremes: Theory and Applications*. Wiley Series in Probability and Statistics, Wiley, 2004.
- [65] L. de Haan and A. Ferreira, *Extreme Value Theory: An Introduction*. Springer Series in Operations Research and Financial Engineering, Springer New York, 2006.
- [66] P. Embrechts, C. Klüppelberg, and T. Mikosch, *Modelling Extremal Events: for Insurance and Finance*. Stochastic Modelling and Applied Probability, Springer Berlin Heidelberg, 2013.
- [67] A. Gasull, M. Jolis, and F. Utzet, “On the norming constants for normal maxima,” *J. Math. Anal. Appl.*, vol. 422, no. 1, pp. 376–396, 2015.
- [68] A. Gasull, J. López-Salcedo, and F. Utzet, “Maxima of Gamma random variables and other Weibull-like distributions and the Lambert W function,” *TEST*, vol. 24, pp. 714–733, December 2015.
- [69] B. P. Abbott *et al.*, “First narrow-band search for continuous gravitational waves from known pulsars in advanced detector data,” *Phys. Rev. D*, vol. 96, no. 12, p. 122006, 2017. [Erratum: *Phys.Rev.D* 97, 129903 (2018)].
- [70] B. P. Abbott *et al.*, “Narrow-band search for gravitational waves from known pulsars using the second LIGO observing run,” *Phys. Rev. D*, vol. 99, no. 12, p. 122002, 2019.
- [71] H. Middleton, P. Clearwater, A. Melatos, and L. Dunn, “Search for gravitational waves from five low mass X-ray binaries in the second Advanced LIGO observing run with an improved hidden Markov model,” *Phys. Rev. D*, vol. 102, no. 2, p. 023006, 2020.
- [72] D. Jones and L. Sun, “Search for continuous gravitational waves from Fomalhaut b in the second Advanced LIGO observing run with a hidden Markov model,” *Phys. Rev. D*, vol. 103, no. 2, p. 023020, 2021.
- [73] D. Beniwal, P. Clearwater, L. Dunn, A. Melatos, and D. Ottaway, “Search for continuous gravitational waves from ten H.E.S.S. sources using a hidden Markov model,” *Phys. Rev. D*, vol. 103, no. 8, p. 083009, 2021.
- [74] M. Isi, S. Mastroianni, M. Pitkin, and O. J. Piccinni, “Establishing the significance of continuous gravitational-wave detections from known pulsars,” *Phys. Rev. D*, vol. 102, no. 12, p. 123027, 2020.
- [75] G. Ashton and R. Prix, “Hierarchical multistage MCMC follow-up of continuous gravitational wave candidates,” *Phys. Rev. D*, vol. 97, no. 10, p. 103020, 2018.
- [76] D. Keitel, R. Tenorio, G. Ashton, and R. Prix, “Pyfstat: a python package for continuous gravitational-wave data analysis,” *J. Open Source Softw.*, vol. 6, no. 60, p. 3000, 2021.

- [77] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, *et al.*, “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python,” *Nat. Methods*, vol. 17, pp. 261–272, 2020.
- [78] B. Krishnan, A. M. Sintes, M. A. Papa, B. F. Schutz, S. Frasca, and C. Palomba, “Hough transform search for continuous gravitational waves,” *Phys. Rev. D*, vol. 70, p. 082001, Oct 2004.
- [79] P. Astone, A. Colla, S. D’Antonio, S. Frasca, and C. Palomba, “Method for all-sky searches of continuous gravitational wave signals using the frequency-Hough transform,” *Phys. Rev. D*, vol. 90, p. 042002, Aug. 2014.
- [80] A. Miller, P. Astone, S. D’Antonio, S. Frasca, G. Intini, I. La Rosa, P. Leaci, S. Mastrogiovanni, F. Muciaccia, C. Palomba, O. J. Piccinni, A. Singhal, and B. F. Whiting, “Method to search for long duration gravitational wave transients from isolated neutron stars using the generalized frequency-Hough transform,” *Phys. Rev. D*, vol. 98, no. 10, p. 102004, 2018.
- [81] M. Oliver, D. Keitel, and A. M. Sintes, “Adaptive transient Hough method for long-duration gravitational wave transients,” *Phys. Rev. D*, vol. 99, no. 10, p. 104067, 2019.
- [82] P. B. Covas and A. M. Sintes, “New method to search for continuous gravitational waves from unknown neutron stars in binary systems,” *Phys. Rev. D*, vol. 99, no. 12, p. 124019, 2019.
- [83] J. T. Whelan, S. Sundaesan, Y. Zhang, and P. Peiris, “Model-Based Cross-Correlation Search for Gravitational Waves from Scorpius X-1,” *Phys. Rev. D*, vol. 91, p. 102005, 2015.
- [84] V. Dergachev, “On blind searches for noise dominated signals: a loosely coherent approach,” *Class. Quant. Grav.*, vol. 27, p. 205017, 2010.
- [85] E. Goetz and K. Riles, “An all-sky search algorithm for continuous gravitational waves from spinning neutron stars in binary systems,” *Class. Quant. Grav.*, vol. 28, p. 215006, 2011.
- [86] V. Dergachev and M. A. Papa, “Sensitivity improvements in the search for periodic gravitational waves using O1 LIGO data,” *Phys. Rev. Lett.*, vol. 123, no. 10, p. 101101, 2019.
- [87] R. Abbott *et al.*, “Search for continuous gravitational waves from 20 accreting millisecond X-ray pulsars in O3 LIGO data,” *Phys. Rev. D*, vol. 105, p. 022002, Jan 2022.
- [88] H. J. Pletsch and B. Allen, “Exploiting global correlations to detect continuous gravitational waves,” *Phys. Rev. Lett.*, vol. 103, p. 181102, 2009.
- [89] K. Wette, S. Walsh, R. Prix, and M. A. Papa, “Implementing a semicoherent search for continuous gravitational waves using optimally-constructed template banks,” *Phys. Rev. D*, vol. 97, no. 12, p. 123016, 2018.
- [90] R. Abbott *et al.*, “Narrowband searches for continuous and long-duration transient gravitational waves from known pulsars in the LIGO-Virgo third observing run.” arXiv:2112.10990 [gr-qc].
- [91] L. M. Modafferi, J. Moragues, and D. Keitel, “Search setup for long-duration transient gravitational waves from glitching pulsars during LIGO-Virgo third observing run,”
- [92] R. Abbott *et al.*, “Open data from the first and second observing runs of advanced ligo and advanced virgo,” *SoftwareX*, vol. 13, p. 100658, 2021.
- [93] Prix, Reinhard, “F-statistic bias due to noise-estimator.” <https://dcc.ligo.org/LIGO-T1100551/public>, 2006.
- [94] B. Steltner, M. A. Papa, H. B. Eggenstein, B. Allen, V. Dergachev, R. Prix, B. Machenschalk, S. Walsh, S. J. Zhu, and S. Kwang, “Einstein@Home All-sky Search for Continuous Gravitational Waves in LIGO O2 Public Data,” *Astrophys. J.*, vol. 909, no. 1, p. 79, 2021.
- [95] J. Ming, M. A. Papa, A. Singh, H. B. Eggenstein, S. J. Zhu, V. Dergachev, Y. Hu, R. Prix, B. Machenschalk, C. Beer, O. Behnke, and B. Allen, “Results from an Einstein@Home search for continuous gravitational waves from Cassiopeia A, Vela Jr. and G347.3,” *Phys. Rev. D*, vol. 100, no. 2, p. 024063, 2019.

- [96] M. A. Papa, H. B. Eggenstein, S. Walsh, I. DiPalma, B. Allen, P. Astone, O. Bock, T. D. Creighton, D. Keitel, B. Machenschalk, R. Prix, X. Siemens, A. Singh, S. J. Zhu, and B. F. Schutz, “Hierarchical follow-up of subthreshold candidates of an all-sky Einstein@Home search for continuous gravitational waves on LIGO sixth science run data,” *Phys. Rev. D*, vol. 94, no. 12, p. 122006, 2016.
- [97] J. Ming, M. A. Papa, H.-B. Eggenstein, B. Machenschalk, B. Steltner, R. Prix, *et al.*, “Results from an einstein@home search for continuous gravitational waves from g347.3 at low frequencies in LIGO o2 data,” *Astrophys. J.*, vol. 925, p. 8, jan 2022.
- [98] P. Hall, “On the rate of convergence of normal extremes,” *Journal of Applied Probability*, vol. 16, no. 2, pp. 433–439, 1979.
- [99] M. F. Bennett, A. Melatos, A. Delaigle, and P. Hall, “Reanalysis of  $\mathcal{F}$ -statistic gravitational-wave searches with the higher criticism statistic,” *Astrophys. J.*, vol. 766, p. 99, 2013.
- [100] M. Sezgin and B. Sankur, “Survey over image thresholding techniques and quantitative performance evaluation,” *J. Electron. Imaging*, vol. 13, no. 1, pp. 146–165, 2004.
- [101] N. Otsu, “A threshold selection method from gray-level histograms,” *IEEE Trans. Syst. Man Cybern.*, vol. 9, no. 1, pp. 62–66, 1979.
- [102] C. Li and C. Lee, “Minimum cross entropy thresholding,” *Pattern Recognit.*, vol. 26, no. 4, pp. 617–625, 1993.
- [103] C. Li and P. Tam, “An iterative algorithm for minimum cross entropy thresholding,” *Pattern Recognit. Lett.*, vol. 19, no. 8, pp. 771–776, 1998.
- [104] J. M. S. Prewitt and M. L. Mendelsohn, “The analysis of cell images,” *Ann. N. Y. Acad. Sci.*, vol. 128, no. 3, pp. 1035–1053, 1966.
- [105] S. van der Walt, J. L. Schönberger, J. Nunez-Iglesias, F. Boulogne, J. D. Warner, N. Yager, E. Gouillart, T. Yu, and the scikit-image contributors, “scikit-image: image processing in Python,” *PeerJ*, vol. 2, p. e453, 6 2014.



## Chapter 6

# All-sky search in early O3 LIGO data for continuous gravitational-wave signals from unknown neutron stars in binary systems

This chapter is an adaptation of the material presented in

*All-sky search in early O3 LIGO data for continuous gravitational-wave signals from unknown neutron stars in binary systems*

R. Abbott et al. (LIGO–Virgo Collaboration)

[Phys. Rev. D 103, 064017 \(2021\)](#) – [arXiv:2012.12128 \[gr-qc\]](#)

DOI: 10.1103/PhysRevD.103.064017

## 6.1 Introduction

Continuous gravitational waves (CWs) are a long-lasting form of gravitational radiation. For ground-based interferometric detectors the canonical sources are rapidly spinning neutron stars (NSs) sustaining a quadrupolar deformation. Several emission mechanisms have been proposed, such as crustal deformations, r-modes, or free precession (see [1] for a recent review). Detecting CWs would probe the physics of such compact objects, leading us to a better understanding of the equation of state of matter under extreme conditions. More exotic types of CW sources are also theorized, such as boson clouds around spinning black holes [2].

Every CW search method assumes certain information about the intended sources. All-sky searches, such as the one reported in this paper, impose the least constraints on the CW emission. The latest results obtained by the LIGO–Virgo collaboration using Advanced LIGO [3] and Advanced Virgo [4] data, covering targeted (known pulsars), directed (known sky locations), and all-sky searches, can be found in [5, 6, 7, 8, 9, 10].

All-sky searches require highly efficient analysis methods because they must account for a Doppler modulation due to the Earth’s movement with respect to the Solar System Barycenter (SSB), an effect that depends on sky position. In principle, one can construct a search pipeline using fully coherent matched filtering; for wide parameter space searches, however, such an approach quickly becomes computationally unaffordable [11]. As a result, semicoherent methods are used, splitting the data stream into smaller time segments that can be coherently analyzed. Then, per-segment results are combined according to the expected frequency evolution of the template under analysis. This method reduces the computational cost of a search while achieving a reasonable sensitivity.

Only a small fraction of the expected population of galactic NSs has been detected electromagnetically [12]. Through gravitational waves we could access these unknown populations of NSs. About half of the NSs detected using electromagnetic means within the most sensitive frequency band of current ground-based detectors are part of a binary system [13, 14]. Searches for CWs from this class of NSs pose an additional, substantial computational challenge compared to standard all-sky searches that target isolated NSs because additional unknown binary orbital parameters increase the search parameter space dimensionality. As a result, one must use specialized methods in order to search for this type of signal.

We present an all-sky search for CWs produced by NSs in binary systems using the semicoherent `BinarySkyHough` pipeline [13]. It builds upon `SkyHough` [15], inheriting its characteristic noise robustness and computational efficiency, and uses Graphics Processing Units (GPUs) to speed up the core part of the search. The concept of `BinarySkyHough` is to compute search statistics over the parameter space and to use those statistics to rank the interesting regions for subsequent follow up using more sensitive, computationally demanding techniques. This balance between sensitivity and computational cost has proven effective in previous searches of the LIGO O2 observing run data using both the isolated `SkyHough` [8] and `BinarySkyHough` [16] flavors of this pipeline.

In Sec. 6.2 we introduce the signal model; Sec. 6.3 describes the early third observing run of the Advanced LIGO and Advanced Virgo detectors; Sec. 6.4 briefly describes the main analysis pipeline; Sec. 6.5 introduces the first post processing stage; and in Sec. 6.6 we estimate the sensitivity of this search. In Sec. 6.7 we further analyze the most significant outliers and rule them out as non-astrophysical candidates. We present our conclusions in Sec. 6.8.

## 6.2 Signal Model

A non-axisymmetric neutron star spinning about one of its principal axes is expected to emit gravitational waves at twice its rotation frequency  $f_0 = 2f_{\text{rot}}$  with a strain amplitude given by [17]

$$h(t) = h_0 [F_+(t; \psi, \hat{n}) \frac{1 + \cos \iota}{2} \cos \phi(t) + F_\times(t; \psi, \hat{n}) \cos \iota \sin \phi(t)], \quad (6.1)$$

where  $F_{+, \times}$  are the antenna patterns of the interferometric detectors, depending on the polarization angle  $\psi$  and the sky position  $\hat{n}$  of the source;  $h_0$  and  $\cos \iota$  are the characteristic CW amplitude and the cosine of the inclination of the source with respect to the line of sight, respectively;  $\phi(t)$  represents the phase of the gravitational wave signal.

The CW amplitude  $h_0$  can be expressed in terms of the physical properties of the source once an emission mechanism has been assumed. The three principal moments of inertia of a non-axisymmetric NS are given by  $I_x$ ,  $I_y$ ,  $I_z$ , and the equatorial ellipticity is given by  $\epsilon = |I_x - I_y|/I_z$ , assuming the spin axis is aligned with  $I_z$ . The gravitational wave amplitude can be expressed as

$$h_0 = \frac{4\pi^2 G}{c^4} \frac{I_z \epsilon}{d} f_0^2, \quad (6.2)$$

where  $d$  denotes the distance to the source from the detector,  $f_0$  the gravitational wave frequency, and  $G$  and  $c$  respectively refer to the gravitational constant and the speed of light. We can further relate this quantity to the mass quadrupole  $Q_{22}$  of the star through the equatorial ellipticity

$$\epsilon = \sqrt{\frac{8\pi}{15}} \frac{Q_{22}}{I_z}. \quad (6.3)$$

We can describe the signal phase via Taylor expansion with respect to a fiducial starting time  $\tau_0$  in the source frame

$$\phi(\tau) = \phi_0 + 2\pi [f_0 \cdot (\tau - \tau_0) + \dots], \quad (6.4)$$

where  $\tau$  is the proper source frame time and  $\phi_0$  represents the initial phase at  $\tau_0$ . The number of higher order terms to include in this expansion depends on the population of NSs under consideration. After analyzing the ATNF pulsar catalog [14], it was argued in [13] that searching for NSs in binary systems need not take into account any spindown parameters when using datasets lasting for less than a few years. As we will discuss in Sec. 6.8, this search remains sensitive to signals up to a certain spindown value, but there is an implicit limit on the astrophysical reach.

Because of the relative motion of the detector around the SSB and the relative motion of the source around the Binary System Barycenter (BSB), the phase as measured by the detector at time  $t$  is Doppler-modulated according to the timing relation

$$\tau + a_p \sin [\Omega (\tau - \tau_{\text{asc}})] = t + \frac{\vec{r}(t) \cdot \hat{n}}{c} - \frac{d}{c}, \quad (6.5)$$

where  $a_p$  represents the semi-major axis of the binary orbit projected onto the line of sight (measured in light-seconds),  $\Omega$  represents the orbital frequency of the source,  $\tau_{\text{asc}}$  represents the time of passage through the ascending node as measured from the source frame and  $\vec{r}$  represents the position of the detector in the SSB. In order to derive this expression, we assumed circular, Keplerian orbits; the search remains sensitive, however, to signals from sources in binary systems up to a certain eccentricity as discussed in Section 6.7.1 and [13].

We define a *template* as  $\lambda = \{f_0, \hat{n}, a_p, \Omega, t_{\text{asc}}\}$ . The parameter space (i.e. the set of all templates searched) will be denoted as  $\mathbb{P}$ . The orbital period is related to the orbital angular frequency by  $P = 2\pi/\Omega$ .

We refer the reader to [18] for a complete derivation of Eq. (6.5) and a discussion about how to express Eq. (6.4) in the detector frame. The gravitational wave frequency evolution associated to a template  $\lambda$  as measured from the detector frame is thus

$$f_\lambda(t) = f_0 \cdot \left( 1 + \frac{\vec{v}(t) \cdot \hat{n}}{c} - a_p \Omega \cos[\Omega(t - t_{\text{asc}})] \right), \quad (6.6)$$

where  $\vec{v}(t)$  refers to the detector velocity and  $t_{\text{asc}}$  is akin to  $\tau_{\text{asc}}$  measured from the detector frame. Equation (6.6) assumes a projected semi-major axis of about  $\mathcal{O}(1 - 10)$  s and orbital period of at least a few hours so that the change in radial position of the source is negligible compared to the wave's travelling time. We choose the initial phase  $t_{\text{asc}}$  to be located within the range  $[t_{\text{mid}} - \frac{P}{2}, t_{\text{mid}} + \frac{P}{2}]$ , where  $t_{\text{mid}}$  represents the mean time between the start and the end of the run measured in GPS seconds.

### 6.3 Data used

The first part of the third observing run of the Advanced LIGO and Advanced Virgo detectors (O3a) comprises six months of data collected from the 1st of April 2019 at 15:00 UTC to the 1st of October 2019 at 15:00 UTC. Data was taken by the Advanced LIGO detectors, located in Hanford (Washington, USA, designated H1) and Livingston (Louisiana, USA, designated L1), together with the Advanced Virgo detector, located in Cascina (Pisa, Italy). We did not make use of Advanced Virgo data because of an unfavorable trade-off between computing cost and expected sensitivity improvement of the search. The detector duty factor (the fraction of the run when the detector is collecting observational-quality data) was 71.2% for H1 and 75.8% for L1. The implementation of instrumental upgrades has allowed the detectors to improve their overall sensitivities with respect to the previous observing run (O2) [19].

For the duration of the run, several artificial signals were injected into both detectors in order to calibrate and monitor their performance. *Calibration lines* are artificial monochromatic signals, injected at different frequencies in each detector to avoid coherent artifacts. They are used to monitor time-varying detector operating parameters. *Hardware injections*, on the other hand, are artificial quasi-monochromatic signals consistently injected into both detectors in order to mimic the effects of an actual CW signal present in both detectors. They are used to verify expected detector response and characterize calibrated data [20]. Both of these artificial signals may interfere with CW searches in general, showing up as significant candidates due to their high strength in the detector spectrum. Spectral artifacts in detector data can be produced by environmental or instrumental noise and also interfere with CW searches [21].

The search was performed using Short Fourier Transforms (SFTs) created from the C00 (initial calibration version) time-domain observing-quality strain data [22]. These SFTs were extracted from SFDB (Short Fourier Data Base) data [23], which incorporates a time-domain cleaning procedure to avoid noise-floor degradation due to glitches and other forms of transient noise. Every SFT lies completely within observing-quality data. Fourier transforms were computed using a Tukey-windowed baseline of  $T_{\text{SFT}} = 1024$  s with tapering parameter  $\beta_{\text{Tukey}} = 0.5$  and a 50% overlap. These values are collected in table 6.1.

Following the same procedure used in the O2 *SkyHough* search [8], SFT data are split into two datasets to be used in two different stages of the search. The first dataset, which we refer to as *non-overlapping*, leaves out overlapping SFTs (i.e. every SFT starts at the end of the previous one). The second dataset, which we refer to as *overlapping*, contains all of the SFTs. Using the non-overlapping set for the first stage of the analysis reduces the computational cost of the search at a manageable loss

Search setup parameter	Value
$T_{\text{SFT}}$	1024 s
$\beta_{\text{Tukey}}$	0.5
$T_{\text{obs}}$	14832675 s
$t_{\text{mid}}$	1245582821.5 s
$\Delta f$	0.125 Hz

TABLE 6.1: Miscellaneous parameters used in the search.  $T_{\text{SFT}}$  denotes the time span employed to compute Short Fourier Transforms (SFTs).  $\beta_{\text{Tukey}}$  refers to the tapering parameter of the Tukey window, denoting the fractional length of the window’s central unitary plateau.  $T_{\text{obs}}$  is the observing time of the run.  $t_{\text{mid}}$  represents the mean time between the start and the end of the run measured in GPS seconds.  $\Delta f$  refers to the bandwidth of the individual sub-bands analyzed by each computing job.

	Non-overlapping	Overlapping
H1	10172	20577
L1	10962	22049
Total	21234	42626

TABLE 6.2: Number of Short Fourier Transforms (SFTs) in each of the datasets. Characteristics of these SFTs are summarized in table 6.1 and section 6.3.

in sensitivity. Table 6.2 lists the number of SFTs in each of the datasets. Datasets contain SFTs from *both* LIGO detectors (i.e. we perform a *multi-detector* search [13]).

## 6.4 The Search Pipeline

We split the search into two main frequency bands: the low-frequency band, from 50 Hz to 100 Hz, and the high-frequency band, from 100 Hz to 300 Hz. These bands are further divided into  $\Delta f = 0.125$  Hz sub-bands, which constitute the basic working unit of our setup: each *computing job* performs an all-sky search over one such sub-band, searching for binary modulated signals within a certain region of the binary parameter space among the ones specified in Fig. 6.1 and Table 6.3. Because of the limited computing power available, the high-frequency search focuses on a single binary parameter space region, denoted as B in Table 6.3; the low-frequency search is performed in all four binary parameter space regions.

The search parameter space is gridded with templates as described in [13]:

$$\begin{aligned} \delta f_0 &= \frac{1}{T_{\text{SFT}}}, & \delta\theta &= \frac{c/v}{T_{\text{SFT}} f_0 P_f}, & \delta a_p &= \frac{\sqrt{6m}}{\pi T_{\text{SFT}} f_0 \Omega}, \\ \delta\Omega &= \frac{\sqrt{72m}}{\pi T_{\text{SFT}} f_0 a_p \Omega T_{\text{obs}}}, & \delta t_{\text{asc}} &= \frac{\sqrt{6m}}{\pi T_{\text{SFT}} f_0 a_p \Omega^2}, \end{aligned} \quad (6.7)$$

where  $\delta\theta$  refers to the angular sky position resolution,  $v = |\vec{v}|$  and  $v/c \sim 10^{-4}$ .  $T_{\text{obs}}$  denotes the observing time of the search, quoted in Table 6.1. The variables  $P_f$  and  $m$  are the so-called *pixel factor* and *mismatch* parameters, which can be used to manually control the parameter space template density. In this search, we tune them in order to adjust the computing cost as we reach higher frequencies, where template spacing naturally becomes finer. Table 6.4 summarizes the choices made for each of the frequency bands.

The pipeline uses the Hough transform to relate tracks in the digitized spectrogram, as explained below, to points in the parameter space. For each point in the parameter space  $\lambda \in \mathbb{P}$  there is a corresponding track [see Eq. (6.6)] of the time-frequency evolution, which denotes the instantaneous frequency of the signal as observed by the detector.




FIGURE 6.1: Binary orbital parameters considered by the present search. Solid color regions denote parameter space regions in which a search was performed; blue dots mark binary orbital parameters corresponding to the known binary pulsar population. Regions A, B, C and D were covered by the low-frequency analysis, while region B was covered by the high-frequency analysis as well. Time of ascending node passage is taken into account according to the orbital frequency, as explained in Sec. 6.2. Pulsar population data was taken from [14] using [24].

### 6.4.1 Ranking statistics

Let us assume the data can be described as a noise background plus a CW signal

$$x(t; \lambda) = n(t) + h(t; \lambda). \quad (6.8)$$

We start by computing the *normalized power* of SFT data

$$\rho_k^\alpha = \frac{|\tilde{x}_k^\alpha|^2}{\langle |\tilde{n}_k^\alpha|^2 \rangle}, \quad (6.9)$$

where tildes represents a Fourier transformed quantity,  $k$  indexes frequency bins,  $\alpha$  indexes SFTs and  $\langle \cdot \rangle$  denotes a running median average using 101 frequency bins, as explained in [13]. Each SFT  $\alpha$  can be related to a certain starting time  $t_\alpha$ , effectively obtaining a spectrogram where each bin  $(\alpha, k)$  corresponds to the normalized power  $\rho_k^\alpha$  present at a certain frequency bin  $k$  in a certain SFT  $\alpha$ . Then, we impose a normalized power threshold  $\rho_{\text{th}} = 1.6$  to digitize the spectrogram, obtaining a discrete spectrogram populated by ones and zeros.

For each template, we follow the corresponding track and define the first ranking statistic, the *number count*, as the weighted sum of ones and zeroes

$$n(\lambda) = \sum_{(\alpha, k) \in f_\lambda} w_k^\alpha \mathcal{H}(\rho_k^\alpha - \rho_{\text{th}}), \quad (6.10)$$

where  $\mathcal{H}$  denotes the Heaviside step function and the weights  $w_k^\alpha$  account for varying noise floor and antenna response effects [25].

The number count statistic can be efficiently computed by means of the Look Up Table (LUT) approach described in [15]. Incidentally, this strategy simplifies the cost by analyzing multiple sky positions (called *sky patches*) together. The approach applies the Doppler modulation used to analyze a particular frequency bin to a neighborhood of frequency bins. The sensitivity loss introduced by this approximation is later compensated by re-analyzing the most significant candidates using their exact time-frequency tracks [26].

Binary Region	$P$ [days]	$a_p$ [s]
A	[15, 45]	[10, 40]
B	[7, 15]	[5, 15]
C	[5, 7]	[2, 10]
D	[3, 5]	[2, 5]

TABLE 6.3: Binary parameter space regions analyzed by the search, corresponding to the four colored regions in Fig. 6.1. Time of passage through the ascending node  $t_{\text{asc}}$  is searched along the interval specified in Sec. 6.2.

Label	Frequency [Hz]	$m$	$P_{\text{f}}$
L	[50, 100)	0.4	1
H1	[100, 125)	0.4	1
H2	[125, 150)	0.6	1
H3	[150, 200)	0.9	0.8
H4	[200, 250)	1.6	0.75
H5	[250, 300)	2	0.7

TABLE 6.4: Mismatch and pixel factor configurations for the different frequency bands of the search. L refers to the low-frequency band; H1-5 refer to each of the five sub-bands into which the high-frequency band was partitioned: 1 and 2 span 25Hz each, while 3 to 5 span 50 Hz each.

The re-analysis uses the *weighted normalized power* statistic,

$$\rho(\lambda) = \sum_{(\alpha,k) \in f_\lambda} w_k^\alpha \rho_k^\alpha. \quad (6.11)$$

Using this new ranking statistic instead of simply re-computing Eq. (6.10) along the exact track yields a 10–20% improvement in detection efficiency for the toplists (ranking of the most significant candidates) based on the number-count statistic used here and discussed below [8].

In order to further select candidates across different sky patches, we compute a *significance* statistic by normalizing Eq. (6.11) to the *expected noise values* derived in [15]:

$$s_\rho(\lambda) = \frac{\rho(\lambda) - \bar{\rho}}{\sigma_{\bar{\rho}}}, \quad (6.12)$$

where  $\bar{\rho}$  and  $\sigma_{\bar{\rho}}$  represent the expected value and standard deviation of weighted normalized power in pure Gaussian noise. This statistic removes any dependency on the sky position of the source due to the weights, being well suited for comparisons across different sky patches.

### 6.4.2 Toplist construction

Toplists are constructed frequency-bin wise across sky patches, as shown schematically in Fig. 6.2. For a given sky patch and frequency bin, the top 5% of parameter space candidates are selected according to the number count statistic Eq. (6.10) using the LUT approach and the non-overlapping set of SFTs. Then, they are re-analyzed computing their corresponding normalized power Eq. (6.11) along the exact time-frequency evolution given by Eq. (6.6) using the overlapping set of SFTs. Finally, top candidates according to Eq. (6.12) are collected into a final toplist. We collect the top 80000 candidates from each 0.125 Hz sub-band.

This approach optimizes the GPU usage in the number count stage (preventing loud spectral artifacts from saturating the toplist) as each frequency bin provides a controlled number of candidates.

---

```

Result: Toplist per 0.125 Hz sub-band
Select 0.125 Hz sub-band;
for SkyPatchIndex  $p$  do
  Initialize per-patch toplist  $p$ ;
  for FrequencyBinIndex  $k$  do
    Rank candidates by number count (10);
    Select top 5 candidates;
    Rank candidates by norm. power (11);
    Select top 0.1 candidates;
    Write candidates into per-patch toplist  $p$ ;
  end
end
Initialize sub-band toplist;
Collect all per-patch toplist into sub-band toplist;
Rank sub-band toplist by significance (12);
Select top 80000 candidates.

```

---

FIGURE 6.2: Explicit description of the `BinarySkyHough` toplist construction process.

## 6.5 Post Processing

Similar to previous searches [8, 16], we apply a clustering algorithm to the resulting candidates in order to look for particularly interesting candidates. Clustering candidates reduces the total number of candidates to follow up since typically many candidates are found to be produced by a single source (either a CW signal or an instance of instrumental noise).

We implement a new clustering algorithm using the frequency evolution of a candidate to define a parameter space distance [27]. This choice allows the algorithm to naturally take into account the parameter space structure, avoiding the usage of ad hoc sky projections or mishandling periodic boundary conditions.

After the cluster selection, we apply the well-known *line veto*, used in previous searches (e.g. [28, 29, 8, 16]) in order to rule out non-astrophysical candidates.

### 6.5.1 Clustering

The clustering algorithm is summarized below; see [27] for further details. Given two candidates with template values  $\lambda, \lambda^* \in \mathbb{P}$ , we define the parameter space distance as

$$d(\lambda^*, \lambda) = \frac{T_{\text{SFT}}}{N_\alpha} \sum_{t_\alpha} |f_{\lambda^*}(t_\alpha) - f_\lambda(t_\alpha)|, \quad (6.13)$$

where  $f_\lambda(t_\alpha)$  represents the instantaneous frequency of a CW produced by a source with parameters  $\lambda$  as measured by the detector at time  $t_\alpha$  and  $N_\alpha$  denotes the number of SFT timestamps used. Essentially, Eq. (6.13) is the average mismatch among time-frequency tracks.

Clusters are formed by grouping together candidates from connected components, i.e. each candidate in a cluster is closer than a maximum distance  $d^{\text{th}} = 1$  to at least one other candidate in the same cluster. Final clusters are ranked according to the significance of their loudest candidate, which we will refer to as the *cluster center*.

For each 0.125 Hz toplist, the top 5 clusters according to their significance are selected. This leaves us with a total of 16000 clusters: 8000 for the high-frequency search and 2000 for each region of the low-frequency search.

### 6.5.2 Line Veto

Before the outlier follow up, we apply the line veto to the obtained cluster centers. Using the list of identified narrow spectral artifacts [30], the veto discards any candidate whose time-frequency track crosses an instrumental line, since such a candidate would likely become significant not because of astrophysical reasons but rather instrumental ones.

For every cluster center with parameters  $\lambda$ , we compute its *bandwidth*

$$\mathcal{BW}(\lambda) = [\min_{t_\alpha} f_\lambda(t_\alpha), \max_{t_\alpha} f_\lambda(t_\alpha)]. \quad (6.14)$$

Regions	LA	LB	LC	LD	H1B	H2B	H3B	H4B	H5B
Initial Clusters	2000	2000	2000	2000	1000	1000	2000	2000	2000
Vetoed by Identified Line	366	359	359	373	44	0	32	30	30
Surviving Clusters	1634	1641	1641	1627	956	1000	1968	1970	1970
Fraction (%)	81.7	82.05	82.05	81.35	95.6	100	98.4	98.5	98.5
Surviving Outliers after $2\hat{\mathcal{F}}_{\text{th}}$ veto	73	72	71	71	7	6	8	3	0

TABLE 6.5: Numbers of cluster centers discarded by the line veto using lines present in [30]. The number of surviving outliers after the follow-up stage (see Sec. 6.7) is here specified for the sake of completeness. Five clusters were collected from each 0.125 Hz band: regions H1B and H2B, being the only ones spanning 25 Hz, yield a lower number of clusters.

If the bandwidth of a candidate contains or overlaps with any of the lines present in [30], then the candidate is discarded because of its likely non-astrophysical origin. This veto reduces the number of clustered candidates by  $\sim 20\%$  in the low-frequency search and by a few percent in the high-frequency search (see Table 6.5). This difference is to be expected, considering the greater amount of instrumental lines present at lower frequencies.

Other narrow spectral artifacts have not yet been identified as clearly non-astrophysical in origin in an unidentified list [31]. Although this list has not been used to veto clustered candidates, some of them are consistent with artifacts in the unidentified list (see appendix 6.A).

## 6.6 Sensitivity

The sensitivity of the search is determined using a similar procedure as for previous all-sky searches [28, 29, 8, 16]. A campaign of adding software-simulated signals to the data in order to estimate the  $h_0$  that corresponds to a 95% average detection rate was carried out. We quantify sensitivity using the *sensitivity depth* [32, 33]

$$\mathcal{D} = \frac{\sqrt{S_n}}{h_0} \quad (6.15)$$

where  $S_n$  represents the single-sided Power Spectral Density of the data (PSD),  $\sqrt{S_n}$  is referred to as the Amplitude Spectral Density (ASD) and  $h_0$  is the previously defined CW amplitude. This figure of merit characterizes the sensitivity of the search to putative signals and accounts for the detector sensitivity as a function of frequency. The actual single-sided PSD in Eq. (6.15) depends on the analysis method being used. `BinarySkyHough` sensitivity is dominated by the first stage using the weighted number count statistic meaning one should use the *inverse squared averaged* PSD as shown in equations (42) to (44) of [34]

$$S_n(f) = \sqrt{\frac{N_\alpha}{\sum_\alpha [S_\alpha(f)]^{-2}}}, \quad (6.16)$$

where  $S_\alpha(f)$  represents the running-median noise floor estimation using 101 bins corresponding to the SFT labeled by starting time  $t_\alpha$  at frequency  $f$ . The goal is to characterize the average detection rate by numerically computing the efficiency distribution with respect to the depth. The result is interpolated to find the estimated sensitivity depth that corresponds to 95% detection efficiency. Using Eq. (6.15) the sensitivity depth is converted to the sensitivity amplitude. It is in this last step where the systematic error of the calibration is potentially relevant.

Systematic error in the amplitude of calibration of C01 data (final calibration version) is estimated to be lower than 7% (68% confidence interval) for both detectors over all frequencies throughout O3a [20]. Relative deviations of ASDs computed using C00 data with respect to ASDs computed using C01 data (used as a proxy for an estimate of systematic error in C00 data calibration which otherwise does not exist for all time or frequencies) are below 7% for all frequency bands except in the [59, 61] Hz sub-band, where the relative deviation is 10%. Assuming the proxy for C00 systematic error is complete, the impact of such 10%-level of systematic error is negligible to the conclusions of this analysis.



Five representative frequency bands are selected across each 25 Hz band and binary parameter space region, and five sensitivity depth values used, namely [18, 20, 22, 24, 26] Hz<sup>-1/2</sup>. Two hundred signals drawn from uniform distributions in phase and amplitude parameters are added to the data at each depth, band and binary parameter space region. For each simulated signal, `BinarySkyHough` analyzes the data again in order to evaluate how many of them are detected. Sensitivity depth values are selected such that the 95% efficiency depth was properly bracketed; regions H4B and H5B required two extra depth values [14, 16] Hz<sup>-1/2</sup> to ensure this. Using a small number of frequency bands drastically reduces the computing cost of the sensitivity estimation procedure while yielding consistent results when compared to an exhaustive injection campaign, as justified in [16].

Three criteria must be fulfilled in order to label a simulated signal as “detected”. First, the toplist obtained from the injection search should contain at least one candidate whose significance Eq. (6.12) is greater than the minimum significance present in the corresponding all-sky toplist. Second, after clustering the injection toplist, at least one cluster with a significance greater than the lowest significance recovered by the corresponding all-sky clustering must be obtained. These two criteria ensure the injection is prominent enough so as not to be discarded by the first stage of the search. Lastly, we require at least one of the top five clusters from the injection toplist to be located closer than two parameter space bins in each of the parameters with respect to the injection parameters. This last criterion takes into account the fact that, in the actual search, a follow up will be done in corresponding regions around each significant cluster center.

After separating detected from non-detected simulated signals, we construct efficiency curves akin to the example shown in Fig. 6.3. Each point is the fraction of simulated signals detected (i.e. detection efficiency) as a function of the sensitivity depth. For each sensitivity depth set of  $N_I = 200$  simulated signals, the uncertainty on detection efficiency  $E$  is given by

$$\delta E = \sqrt{\frac{E \cdot (1 - E)}{N_I}}. \quad (6.17)$$

Then, using SciPy’s `curve_fit` function [35], we fit a sigmoid curve to the data given by

$$S(\mathcal{D}; \vec{p}) = 1 - \frac{1}{1 + e^{-p_0(\mathcal{D} - p_1)}}, \quad (6.18)$$

with fitted parameters  $\vec{p} = (p_0, p_1)$ . This expression can be inverted in order to find the 95% sensitivity depth.

The interpolations are accompanied by a corresponding uncertainty, obtained through the covariance matrix of the fit  $\mathcal{C}(\mathcal{D})$  as

$$\delta \mathcal{D}^{95\%} = \sqrt{\nabla_{\vec{p}} S(\mathcal{D}; \vec{p})^T \cdot \mathcal{C}(\mathcal{D}) \cdot \nabla_{\vec{p}} S(\mathcal{D}; \vec{p})} \Big|_{\mathcal{D} = \mathcal{D}^{95\%}}, \quad (6.19)$$

where

$$\nabla_{\vec{p}} S(\mathcal{D}; \vec{p})^T = \left( \frac{\partial S(\mathcal{D}; \vec{p})}{\partial p_0}, \frac{\partial S(\mathcal{D}; \vec{p})}{\partial p_1} \right) \quad (6.20)$$

and the superscript T denotes matrix transposition. The resulting interpolated depths per frequency band are shown in Fig. 6.4. The high-frequency search shows a clear degradation of depth values as frequency increases. This is related to the decaying density of parameter space templates: the higher the frequency, the finer one must construct a template bank in order to achieve a comparable level of sensitivity.

Finally, we compute an average 95% sensitivity depth for each of the regions quoted in Table 6.6. We also quote a corresponding  $3\sigma$  uncertainty, which previous studies have proven to deliver a good coverage of the actual 95% efficiency sensitivity depth [16]. These values are translated to CW amplitude  $h_0$  via Eq. (6.15) and shown in Fig. 6.5.




FIGURE 6.3: Example of 95% sensitivity depth interpolation. Five sensitivity depths were selected at 124.625 Hz in region H1B. 200 injections were injected at each of these depths, applying the criteria exposed in the text in order to label injections as detected/not detected. Blue dots represent the fraction of detected injections; the sigmoid fit is represented by a blue line; fit uncertainties at one, two and three sigmas are represented by pale yellow shades. The interpolated 95% sensitivity depth  $\mathcal{D}^{95\%} = 21 \pm 0.4$  is marked using a star.

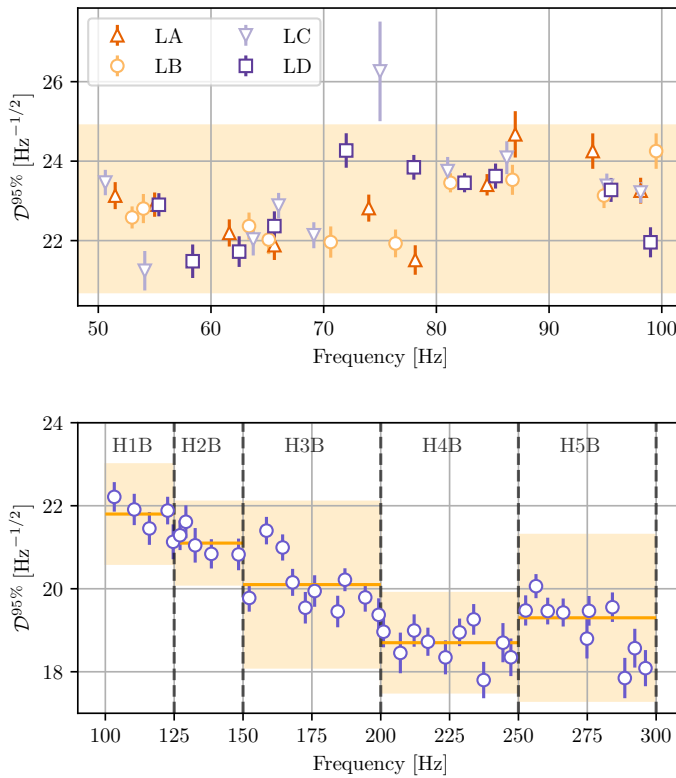


FIGURE 6.4: Average 95% sensitivity depths obtained in the low-frequency (top panel) and high-frequency (bottom panel) bands. Data points correspond to the interpolated results obtained through the sigmoid fit of the efficiencies at the selected frequency bands (5 bands randomly selected in each 25 Hz). Error bars correspond to 95% efficiency uncertainties towards low depth values. Shaded regions show the averaged results with their uncertainties, as summarized in Table 6.6. In the top panel, shading is only shown for the results obtained for the binary parameter space region B.

Region	$\langle \mathcal{D}^{95\%} \rangle \pm 3\sigma$ [ $\text{Hz}^{-1/2}$ ]
LA	$22.9 \pm 2.5$
LB	$22.8 \pm 2.1$
LC	$23.0 \pm 2.5$
LD	$23.0 \pm 2.5$
H1B	$21.8 \pm 1.2$
H2B	$21.1 \pm 1.0$
H3B	$20.1 \pm 2.0$
H4B	$18.7 \pm 1.2$
H5B	$19.3 \pm 2.0$

TABLE 6.6: Average 95% sensitivity depths for the parameter space regions analyzed in this search. Region labels are defined in Tables 6.3 and 6.4.




FIGURE 6.5: Implied 95% efficiency amplitude from the obtained sensitivity depth values. The  $h_0^{95\%}$  amplitude estimates are obtained from the 95% efficiency depth values shown in Table 6.6 and the *inverse squared averaged* PSD using Eq. (6.15). Low-frequency results are shown for binary parameter space region B.

Hyperparameter	Value
Parallel chains	3
Walkers per chain	100
Burn-in & Production steps	100 + 100

TABLE 6.7: MCMC hyperparameter choice for the first stage of the follow up. Each of the *parallel chains* sample the likelihood at a different temperature, as explained in [41].

## 6.7 Follow Up

Remaining candidates from the main search are followed up applying a more sensitive method to the data. Longer coherence times constrain the phase evolution of the candidate under consideration and would yield higher significance for a true continuous wave signal. A potential downside remains that a true signal could be discarded if it is not well modeled by the assumed phase evolution. Moreover, increasing the coherence time also requires increasing the density of templates so as not to overlook a putative signal.

An effective way to cover small parameter regions is through Markov Chain Monte Carlo (MCMC) methods, which, rather than following a prescribed parameter space grid, sample the parameter space following a certain probability density function. Reference [36] describes how this can be implemented in a search for continuous waves by using the so-called  $\mathcal{F}$ -statistic, a well-established CW analysis technique, as a likelihood function. We refer to [36] and references therein for an in-depth explanation of this method.

The  $\mathcal{F}$ -statistic is a coherent statistic, usually referred to as  $2\tilde{\mathcal{F}}$ , which compares data against templates by matched filtering. A semicoherent  $\mathcal{F}$ -statistic  $2\hat{\mathcal{F}}$  can be defined by adding individual  $2\tilde{\mathcal{F}}$  values computed over  $N_{\text{seg}}$  segments spanning  $T_{\text{coh}}$  each, in the same way as weighted normalized power was computed from weighted power in Eq. (6.11):

$$2\hat{\mathcal{F}}(\lambda) = \sum_{s=0}^{N_{\text{seg}}-1} 2\tilde{\mathcal{F}}_s(\lambda), \quad (6.21)$$

where the index  $s$  indicates the coherent quantity has been computed for a certain segment spanning  $T_{\text{coh}}$ .

We use software injections in order to calibrate a threshold  $2\hat{\mathcal{F}}_{\text{th}}$ . Candidates such that  $2\hat{\mathcal{F}}(\lambda) < 2\hat{\mathcal{F}}_{\text{th}}$  will be deemed as non-significant and consequently discarded.

This algorithm is implemented in `PyFstat` [37, 38]. It builds on top of `LALSuite` [39], which provides the CW data analysis functionality, and `ptemcee` [40, 41], which implements the MCMC algorithms.

### 6.7.1 MCMC follow-up configuration

The MCMC follow up employed is not intended to describe the posterior distribution of parameters defining a candidate. Rather, we only require enough convergence such that the sampled  $\mathcal{F}$ -statistic values are close enough to the local maximum to establish a reliable veto threshold.

#### Sampler configuration

The `ptemcee` package implements an ensemble-based sampler that uses several walker chains to sample multimodal distributions. Expensive setups are not required in order to perform a first-stage follow up using a threshold-based approach. The reason for this is two-fold: we are increasing the coherence time with respect to the search, and we do not require extensive convergence to be achieved. No second-stage follow up was required because all of the first-stage outliers were attributed to instrumental causes. If this was not the case, we would have applied a second follow-up stage using a more expensive setup. The number of parallel chains, walkers per chain, and number of steps to take are summarized in Table 6.7.

We choose to use  $N_{\text{seg}} = 260$ , which corresponds to  $T_{\text{coh}} \simeq 17$  h. This is a longer coherence time with respect to that of the initial stage of the search, and a choice used in previous searches [16].

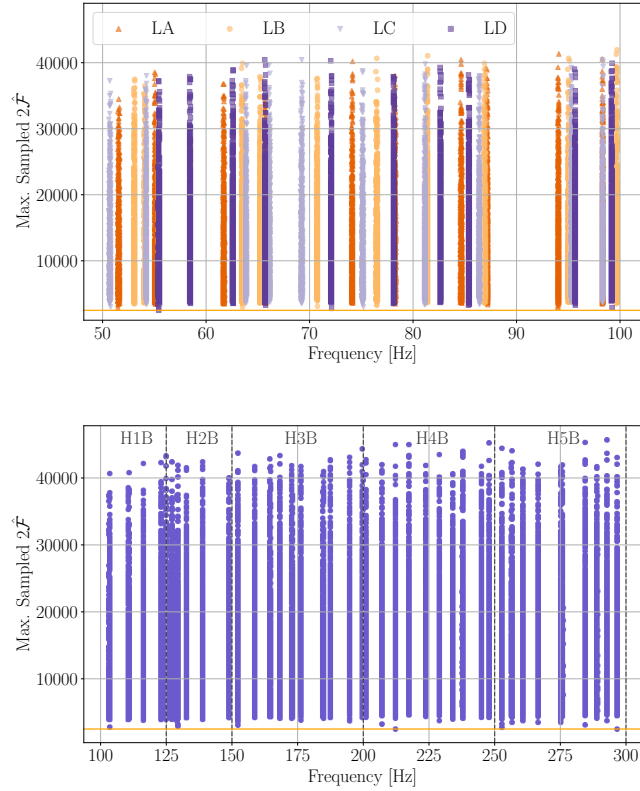


FIGURE 6.6: Recovered  $2\hat{\mathcal{F}}$  values using a set of injections labeled as *detected* by the sensitivity estimation criteria. The amplitudes of these injections were distributed using the five sensitivity depth values explained in Sec. 6.6. The top panel shows the results for the four regions of the low-frequency search, involving 33757 injections; the bottom panel shows the same result for the high-frequency search, using 37549 injections. The horizontal orange line marks the threshold  $2\hat{\mathcal{F}}_{\text{th}} = 2500$ .

### Prior choice

Following a similar prescription as the one given in [36], we set up uniform priors in each parameter space dimension, forming a box centered on each cluster center. Each edge of this box spans two parameter space bins according to the spacing given in Eq. (6.7), where the parameter-space-dependent quantities are computed at the center of the cluster. This is in agreement with the detection criteria imposed to perform the sensitivity estimation.

Although `BinarySkyHough` targets CW sources in circular orbits, it is still sensitive to signals with eccentricities up to a certain value, as long as the Doppler modulation derived from eccentricity is smaller than half a frequency bin. The upper bound for the maximum allowable eccentricity according to this argument was derived in [13]

$$e^{\text{m.a.}} = [2 T_{\text{SFT}} f_0 a_p \Omega]^{-1} . \quad (6.22)$$

Therefore, uniform priors on eccentricity,  $[0, e^{\text{m.a.}}]$ , and argument of periastron,  $[0, 2\pi]$  are included as MCMC parameters. Maximum eccentricities range from 0.2 – 0.5 at 50 Hz to less than 0.1 at 300 Hz.

### 6.7.2 Setting up a threshold

We use the `BinarySkyHough` and the MCMC follow up on a total of 71306 software injections in order to calibrate a significance threshold. The employed injections are consistent with the ones used for the sensitivity estimation, focusing on those detected according to the three criteria (see Sec. VI). This implies a significant fraction of the injections will possess an amplitude below the obtained 95% sensitivity amplitude, as they will be distributed according to the five original depths. The threshold

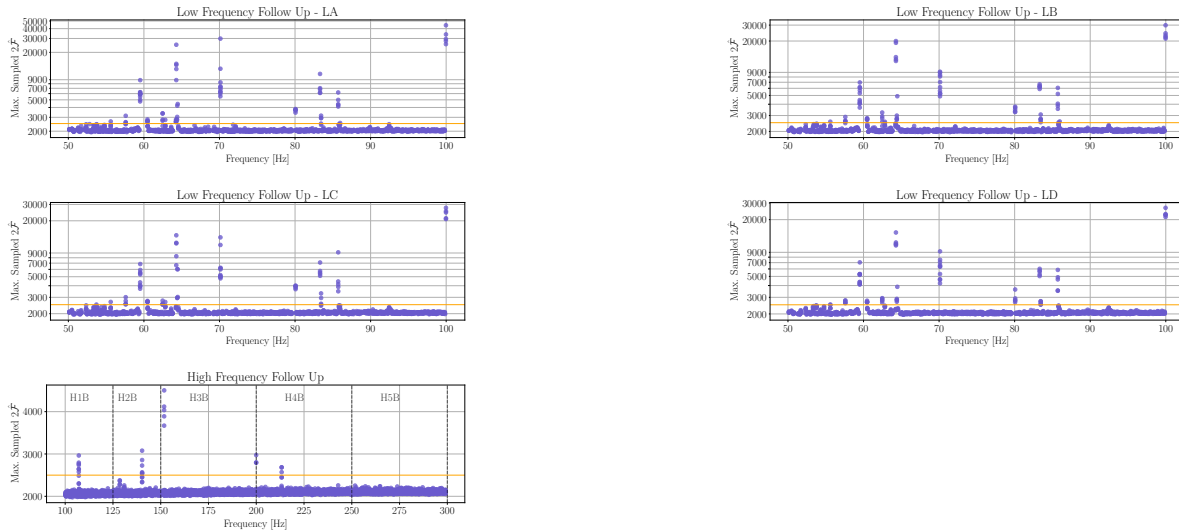


FIGURE 6.7: MCMC follow-up results. Dots represent cluster centers not vetoed by the identified line veto, and a horizontal line represents the imposed threshold  $2\hat{\mathcal{F}}_{\text{th}} = 2500$ , calibrated by software injections (see Fig. 6.6). The horizontal axis represents the frequency value associated to each cluster center, and the vertical axis represents the maximum  $2\hat{\mathcal{F}}$  value sampled by the MCMC run.

obtained using this calibration strategy will have a low false-dismissal rate ( $\lesssim 1/71306 \simeq 1.5 \times 10^{-5}$ ) against signals detectable by this pipeline.

We run the MCMC algorithm in order to sample  $2\hat{\mathcal{F}}$  values, retrieving the maximum value for each of the injections. Resulting  $2\hat{\mathcal{F}}$  values for these simulations are plotted in Fig. 6.6. These results support the choice of  $2\hat{\mathcal{F}}_{\text{th}} = 2500$  as the threshold value, with all detected injections above this threshold.

### 6.7.3 Surviving Outliers

After executing the MCMC follow up and imposing the  $2\hat{\mathcal{F}}_{\text{th}} = 2500$  threshold, 287 outliers remain in the low-frequency band and 24 outliers remain in the high-frequency band, as shown in Fig. 6.7. It is clear from the figure that low-frequency outliers mostly belong to the same frequency bands across the four binary parameter space regions. We next analyze each candidate using a cumulative semicoherent  $\mathcal{F}$ -statistic, defined as

$$2\hat{\mathcal{F}}(\lambda; t) = \sum_{\alpha: t_\alpha < t} 2\tilde{\mathcal{F}}_\alpha(\lambda), \quad (6.23)$$

in order to discern those candidates originating from instrumental noise.

We use three flavors of Eq. (6.23), one for each of the detectors (H1 and L1) and another one using a multi-detector approach (H1 + L1). These statistics lead to the rejection of the remaining outliers, as described below.

#### Line-crossing outliers

CWs are expected to accumulate a  $2\hat{\mathcal{F}}$  value linearly with respect to the observing time. We find that 263 outliers surpass the  $2\hat{\mathcal{F}}_{\text{th}} = 2500$  threshold due to the presence of prominent values of segment-wise  $\mathcal{F}$ -statistic at certain times of the run in one of the detectors (260 in H1 and 3 in L1), as exemplified in Fig. 6.8. The higher number of outliers in H1 arise from the greater number of instrumental lines present in that detector [30, 31].

This behavior would be expected from a candidate whose frequency evolution track crosses a narrow instrumental artifact (line) for a limited duration, either because of the frequency track drifting away from the line or the transient nature of the line itself. Most strong persistent instrumental disturbances are already discarded using the known lines list [30], but weaker lines or transient disturbances (lasting

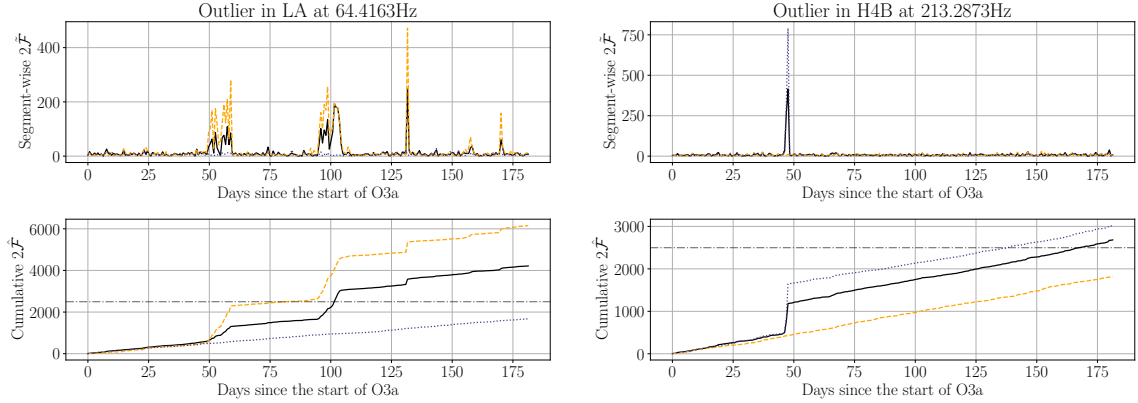


FIGURE 6.8: Example of outliers produced by instrumental noise in one of the Advanced LIGO detectors. Each pair of panels corresponds to a different outlier; the segment-wise  $2\hat{\mathcal{F}}_\alpha$  statistic is shown at the top of each pair; the cumulative semicoherent  $2\hat{\mathcal{F}}$  as described in Eq. (6.23) is shown at the bottom of each pair. Dashed lines denote an H1-only analysis, dotted blue lines an L1-only analysis, and solid black lines a multi-detector analysis. The  $2\hat{\mathcal{F}}_{\text{th}} = 2500$  threshold is shown as a horizontal line.




FIGURE 6.9: Example of an outlier vetoed by the multi-detector consistency veto. Dashed lines denote an H1-only analysis, dotted blue lines an L1-only analysis, and solid black lines a multi-detector analysis. The  $2\hat{\mathcal{F}}_{\text{th}} = 2500$  threshold is shown as a horizontal line.

hours to days), which are more difficult to identify in the run-averaged spectra, could still affect our searches [42, 21].

In Table 6.8 in Appendix 6.A, we present a list of frequency bands containing these 263 candidates whose behavior suggests a brief line crossing or the presence of a transient instrumental disturbance. Outliers were selected as belonging to this category if they have at least one per-segment  $\mathcal{F}$ -statistic value greater than 100, and for each frequency band listed in the table, the first/last timestamps bracket the data segments where  $\mathcal{F}$ -statistic values greater than 50 were observed for at least one of those candidates. Overlapping frequency bands were merged together for the sake of clarity.

### Detector consistency veto

A second set of 28 outliers is discarded by the detector consistency veto (see e.g. [28]). During O3a, the L1 detector presents a better sensitivity than H1 at low frequencies [19]. A CW candidate would be expected to behave consistently, i.e.  $2\hat{\mathcal{F}}_{\text{L1}} > 2\hat{\mathcal{F}}_{\text{H1}}$  for most signals. We calibrate this veto using the aforementioned set of software injections in order to take detector sensitivity anisotropies due to the antenna pattern functions into account, obtaining a maximum relative 5% excess of  $2\hat{\mathcal{F}}_{\text{H1}}$  with respect to  $2\hat{\mathcal{F}}_{\text{L1}}$ .

The 28 outliers rejected with this veto show more than a 30% relative excess of  $2\hat{\mathcal{F}}_{\text{H1}}$  with respect to  $2\hat{\mathcal{F}}_{\text{L1}}$ . Hence, we discard them as being inconsistent with an astrophysical signal. Figure 6.9 shows an example of these outliers. After computing the bandwidth covered by each of these candidates, we obtain seven distinct frequency bands affected by instrumental disturbances of this type, summarized in Table 6.9 in appendix 6.A.




FIGURE 6.10: Example of an outlier surpassing the  $2\hat{\mathcal{F}}_{\text{th}} = 2500$  threshold later vetoed by inspection of the detector spectra. The combined ASD of both detectors around the frequency of this group of outliers is shown in Fig. 6.11. Candidates related to this signal saturated the toplist in the four regions LA, LB, LC, LD. Dashed lines denote an H1-only analysis, dotted blue lines an L1-only analysis, and solid black lines a multi-detector analysis. The  $2\hat{\mathcal{F}}_{\text{th}} = 2500$  threshold is shown as a horizontal line.




FIGURE 6.11: Location of the manually inspected outliers (vertical dashed line) with respect to the amplitude spectral density of the detector, here represented as a multi-detector *inverse squared average* (orange line). A diamond and a square mark the twin peaks' frequency, 59.53Hz and 60.47Hz respectively.

### Powerline sidebands

The last 20 outliers were consistently present in each one of the four parameter space regions within the  $[60.46, 60.48]$  Hz sub-band. These outliers were not vetoed by any of the previous stages. As shown in Fig. 6.10,  $2\hat{\mathcal{F}}$  was accumulated in a fairly linear fashion, achieving greater values in L1 than H1. The detector ASD (Fig. 6.11), however, shows that these candidates were caused by sidebands of the 60 Hz power supply artifact. These sidebands can be explained by a non-linear coupling between the main power supply frequency and a low-frequency noise. They do not appear in the line lists [30, 31] as they do not correspond to narrow spectral artifacts and their effect on CW searches is highly dependent on the search method. Due to the presence of said artifact in the data and the wide spread of the candidates obtained by our search across these bands, we deem this final set of candidates as non-astrophysical.

## 6.8 Conclusion

We report on a search for continuous gravitational wave signals from unknown sources in binary systems using LIGO data from the first six months of the third Advanced LIGO and Advanced Virgo observing run. Four different binary parameter space regions, spanning orbital periods of 3–45 days and projected semimajor axes of 2–40 light-seconds, are searched across the 50–300 Hz frequency band. We claim no detections and estimate the sensitivity of the search in terms of the gravitational wave amplitude corresponding to the interpolated 95% detection efficiency using a simulated population of signals.

The minimum amplitude sensitivity attains an average value of  $h_0^{95\%} = (2.4 \pm 0.1) \times 10^{-25}$  in the  $f_0 = 149.5$  Hz sub-band. This is a factor of  $\sim 1.6$  lower than the lowest amplitude sensitivity obtained by a previous search performed on data from the second Advanced LIGO observing run [16]. The estimated amplitude sensitivity can be interpreted in terms of astrophysical reach and equatorial ellipticity by means of equation Eq. (6.2), as shown in Fig. 6.12.



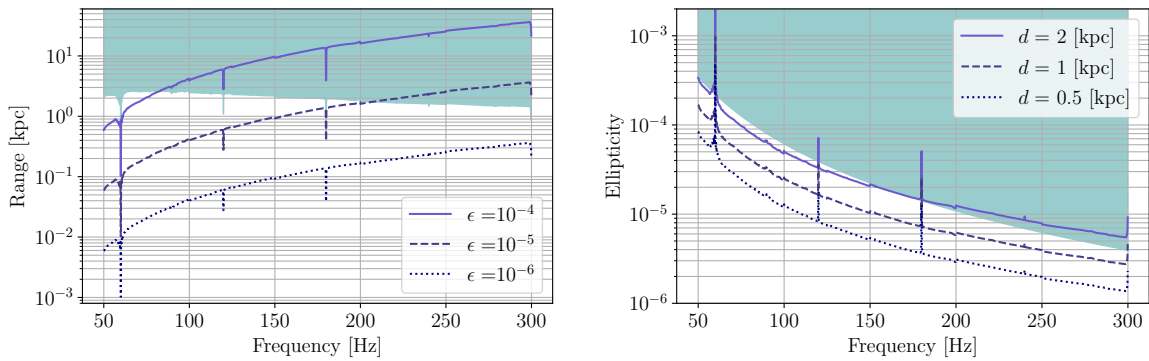


FIGURE 6.12: Astrophysical reach (left) and equatorial ellipticity (right) implied by the search sensitivity  $h_0^{95\%}$ . These results were obtained assuming a canonical neutron star moment of inertia  $I_z = 10^{38} \text{kg} \cdot \text{m}^2$ . Shaded areas denote regions excluded due to the spindown limit implied by the maximum spindown value considered in this search.

The validity of this estimation must be discussed in terms of the *spindown limit*, which corresponds to the maximum gravitational wave amplitude achievable by a neutron star assuming its rotational energy is solely lost via gravitational waves. We refer the reader to Appendix A of [6] for its definition and the relevant conversion equations.

The maximum spindown value probed by our search is  $|\dot{f}_0| \equiv (T_{\text{SFT}} \cdot T_{\text{obs}})^{-1} \simeq 6.5 \times 10^{-11} \text{Hz/s}$  [13], meaning sources braking at higher rates would not be detected by our pipeline (see Table 6.1 for the definition of  $T_{\text{SFT}}$  and  $T_{\text{obs}}$ ). Assuming the canonical emission model of a deformed NS as in Eq. (6.2), this implies the existence of a distance beyond which the required ellipticity to emit a detectable amplitude would imply a greater spindown than the one probed by the search, as long as no processes balancing the rotational energy loss are in place<sup>1</sup>. Regions excluded by the spindown limit correspond to shaded areas in Fig. 6.12.

Equatorial ellipticity values can be constrained below  $\epsilon = 10^{-5}$  for sources in binary systems such as the ones analyzed by this search located at 1 kpc emitting within the 150 – 300 Hz band. Constraints below  $\epsilon = 10^{-4}$  can be set for sources located at 2 kpc emitting within the 75 – 150 Hz band. These sensitivities approach the expected allowed maximum ellipticities of relativistic stars, which range from the order of  $10^{-6} - 10^{-7}$  to values around  $10^{-5}$  for more exotic equations of state [47].

Future enhancements of the terrestrial gravitational wave detector network will improve our sensitivity to fainter gravitational wave signals, providing a valuable tool to prospect the expected population of galactic NSs in binary systems [48, 49, 50, 51, 52].

## 6.A Frequency bands containing outliers

We provide a list of frequency bands in which outliers surviving the follow up were found. These outliers were discarded due to their inconsistent behavior with respect to an astrophysical signal, as discussed in Sec. 6.7. Table 6.8 lists frequency bands where line-crossing outliers were found. Table 6.9 corresponds to frequency bands presenting outliers discarded by the detector consistency veto. In both tables, overlapping frequency bands are merged together for the sake of compactness.

<sup>1</sup>An accretion-driven torque balance [43] could be subject to fluctuating accretion [44], leading to long-term phase wandering. This effect is unlikely to affect a semicoherent search like the one here reported, but it could have a significant impact during the follow-up stage, where longer coherence times are used [45, 46].

Min. frequency [Hz]	Max. frequency [Hz]	First timestamp [GPS]	Last timestamp [GPS]	Duration [days]	Detector	Listed
55.605	55.606	1239194624	1253831680	169	H1	Yes
57.588	57.592	1243670016	1253770240	116	H1	Yes
59.501	59.513	1238290944	1253831680	179	H1	Yes
62.474	62.478	1244523008	1251895296	85	H1	Yes
64.257	64.266	1241190400	1252681728	133	H1	No
64.284	64.291	1238955520	1253831680	172	H1	Yes
64.470	64.475	1239557632	1253770240	164	H1	Yes
64.403	64.408	1240341504	1253588992	153	H1	Yes
64.364	64.375	1238529536	1253831680	177	H1	Yes
64.415	64.417	1239375872	1253831680	167	H1	Yes
70.124	70.128	1238408192	1245484544	81	H1	Yes
80.067	80.070	1238351360	1238351360	0	H1	No
83.307	83.316	1238351360	1253831680	179	H1	Yes
83.446	83.448	1238831616	1253831680	173	H1	Yes
85.712	85.715	1239194624	1253831680	169	H1	Yes
85.964	85.965	1239738880	1253165568	155	H1	Yes
99.966	99.979	1238290944	1253831680	179	H1	Yes
107.113	107.119	1241129984	1253709824	145	H1	Yes
140.253	140.254	1238955520	1245424128	74	H1	Yes
151.800	151.800	1253105152	1253105152	0	H1	Yes
199.946	199.955	1242149888	1253770240	134	H1	Yes
213.301	213.301	1242339840	1242339840	0	L1	Yes

TABLE 6.8: Frequency bands containing line-crossing outliers. As described in Sec. 6.7.3, these could be produced because of the presence of a transient instrumental artifact or the frequency evolution of a candidate drifting away from the spectral disturbance. Overlapping frequency bands were grouped together for the sake of simplicity. Outliers belonging to this category show at least one per-segment  $\mathcal{F}$ -statistic value greater than 100. Timestamps refer to the first and last coherent segments ( $T_{\text{coh}} \simeq 17\text{h}$ ) for which at least one of those candidates showed an  $\mathcal{F}$ -statistic value over 50. The last column relates these bands to the list of unidentified lines [31].

Min. frequency [Hz]	Max. frequency [Hz]	Listed
53.709	53.721	No
55.603	55.609	Yes
57.583	57.600	Yes
62.823	62.828	Yes
64.400	64.411	Yes
83.442	83.453	Yes
85.815	85.824	No

TABLE 6.9: Frequency bands containing outliers discarded by the detector consistency veto as described in Sec. 6.7.3. Overlapping frequency bands were grouped together for the sake of simplicity. The last column relates these bands to the list of unidentified lines of the H1 detector [31].

# Bibliography

- [1] M. Sieniawska and M. Bejger, “Continuous gravitational waves from neutron stars: current status and prospects,” *Universe*, vol. 5, no. 11, p. 217, 2019.
- [2] A. Arvanitaki and S. Dubovsky, “Exploring the string axiverse with precision black hole physics,” *Physical Review D*, vol. 83, p. 044026, Feb 2011.
- [3] J. Aasi *et al.*, “Advanced LIGO,” *Classical and Quantum Gravity*, vol. 32, p. 074001, mar 2015.
- [4] F. Acernese *et al.*, “Advanced virgo: a second-generation interferometric gravitational wave detector,” *Classical and Quantum Gravity*, vol. 32, p. 024001, dec 2014.
- [5] B. P. Abbott *et al.*, “Narrow-band search for gravitational waves from known pulsars using the second ligo observing run,” *Physical Review D*, vol. 99, p. 122002, Jun 2019.
- [6] B. P. Abbott *et al.*, “Searches for gravitational waves from known pulsars at two harmonics in 2015–2017 LIGO data,” *The Astrophysical Journal*, vol. 879, p. 10, jun 2019.
- [7] B. P. Abbott *et al.*, “Search for gravitational waves from scorpius x-1 in the second advanced ligo observing run with an improved hidden markov model,” *Physical Review D*, vol. 100, p. 122002, Dec 2019.
- [8] B. P. Abbott *et al.*, “All-sky search for continuous gravitational waves from isolated neutron stars using advanced ligo o2 data,” *Physical Review D*, vol. 100, p. 024004, Jul 2019.
- [9] R. Abbott *et al.*, “Gravitational-wave Constraints on the Equatorial Ellipticity of Millisecond Pulsars,” *The Astrophysical Journal Letters*, vol. 902, no. 1, p. L21, 2020.
- [10] R. Abbott *et al.*, “Diving below the spin-down limit: Constraints on gravitational waves from the energetic young pulsar PSR J0537-6910.” arXiv:2012.12926 [astro-ph.HE].
- [11] P. R. Brady, T. Creighton, C. Cutler, and B. F. Schutz, “Searching for periodic sources with ligo,” *Physical Review D*, vol. 57, pp. 2101–2116, Feb 1998.
- [12] K. Rajwade, J. Chennamangalam, D. Lorimer, and A. Karastergiou, “The Galactic halo pulsar population,” *Monthly Notices of the Royal Astronomical Society*, vol. 479, no. 3, pp. 3094–3100, 2018.
- [13] P. B. Covas and A. M. Sintes, “New method to search for continuous gravitational waves from unknown neutron stars in binary systems,” *Physical Review D*, vol. 99, p. 124019, Jun 2019.
- [14] R. N. Manchester, G. B. Hobbs, A. Teoh, and M. Hobbs, “The Australia Telescope National Facility pulsar catalogue,” *Astronomical Journal*, vol. 129, p. 1993, 2005.
- [15] B. Krishnan, A. M. Sintes, M. A. Papa, B. F. Schutz, S. Frasca, and C. Palomba, “Hough transform search for continuous gravitational waves,” *Physical Review D*, vol. 70, p. 082001, Oct 2004.
- [16] P. B. Covas and A. M. Sintes, “First all-sky search for continuous gravitational-wave signals from unknown neutron stars in binary systems using advanced ligo data,” *Physical Review Letters*, vol. 124, p. 191102, May 2020.
- [17] P. Jaranowski and A. Królak, “Data analysis of gravitational-wave signals from spinning neutron stars. ii. accuracy of estimation of parameters,” *Physical Review D*, vol. 59, p. 063003, Feb 1999.

- [18] P. Leaci and R. Prix, “Directed searches for continuous gravitational waves from binary systems: Parameter-space metrics and optimal scorpius x-1 sensitivity,” *Physical Review D*, vol. 91, p. 102003, May 2015.
- [19] A. Buikema *et al.*, “Sensitivity and performance of the advanced ligo detectors in the third observing run,” *Physical Review D*, vol. 102, p. 062003, Sep 2020.
- [20] L. Sun *et al.*, “Characterization of systematic error in advanced LIGO calibration,” *Classical and Quantum Gravity*, vol. 37, p. 225008, oct 2020.
- [21] P. B. Covas *et al.*, “Identification and mitigation of narrow spectral artifacts that degrade searches for persistent gravitational waves in the first two observing runs of advanced ligo,” *Physical Review D*, vol. 97, p. 082002, Apr 2018.
- [22] A. Viets *et al.*, “Reconstructing the calibrated strain signal in the Advanced LIGO detectors,” *Classical and Quantum Gravity*, vol. 35, no. 9, p. 095015, 2018.
- [23] P. Astone, S. Frasca, and C. Palomba, “The short FFT database and the peak map for the hierarchical search of periodic sources,” *Classical and Quantum Gravity*, vol. 22, pp. S1197–S1210, sep 2005.
- [24] M. Pitkin, “psrqpy: a python interface for querying the ATNF pulsar catalogue,” *Journal of Open Source Software*, vol. 3, p. 538, Feb. 2018.
- [25] A. M. Sintes and B. Krishnan, “Improved hough search for gravitational wave pulsars,” *Journal of Physics: Conference Series*, vol. 32, pp. 206–211, mar 2006.
- [26] L. Sancho de la Jordana, “Hierarchical Hough all-sky search for periodic gravitational waves in LIGO S5 data,” *Journal of Physics: Conference Series*, vol. 228, p. 012004, 2010.
- [27] R. Tenorio, D. Keitel, and A. M. Sintes, “Time-frequency track distance for comparing continuous gravitational wave signals.” arXiv:2012.05752 [gr-qc].
- [28] B. P. Abbott *et al.*, “All-sky Search for Periodic Gravitational Waves in the O1 LIGO Data,” *Physical Review D*, vol. 96, no. 6, p. 062002, 2017.
- [29] B. P. Abbott *et al.*, “Full Band All-sky Search for Periodic Gravitational Waves in the O1 LIGO Data,” *Physical Review D*, vol. 97, no. 10, p. 102003, 2018.
- [30] E. Goetz, A. Neunzert, K. Riles, A. Mateas, S. Kandhasamy, J. Tasson, *et al.*, “O3a lines and combs in C00 data.” <https://dcc.ligo.org/T2000719/public>, 2020.
- [31] E. Goetz, A. Neunzert, K. Riles, A. Mateas, S. Kandhasamy, J. Tasson, *et al.*, “Unidentified O3a lines and combs in C00 data.” <https://dcc.ligo.org/T2000720/public>, 2020.
- [32] B. Behnke, M. A. Papa, and R. Prix, “Postprocessing methods used in the search for continuous gravitational-wave signals from the galactic center,” *Physical Review D*, vol. 91, p. 064007, Mar 2015.
- [33] C. Dreissigacker, R. Prix, and K. Wette, “Fast and accurate sensitivity estimation for continuous-gravitational-wave searches,” *Physical Review D*, vol. 98, p. 084058, Oct 2018.
- [34] B. Abbott *et al.*, “All-sky search for periodic gravitational waves in ligo s4 data,” *Physical Review D*, vol. 77, p. 022001, Jan 2008.
- [35] P. Virtanen *et al.*, “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python,” *Nature Methods*, vol. 17, pp. 261–272, 2020.
- [36] G. Ashton and R. Prix, “Hierarchical multistage MCMC follow-up of continuous gravitational wave candidates,” *Physical Review D*, vol. 97, no. 10, p. 103020, 2018.
- [37] D. Keitel, R. Tenorio, G. Ashton, and R. Prix, “PyFstat: a Python package for continuous gravitational-wave data analysis.” arXiv:2101.10915 [gr-qc].

- [38] G. Ashton, D. Keitel, R. Prix, and R. Tenorio, “PyFstat v1.7.3,” 2020. 10.5281/zenodo.4041983.
- [39] LIGO Scientific Collaboration, “LIGO Algorithm Library - LALSuite.” 10.7935/GT1W-FZ16, 2018.
- [40] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, “emcee: The MCMC hammer,” *Publications of the Astronomical Society of the Pacific*, vol. 125, pp. 306–312, mar 2013.
- [41] W. D. Vousden, W. M. Farr, and I. Mandel, “Dynamic temperature selection for parallel tempering in Markov chain Monte Carlo simulations,” *Monthly Notices of the Royal Astronomical Society*, vol. 455, pp. 1919–1937, 11 2015.
- [42] D. Keitel, “Robust semicoherent searches for continuous gravitational waves with noise and signal models including hours to days long transients,” *Physical Review D*, vol. 93, no. 8, p. 084024, 2016.
- [43] L. Bildsten, “Gravitational radiation and rotation of accreting neutron stars,” *Astrophysical Journal Letters*, vol. 501, p. L89, 1998.
- [44] A. Watts, B. Krishnan, L. Bildsten, and B. F. Schutz, “Detecting gravitational wave emission from the known accreting neutron stars,” *Monthly Notices of the Royal Astronomical Society*, vol. 389, pp. 839–868, 2008.
- [45] A. Mukherjee, C. Messenger, and K. Riles, “Accretion-induced spin-wandering effects on the neutron star in Scorpius X-1: Implications for continuous gravitational wave searches,” *Physical Review D*, vol. 97, no. 4, p. 043016, 2018.
- [46] G. Ashton, R. Prix, and D. I. Jones, “Statistical characterization of pulsar glitches and their potential impact on searches for continuous gravitational waves,” *Physical Review D*, vol. 96, no. 6, p. 063004, 2017.
- [47] N. K. Johnson-McDaniel and B. J. Owen, “Maximum elastic deformations of relativistic stars,” *Physical Review D*, vol. 88, p. 044004, 2013.
- [48] T. Akutsu *et al.*, “KAGRA: 2.5 generation interferometric gravitational wave detector,” *Nature Astronomy*, vol. 3, pp. 35–40, 2019.
- [49] M. Maggiore *et al.*, “Science Case for the Einstein Telescope,” *Journal of Cosmology and Astroparticle Physics*, vol. 03, p. 050, 2020.
- [50] D. Reitze *et al.*, “Cosmic Explorer: The U.S. Contribution to Gravitational-Wave Astronomy beyond LIGO,” *Bulletin of the American Astronomical Society*, vol. 51, p. 035, 7 2019.
- [51] J. Miller, L. Barsotti, S. Vitale, P. Fritschel, M. Evans, and D. Sigg, “Prospects for doubling the range of Advanced LIGO,” *Physical Review D*, vol. 91, p. 062005, 2015.
- [52] B. P. Abbott *et al.*, “Prospects for observing and localizing gravitational-wave transients with Advanced LIGO, Advanced Virgo and KAGRA,” *Living Reviews in Relativity*, vol. 23, no. 1, p. 3, 2020.



## Chapter 7

# All-sky search for continuous gravitational waves from isolated neutron stars using Advanced LIGO and Advanced Virgo O3 data

This chapter is an adaptation of the material presented in

*All-sky search for continuous gravitational waves from isolated neutron stars using Advanced LIGO and Advanced Virgo O3 data*

R. Abbott et al. (LIGO–Virgo–KAGRA Collaboration)

[Phys. Rev. D 106, 102008 \(2022\)](#) – [arXiv:2201.00697 \[gr-qc\]](#)

DOI: 10.1103/PhysRevD.106.102008

where details about search pipelines other than `SkyHough` and the interpretation of upper limits in terms of primordial-black-hole abundances have been removed as they were produced by other authors and thus are not part of this thesis. The text has been adapted to focus the discussion of the setup on the `SkyHough` pipeline but include all the involved searches whenever introducing the topic and discussing the results.

### 7.1 Introduction

The Advanced LIGO [1] and Advanced Virgo [2] detectors have made numerous detections of gravitational waves (GW), to date consisting of short-duration (transient) GW emitted during the inspirals and mergers of compact binary systems of black holes (BH), neutron stars (NS), [3, 4], as well as mixed NS-BH binaries [5]. Among still undiscovered types of GW radiation are long-lasting, almost-monochromatic continuous waves (CW), whose amplitudes and frequencies change much more slowly compared to those of transient sources (on the timescale of years rather than seconds). Astrophysically, promising sources of CW are rotating, non-axisymmetric NS, emitting GW at a frequency close to, or related to, their spin frequency. Deviations from the symmetry (a NS ‘deformation’) may be caused by fluid instabilities, such as in the case of r-modes, or by elastic, thermal or magnetic stresses in the crust and/or core of NS, and may be acquired at various stages of stars’ isolated evolution, or during an interaction with a companion in a binary system (for recent reviews on sources of CW, see e.g., [6, 7, 8]). Discovery of CW emitted by NS would allow to probe their still mysterious interiors, study properties of dense matter in conditions distinct from those occurring in inspirals and mergers of binary NS systems, as well as carry out additional tests of the theory of gravity [9]. Due to intrinsically smaller GW amplitude of CW in comparison to the already-detected transient sources, searches for CW from rotating non-axisymmetric NS are essentially limited to the Galaxy.

Searches for continuous waves are usually split in three different domains: *targeted searches* look for signals from known pulsars; *directed searches* look for signals from known sky locations; *all-sky searches* look for signals from unknown sources. All-sky searches for *a priori* unknown CW sources have been carried out in the Advanced LIGO and Advanced Virgo data previously [10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 22]. A recent review on pipelines for wide parameter-space searches can be found in [23].

Here we report on results from an all-sky, broad frequency range search by the LIGO–Virgo–KAGRA collaboration using the most-sensitive data to date, the LIGO–Virgo O3 observing run, employing four different search pipelines: the `FrequencyHough` [24], `SkyHough` [25], `Time-domain  $\mathcal{F}$ -statistic` [26, 27], and `SOAP` [28]. The main contribution of this thesis is the setup and production of the results obtained by the `SkyHough` pipeline. Each pipeline uses different data analysis methods and covers different regions of the frequency and frequency time derivative parameter space, although there exist overlaps between them (see Table 7.1 and Fig. 7.1 for details). The search is performed for frequencies between 10 Hz and 2048 Hz and for a range of frequency time derivative between  $-10^{-8}$  Hz/s and  $10^{-9}$  Hz/s, covering the whole sky. We note here that the search is generally-agnostic to the type of the GW source, so the results are not actually limited to signals from non-axisymmetric rotating NS in our Galaxy. A comprehensive multi-stage analysis of the signal outliers obtained by the four pipelines has not revealed any viable candidate for a continuous GW signal. However we improve the broad-range frequency upper limits with respect to previous O1 and O2 observing run and also with respect to the recent analysis of the first half of the O3 run [18]. This is also the first all-sky search for CW sources that uses the Advanced Virgo detector’s data.

The article is organized as follows: in Section 7.2 we describe the O3 observing run and provide details about the data used. Section 7.3 we present an overview of common aspects among the different pipelines used in this search. Section 7.4, details and results of the `SkyHough` pipeline are discussed. Section 7.5 contains a discussion of the astrophysical implications of our results.

## 7.2 Data sets used

The data set used in this analysis was the third observing run (O3) of the Advanced LIGO and Advanced Virgo GW detectors [1, 2]. LIGO is made up of two laser interferometers, both with 4 km long arms. One is at the LIGO Livingston Observatory (L1) in Louisiana, USA and the other is at the LIGO Hanford Observatory (H1) in Washington, USA. Virgo (V1) consists of one interferometer with 3 km arms located at European Gravitational Observatory (EGO) in Cascina, Italy. The O3 run took place between the 2019 April 1 and the 2020 March 27. The run was divided into two parts, O3a and O3b, separated by one month commissioning break that took place in October 2019. The duty factors for this run were  $\sim 76\%$ ,  $\sim 71\%$ ,  $\sim 76\%$  for L1, H1, V1 respectively. The maximum uncertainties (68% confidence interval) on the calibration of the LIGO data were of 7%/11% in magnitude and 4 deg/9 deg in phase for O3a/O3b data ([29, 30]). For Virgo, it amounted to 5% in amplitude and 2 deg in phase, with the exception of the band 46 - 51 Hz, for which the maximum uncertainty was estimated as 40% in amplitude and 34 deg in phase during O3b. For the smaller range 49.5 - 50.5 Hz, the calibration was unreliable during the whole run [31].

## 7.3 Common Aspects of Search Pipelines

### 7.3.1 Signal model

The GW signal in the detector frame from an isolated, asymmetric NS spinning around one of its principal axis of inertia is given by [26]:

$$h(t) = h_0 [F_+(t, \alpha, \delta, \psi) \frac{1 + \cos^2 \iota}{2} \cos \phi(t) + F_\times(t, \alpha, \delta, \psi) \cos \iota \sin \phi(t)], \quad (7.1)$$

where  $F_+$  and  $F_\times$  are the antenna patterns of the detectors dependent on right ascension  $\alpha$ , declination  $\delta$  of the source and polarization angle  $\psi$ ,  $h_0$  is the amplitude of the signal,  $\iota$  is the angle between the total angular momentum vector of the star and the direction from the star to the Earth, and  $\phi(t)$  is the phase of the signal. The amplitude of the signal is given by:

$$h_0 = \frac{4\pi^2 G}{c^4} \frac{\epsilon I_{zz} f^2}{d} \approx 1.06 \times 10^{-26} \left( \frac{\epsilon}{10^{-6}} \right) \times \left( \frac{I_{zz}}{10^{38} \text{ kg m}^2} \right) \left( \frac{f}{100 \text{ Hz}} \right)^2 \left( \frac{1 \text{ kpc}}{d} \right), \quad (7.2)$$



where  $d$  is the distance from the detector to the source,  $f$  is the GW frequency (assumed to be twice the rotation frequency of the NS),  $\epsilon$  is the ellipticity or asymmetry of the star, given by  $(I_{xx} - I_{yy})/I_{zz}$ , and  $I_{zz}$  is the moment of inertia of the star with respect to the principal axis aligned with the rotation axis.

We assume that the phase evolution of the GW signal can be approximated with a second order Taylor expansion around a fiducial reference time  $\tau_r$ :

$$\phi(\tau) = \phi_o + 2\pi[f(\tau - \tau_r) + \frac{\dot{f}}{2!}(\tau - \tau_r)^2], \quad (7.3)$$

where  $\phi_o$  is an initial phase and  $f$  and  $\dot{f}$  are the frequency and first frequency derivative at the reference time. The relation between the time at the source  $\tau$  and the time at the detector  $t$  is given by:

$$\tau(t) = t + \frac{\vec{r}(t) \cdot \vec{n}}{c} + \Delta_{E\odot} - \Delta_{S\odot}, \quad (7.4)$$

where  $\vec{r}(t)$  is the position vector of the detector in the Solar System Barycenter (SSB) frame, and  $\vec{n}$  is the unit vector pointing to the NS;  $\Delta_{E\odot}$  and  $\Delta_{S\odot}$  are respectively the relativistic Einstein and Shapiro time delays. In standard equatorial coordinates with right ascension  $\alpha$  and declination  $\delta$ , the components of the unit vector  $\vec{n}$  are given by  $(\cos \alpha \cos \delta, \sin \alpha \cos \delta, \sin \delta)$ .

### 7.3.2 Parameter space analyzed

All the pipelines involved in this analysis perform an all-sky search, however the frequency and frequency derivative ranges analyzed are different for each pipeline. The detailed ranges analyzed by the four pipelines are summarized in Table 7.1 and presented in Fig. 7.1. The `FrequencyHough` pipeline analyzes a broad frequency range between 10 Hz and 2048 Hz and a broad frequency time derivative range between  $-10^{-8}$  Hz/s and  $10^{-9}$  Hz/s. A very similar range of  $f$  and  $\dot{f}$  is analyzed by `SOAP` pipeline. The `SkyHough` pipeline analyzes a narrower frequency range where the detectors are most sensitive whereas `Time-domain  $\mathcal{F}$ -statistic` pipeline analyzes  $f$  and  $\dot{f}$  ranges of the bulk of the observed pulsar population.

### 7.3.3 Detection statistics

As all-sky searches cover a large parameter space they are computationally very expensive and it is computationally prohibitive to analyze coherently the data from the full observing run using optimal matched-filtering. As a result each of the pipelines developed for the analysis uses a semi-coherent method. Moreover to reduce the computer memory and to parallelize the searches the data are divided into narrow bands. Each analysis begins with sets of *short Fourier transforms* (SFTs) that span the observation period, with coherence times ranging from 1024s to 8192s. The `FrequencyHough`, `SkyHough` and `SOAP` pipelines compute measures of strain power directly from the SFTs and create detection statistics by stacking those powers with corrections for frequency evolution applied. The `FrequencyHough` and `SkyHough` pipelines use *Hough* transform to do the stacking whereas `SOAP` pipeline uses the *Viterbi* algorithm. The `Time-domain  $\mathcal{F}$ -statistic` pipeline extracts band-limited 6-day long time-domain data segments from the SFT sets and applies frequency evolution corrections coherently to obtain the  $\mathcal{F}$ -statistic [26]. Coincidences are then required among multiple data segments with no stacking.

### 7.3.4 Outlier follow-up

All four pipelines perform a follow-up analysis of the statistically significant candidates (outliers) obtained during the search. All pipelines perform vetoing of the outliers corresponding to narrow, instrumental artifacts (lines) in the advanced LIGO detectors [32]. Several other consistency vetoes are also applied to eliminate outliers. The `FrequencyHough`, `SkyHough`, and `Time-domain  $\mathcal{F}$ -statistic` pipelines perform follow-up of the candidates by processing the data with increasing long coherence times whereas `SOAP` pipeline use *convolutional neural networks* to do the post processing.




FIGURE 7.1: Frequency and frequency derivative search ranges of the four pipelines: the FrequencyHough pipeline ranges marked in grey, SkyHough in red, Time-domain  $\mathcal{F}$ -statistic in blue, and SOAP in magenta. See Table 7.1 for details.

Pipeline	$f$ [Hz]	$\dot{f}$ [Hz/s]
FrequencyHough	10 – 2048	$-10^{-8} - 10^{-9}$
SkyHough	65 – 350	$-10^{-9} - 5 \times 10^{-12}$
SOAP	40 – 1000	$-10^{-9} - 10^{-9}$
	1000 – 2000	$-10^{-8} - 10^{-8}$
Time-domain	20 – 200	$-3.2 \times 10^{-9} f/100 - 0$
$\mathcal{F}$ -statistic	200 – 750	$-2 \times 10^{-10} - 2 \times 10^{-11}$

TABLE 7.1: Frequency and frequency derivative search ranges of the four pipelines.

Parameter	Resolution
$\delta f$	$1.4 \times 10^{-4}$ Hz
$\delta \dot{f}$	$5 \times 10^{-12}$ Hz/s
$\delta \theta$	$0.69$ Hz/ $f$

TABLE 7.2: Parameter-space resolutions employed by the SkyHough pipeline.

### 7.3.5 Upper limits

No periodic gravitational wave signals were observed by any of the four pipelines and all the pipelines obtain upper limits on their strength. The three pipelines SkyHough, Time-domain  $\mathcal{F}$ -statistic and SOAP obtain the upper limits by injections of the signals according to the model given in Section 7.3.1 above for an array of signal amplitudes  $h_0$  and randomly choosing the remaining parameters. The FrequencyHough pipeline obtains upper limits using an analytic formula that depends on the spectral density of the noise of the detector. The formula was validated by a number of tests consisting of injecting signals to the data.

## 7.4 The SkyHough search

SkyHough [25, 33] is a semicoherent pipeline based on the Hough transform to look for CW signals from isolated neutron stars. Several versions of this pipeline have been used throughout the initial [34, 35] and advanced [10, 11] detector era, as well as to look for different kinds of signals such as CW from neutron stars in binary systems [36, 19, 20] or long-duration GW transients [37]. The current implementation of SkyHough closely follows that of [11] and includes an improved suite of post-processing and follow-up stages [38, 39, 40].

### 7.4.1 Parameter space

The SkyHough pipeline searches over the standard four parameters describing a CW signal from isolated NS: frequency  $f$ , spin-down  $\dot{f}$  and sky position, parametrized using equatorial coordinates  $\alpha, \delta$ .

Parameter-space resolutions are given in [25]

$$\delta f = \frac{1}{T_{\text{SFT}}}, \quad \delta \dot{f} = \frac{\delta f}{T_{\text{obs}}}, \quad \delta \theta = \frac{c/v}{T_{\text{SFT}} P_f f}, \quad (7.5)$$

where  $\theta$  represents either of the sky angles,  $v/c \simeq 10^{-4}$  represents the average detector velocity as a fraction of the speed of light, and the pixel factor  $P_f = 2$  is a tunable overresolution parameter. Table 7.2 summarizes the numerical values employed in this search.

The SkyHough all-sky search covers the most sensitive frequency band of the advanced LIGO detectors, between 65 Hz and 350 Hz. This band is further sub-divided into  $\Delta f = 0.025$  Hz sub-bands, resulting in a total of 11400 frequency bands. Spin-down values are covered from  $-1 \times 10^{-9}$  Hz/s to  $5 \times 10^{-12}$  Hz/s, which include typical spin-up values associated to CW emission from the evaporation of boson clouds around black holes [41].

### 7.4.2 Description of the search

The first stage of the SkyHough pipeline performs a multi-detector search using H1 and L1 SFTs with  $T_{\text{SFT}} = 7200s$ <sup>1</sup>. Each 0.025 Hz sub-band is analyzed separately using the same two step strategy as in [11, 20]: parameter-space is efficiently analyzed using SkyHough's look-up table approach; the top 0.1% most significant candidates are further analyzed using a more sensitive statistic. The result for each frequency sub-band is a toplist containing the  $10^5$  most significant candidates across the sky and spin-down parameter-space.

<sup>1</sup>This choice of  $T_{\text{SFT}}$  is beyond the standard prescription proposed in [25], which however uses rather conservative arguments to establish a safe bound.

Stage	0	1	2	3	4	5
$N_{\text{seg}}$	660	330	92	24	4	1
$T_{\text{coh}}$	0.5 day	1 day	4 days	15 days	90 days	360 days

TABLE 7.3: Coherence-time configuration of the multi-stage follow-up employed by the SkyHough pipeline. The data stream is divided into a fix number of segments of the same length; the reported coherence time is an approximate value obtained by dividing the observation time by the number of segments at each stage.

Each toplist is then clustered using a novel approach presented in [38] and firstly applied in [20]. A parameter-space distance is defined using the average mismatch in frequency evolution between two different parameter-space templates

$$d(\vec{\lambda}, \vec{\lambda}_*) = \frac{T_{\text{SFT}}}{N_{\text{SFT}}} \sum_{\alpha=0}^{N_{\text{SFT}}} \left| f(t_\alpha; \vec{\lambda}) - f(t_\alpha; \vec{\lambda}_*) \right|, \quad (7.6)$$

where  $f(t; \vec{\lambda})$  is defined as

$$f(t; \vec{\lambda}) = \left[ f + (t - t_{\text{ref}}) \cdot \dot{f} \right] \cdot \left[ 1 + \frac{\vec{v}(t) \cdot \vec{n}}{c} \right] \quad (7.7)$$

and  $\vec{\lambda} = \{f, \dot{f}, \alpha, \delta\}$  refers to the phase-evolution parameters of the template.

Clusters are constructed by pairing together templates in consecutive frequency bins such that  $d(\vec{\lambda}, \vec{\lambda}_*) \leq 1$ . Each cluster is characterized by its most significant element (the *loudest* element). From each 0.025 Hz sub-band, we retrieve the forty most significant clusters for further analysis. This results in a total of 456000 candidates to follow-up.

The loudest cluster elements are first sieved through the *line veto*, a standard tool to discard clear instrumental artifacts using the list of known, narrow, instrumental artifacts (lines) in the advanced LIGO detectors [32]: If the instantaneous frequency of a candidate overlaps with a frequency band containing an instrumental line of known origin, the candidate is ascribed an instrumental origin and consequently ruled out.

Surviving candidates are then followed-up using PyFstat, a Python package implementing a Markov-chain Monte Carlo (MCMC) search for CW signals [42, 39]. The follow-up uses the  $\mathcal{F}$ -statistic as a (log) Bayes factor to sample the posterior probability distribution of the phase-evolution parameters around a certain parameter-space region

$$P(\vec{\lambda}|x) \propto e^{\mathcal{F}(\vec{\lambda};x)} \cdot P(\vec{\lambda}), \quad (7.8)$$

where  $P(\vec{\lambda})$  represents the prior probability distribution of the phase-evolution parameters. The  $\mathcal{F}$ -statistic, as opposed to the SkyHough number count, allows us to use longer coherence times, increasing the sensitivity of the follow-up with respect to the main search stage.

As initially described in [42], the effectiveness of an MCMC follow-up is tied to the number of templates covered by the initial prior volume, suggesting a hierarchical approach: coherence time should be increased following a ladder so that the follow-up is able to converge to the true signal parameters at each stage. We follow the proposal in [40] and compute a coherence-time ladder using  $\mathcal{N}^* = 10^3$  (see Eq. (31) of [42]) starting from  $T_{\text{coh}} = 1$  day including an initial stage of  $T_{\text{coh}} = 0.5$  days. The resulting configuration is collected in Table 7.3.

The first follow-up stage is similar to that employed in [19, 20]: an MCMC search around the loudest candidate of the selected clusters is performed using a coherence time of  $T_{\text{coh}} = 0.5$  days. Uniform priors containing 4 parameter-space bins in each dimension are centered around the loudest candidate. A threshold is calibrated using an injection campaign: any candidate whose loudest  $2\mathcal{F}$  value over the MCMC run is lower than  $2\mathcal{F} = 3450$  is deemed inconsistent with CW signal.

The second follow-up stage is a variation of the method described in [40], previously applied to [43, 44]. For each outlier surviving the initial follow-up stage (stage 0 in Table 7.3), we construct a Gaussian prior using the median and inter-quartile range of the posterior samples and run the next-stage MCMC follow-up. The resulting maximum  $2\mathcal{F}$  is then compared to the expected  $2\mathcal{F}$  inferred from

Comparing stages	$(2\mathcal{F} - \mu)/\sigma$ bracket
Stage 0 vs. Stage 1	(-1.79, 1.69)
Stage 1 vs. Stage 2	(-1.47, 1.35)
Stage 2 vs. Stage 3	(-0.94, 0.80)
Stage 3 vs. Stage 4	(-0.63, 0.42)
Stage 4 vs. Stage 5	(-0.34, 0.11)

TABLE 7.4:  $2\mathcal{F}$  consistency brackets employed in the multi-stage follow-up of the SkyHough pipeline. Brackets were computed using a campaign of 500 software-injected signals representing an isotropic population of uniformly sky-distributed NS at 150 representative frequency bands with an amplitude corresponding to the  $h_0^{95\%}$  sensitivity estimation. The implied false dismissal probability is  $\lesssim 1/(150 \times 500) \simeq 1.3 \times 10^{-5}$ . Stages correspond to those described in Table 7.3.

the previous MCMC follow-up stage. Highly-discrepant candidates are deemed inconsistent with a CW signal and hence discarded.

Given an MCMC stage using  $\hat{N}$  segments from which a value of  $2\hat{\mathcal{F}}$  is recovered, the distribution of  $2\mathcal{F}$  values using  $N$  segments is well approximated by

$$P(2\mathcal{F}|N, 2\hat{\mathcal{F}}, \hat{N}) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{1}{2}\left(\frac{2\mathcal{F}-\mu}{\sigma}\right)^2}, \quad (7.9)$$

where

$$\mu = \rho_0^2 + 4N, \quad (7.10)$$

$$\sigma^2 = 8 \cdot (N + \hat{N} + \rho_0^2), \quad (7.11)$$

and  $\rho_0^2 = 2\hat{\mathcal{F}} - 4\hat{N}$  is a proxy for the (squared) SNR [45]. Equation (7.9) is exact in the limit of  $N, \hat{N} \gg 1$  or  $\rho_0^2 \gg 1$ . In this search, however, we calibrate a bracket on  $(2\mathcal{F} - \mu)/\sigma$  for each follow-up stage using an injection campaign, shown in Table 7.4. Candidates outside of the bracket are deemed inconsistent with a CW signal.

Any surviving candidates are subject to manual inspection in search for obvious instrumental causes such as hardware-injected artificial signals or narrow instrumental artifacts.

### 7.4.3 Candidate follow-up

Table 7.5 summarizes the number of outliers discarded by each of the veto and follow-up stages employed in this search. A total of 36 candidates survive the complete suite of veto and follow-up stages of the SkyHough pipeline. Candidates can be grouped into two sets according to their corresponding  $\mathcal{F}$ -statistic value: 31 candidates present a value of  $2\hat{\mathcal{F}} \sim \mathcal{O}(10^3)$ , while the remaining 5 candidate only achieve  $2\hat{\mathcal{F}} \sim \mathcal{O}(30)$ . Their corresponding parameters are collected in Table 7.6.

The 31 strong candidates present consistent values with the only two hardware injections within the SkyHough search range: 24 candidates are ascribed to the hardware injection ip0, while 7 candidates are ascribed to the hardware injection ip3. Parameter deviation of the loudest candidate associated to each injection are reported in Table 7.7.

The five weaker candidates are manually inspected using the segment-wise  $\mathcal{F}$ -statistic on 660 coherent segments, in a similar manner to that in [18, 40].

The first pair of candidates is found around 85.850 Hz, where the H1 detector presents a broad spectral feature. As shown in Fig 7.2, their single-detector  $\mathcal{F}$ -statistic is more prominent in the H1 detector rather than the L1 detector, and scores over the multi-detector  $\mathcal{F}$ -statistic. These characteristics point towards an instrumental, rather than astrophysical, origin.

A second pair of candidates is found around 95.7 Hz. This frequency band is populated by narrow spectral artifacts of unknown origin in the H1 detector. Correspondingly, as shown in Fig. 7.3, the single-detector  $\mathcal{F}$  statistic is prominent in the H1 detector rather than the L1 detector. Due to the narrowness of the feature, in this case the accumulation is better localized around a fraction of the run. As in the previous case, the single-detector  $\mathcal{F}$ -statistic scores over the multi-detector  $\mathcal{F}$ -statistic. These characteristics point towards an instrumental origin.

Search stage	Candidates	% removed
Clustering	456000	
Line veto	414459	9%
$2\mathcal{F}$ threshold	3767	99%
Stage 0 v.s. Stage 1	697	18%
Stage 1 v.s. Stage 2	172	75%
Stage 3 v.s. Stage 3	90	48%
Stage 3 v.s. Stage 4	48	47%
Stage 4 v.s. Stage 5	36	25%

TABLE 7.5: Summary of candidates processed by each of the veto and follow-up stages of the SkyHough search.

Band	Candidate	$f$ [Hz]	$\dot{f}$ [nHz/s]	$\alpha$ [rad]	$\delta$ [rad]	$2\tilde{\mathcal{F}}$	Comment
834	4	85.872761414	$2.41584 \cdot 10^{-3}$	3.143782737	1.165116066	30.54	Broad spectral feature in H1
834	9	85.873653124	$-9.35774 \cdot 10^{-2}$	3.409549407	1.385107830	36.25	Broad spectral feature in H1
1227	35	95.697667346	$-4.89489 \cdot 10^{-2}$	1.593327050	-1.292111453	31.53	Narrow spectral feature in H1
1229	5	95.725474979	$-9.63949 \cdot 10^{-1}$	0.260240661	-1.008336167	30.87	Narrow spectral feature in H1
1754	1	108.857159405	$-8.04825 \cdot 10^{-7}$	3.113189707	-0.583577133	1055.70	Hardware injection ip3
1754	2	108.857159406	$-8.29209 \cdot 10^{-7}$	3.113189734	-0.583577139	1055.69	Hardware injection ip3
1754	5	108.857159404	$-7.43862 \cdot 10^{-7}$	3.113189647	-0.583577277	1055.71	Hardware injection ip3
1754	10	108.857159405	$-7.92726 \cdot 10^{-7}$	3.113189663	-0.583577189	1055.71	Hardware injection ip3
1754	13	108.857159406	$-8.38377 \cdot 10^{-7}$	3.113189745	-0.583577097	1055.69	Hardware injection ip3
1754	14	108.857159405	$-8.14434 \cdot 10^{-7}$	3.113189656	-0.583577155	1055.69	Hardware injection ip3
1754	34	108.857159404	$-7.09929 \cdot 10^{-7}$	3.113189613	-0.583577327	1055.69	Hardware injection ip3
7251	10	246.297680589	$-2.24806 \cdot 10^{-2}$	1.425124776	-1.242786654	35.79	Narrow spectral feature in H1
8022	0	265.575086278	$-4.14962 \cdot 10^{-3}$	1.248816426	-0.981180252	1543.70	Hardware injection ip0
8022	1	265.575086279	$-4.14969 \cdot 10^{-3}$	1.248816468	-0.981180265	1543.68	Hardware injection ip0
8022	2	265.575086278	$-4.14961 \cdot 10^{-3}$	1.248816419	-0.981180239	1543.69	Hardware injection ip0
8022	3	265.575086278	$-4.14964 \cdot 10^{-3}$	1.248816434	-0.981180252	1543.69	Hardware injection ip0
8022	4	265.575086278	$-4.14964 \cdot 10^{-3}$	1.248816444	-0.981180252	1543.70	Hardware injection ip0
8022	5	265.575086277	$-4.14958 \cdot 10^{-3}$	1.248816405	-0.981180243	1543.70	Hardware injection ip0
8022	7	265.575086279	$-4.14968 \cdot 10^{-3}$	1.248816456	-0.981180263	1543.69	Hardware injection ip0
8022	28	265.575086278	$-4.14965 \cdot 10^{-3}$	1.248816441	-0.981180257	1543.69	Hardware injection ip0
8023	0	265.575086278	$-4.14964 \cdot 10^{-3}$	1.248816439	-0.981180255	1543.70	Hardware injection ip0
8023	1	265.575086278	$-4.14961 \cdot 10^{-3}$	1.248816417	-0.981180250	1543.70	Hardware injection ip0
8023	3	265.575086278	$-4.14966 \cdot 10^{-3}$	1.248816464	-0.981180249	1543.68	Hardware injection ip0
8023	4	265.575086279	$-4.14969 \cdot 10^{-3}$	1.248816466	-0.981180264	1543.68	Hardware injection ip0
8023	7	265.575086279	$-4.14967 \cdot 10^{-3}$	1.248816448	-0.981180256	1543.69	Hardware injection ip0
8023	8	265.575086279	$-4.14966 \cdot 10^{-3}$	1.248816453	-0.981180260	1543.71	Hardware injection ip0
8023	9	265.575086278	$-4.14963 \cdot 10^{-3}$	1.248816431	-0.981180254	1543.70	Hardware injection ip0
8023	10	265.575086275	$-4.14945 \cdot 10^{-3}$	1.248816284	-0.981180203	1543.26	Hardware injection ip0
8023	11	265.575086278	$-4.14962 \cdot 10^{-3}$	1.248816419	-0.981180255	1543.69	Hardware injection ip0
8023	12	265.575086278	$-4.14963 \cdot 10^{-3}$	1.248816435	-0.981180249	1543.70	Hardware injection ip0
8023	13	265.575086277	$-4.14956 \cdot 10^{-3}$	1.248816392	-0.981180234	1543.66	Hardware injection ip0
8023	14	265.575086278	$-4.14966 \cdot 10^{-3}$	1.248816450	-0.981180252	1543.70	Hardware injection ip0
8023	16	265.575086278	$-4.14962 \cdot 10^{-3}$	1.248816403	-0.981180252	1543.65	Hardware injection ip0
8023	18	265.575086278	$-4.14962 \cdot 10^{-3}$	1.248816430	-0.981180248	1543.66	Hardware injection ip0
8023	19	265.575086278	$-4.14963 \cdot 10^{-3}$	1.248816436	-0.981180254	1543.72	Hardware injection ip0
8023	34	265.575086278	$-4.14965 \cdot 10^{-3}$	1.248816452	-0.981180250	1543.72	Hardware injection ip0

TABLE 7.6: Surviving candidates of the SkyHough multi-stage MCMC follow-up using PyFstat.  $2\tilde{\mathcal{F}}$  corresponds to the loudest fully-coherent  $\mathcal{F}$ -statistic value of the MCMC run. Band index corresponds to a frequency of  $(65 + 0.025 \times \text{Band})$  Hz. Reference time is GPS 1238166018.

Injection	$2\tilde{\mathcal{F}}$	$\Delta f$ [Hz]	$\Delta \dot{f}$ [nHz/s]	$\Delta \alpha$ [rad]	$\Delta \delta$ [rad]	$\Delta \alpha$ [deg]	$\Delta \delta$ [deg]
ip0	1543.72	$-4.80 \times 10^{-9}$	$3.52 \times 10^{-7}$	$-2.82 \times 10^{-7}$	$-2.49 \times 10^{-8}$	$-1.62 \times 10^{-5}$	$-1.43 \times 10^{-5}$
ip3	1055.71	$1.16 \times 10^{-8}$	$-7.29 \times 10^{-7}$	$9.35 \times 10^{-7}$	$1.53 \times 10^{-6}$	$5.35 \times 10^{-5}$	$8.74 \times 10^{-5}$

TABLE 7.7: Hardware injection recovery by the **SkyHough** pipeline. For each hardware injection within search range we report the dimension-wise errors with respect to loudest surviving candidate of the follow-up.

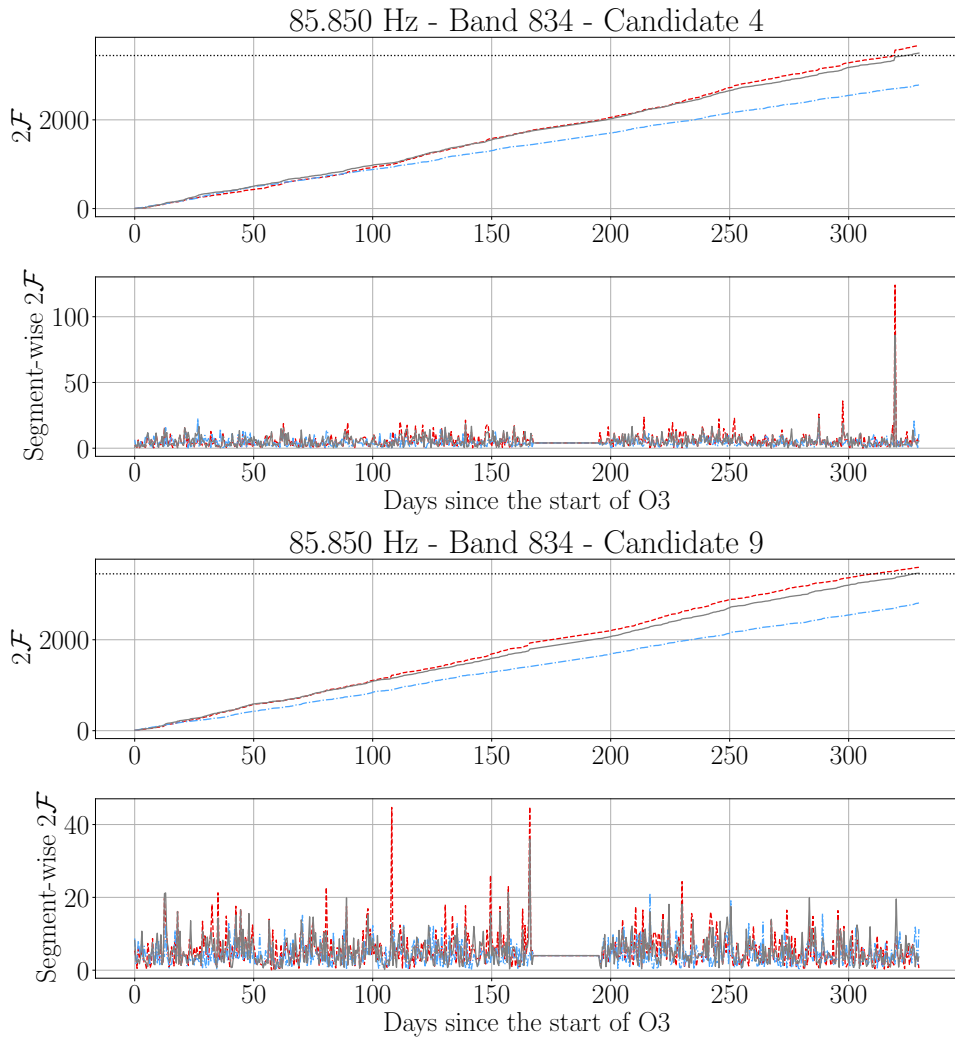


FIGURE 7.2: **SkyHough** candidates consistent with a broad spectral artifact in the H1 detector. Upper panel shows the cumulative semicoherent  $\mathcal{F}$ -statistic using 660 coherent segments ( $T_{\text{coh}} = 0.5$  days). Lower panel shows the segment-wise  $\mathcal{F}$ -statistic. Dashed red line represents the single-detector  $\mathcal{F}$ -statistic using H1-only data; dot-dashed blue line represents the single-detector  $\mathcal{F}$ -statistic using L1-only data. Solid gray line represents the multi-detector  $\mathcal{F}$ -statistic. Dotted horizontal line represents the threshold of  $2\mathcal{F} = 3450$  set at the initial follow-up stage.

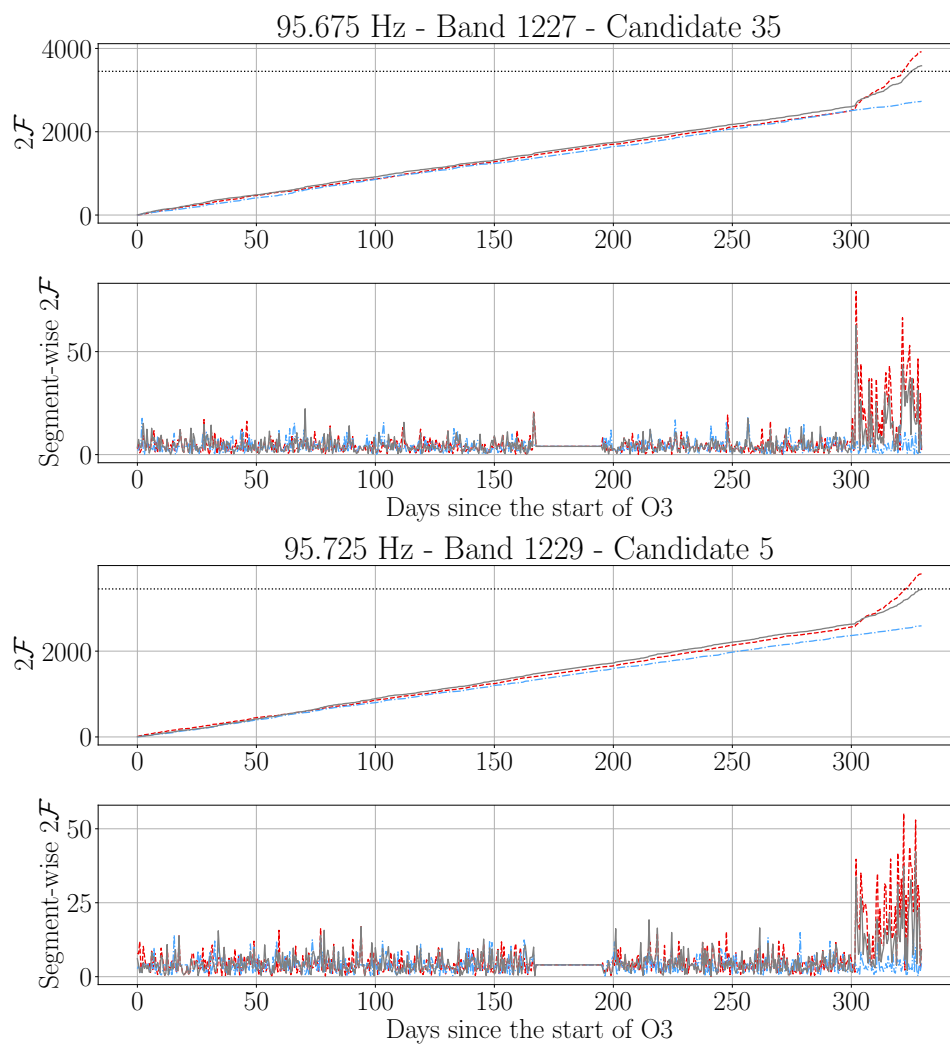


FIGURE 7.3: SkyHough candidates consistent with two narrow spectral artifacts of unknown origin in the H1 detector. The legend is equivalent to that of Fig. 7.2.






FIGURE 7.4: SkyHough candidates consistent with a narrow spectral artifact of unknown origin in the H1 detector. The legend is equivalent to that of Fig. 7.2.

The last weak candidate in the vicinity of 246.275 Hz, where the H1 detector presents another narrow spectral artifact of unknown origin, is shown in Fig. 7.4. The single-detector  $\mathcal{F}$ -statistic is more prominent in the H1 detector than in the L1 detector, and accumulates rapidly at the beginning of the run. As in the previous cases, this behavior is consistent with that of an instrumental artifact.

This concludes the analysis of surviving candidates of the SkyHough pipeline. Every single one of them could be related to an instrumental feature.

#### 7.4.4 Sensitivity estimation

We estimate the search sensitivity following the same procedure as previous searches [10, 11, 13, 19, 20]. Search sensitivity is quantified using the *sensitivity depth* [46, 47]

$$\mathcal{D} = \frac{\sqrt{S_n}}{h_0}, \quad (7.12)$$

where  $S_n$  represents the power spectral density (PSD) of the data, computed as the inverse squared average of the individual SFT's running-median PSD [34, 20]

$$S_n(f) = \sqrt{\frac{N_\alpha}{\sum_\alpha [S_\alpha(f)]^{-2}}}. \quad (7.13)$$

where  $S_\alpha$  represents the running-median noise floor estimation using 101 bins from the SFT labeled by starting time  $t_\alpha$  (including SFTs from both the H1 and L1 detectors) and  $N_\alpha$  represents the total number of SFTs. The resulting amplitude spectral density (ASD)  $\sqrt{S_n}$  is shown in Fig. 7.5.

The sensitivity depth  $\mathcal{D}^{95\%}$  corresponding to a 95% average detection rate is characterized by adding a campaign of software-simulated signals into the data. Simulated signals are added into 150 representative frequency bands at several sensitivity depth values bracketing the  $\mathcal{D}^{95\%}$  value in each band, as represented in Fig. 7.6. For each sensitivity depth, 200 simulated signals drawn from uniform distribution in phase and amplitude parameters are added into the data. The SkyHough pipeline is run on each of these signals in order to evaluate how many of them are detected, and the resulting toplists are clustered using the same configuration as in the main stage of the search.

For each simulated signal, we retrieve the best forty resulting clusters. The following two criteria must be fulfilled in order to label a simulated signal as “detected”. First, the loudest significance of at least one of the selected clusters must be higher than the minimum significance recovered by the corresponding all-sky clustering; this ensures the signal is significant enough to be selected for a follow-up stage. Second, the parameters of the loudest candidate in said clusters must be closer than two




FIGURE 7.5: ASD employed by the `SkyHough` pipeline to estimate the sensitivity of the search. ASD is computed as the square root of the single-sided inverse-square averaged PSD using data from both the H1 and L1 advanced LIGO detectors, as explained in the text surrounding Eq. (7.13).

parameter-space bins (see Eq. (7.5) and Table 7.2) from the simulated-signal’s parameter, as otherwise the follow-up would have missed the signal.

The efficiency associated to each sensitivity depth  $E$  is computed as the fraction of simulated signals labeled as detected. A binomial uncertainty  $\delta E$  is associated to each efficiency

$$\delta E = \sqrt{\frac{E \cdot (1 - E)}{N_I}}, \quad (7.14)$$

where  $N_I = 200$  represents the number of signals. Then, we use `scipy`’s `curve_fit` function [48] to fit a sigmoid curve to the data given by

$$S(\mathcal{D}; a, b) = 1 - \frac{1}{1 + \exp(-a\mathcal{D} + b)} \quad (7.15)$$

where  $a, b$  represent the parameters to adjust. After fitting, this expression can be numerically inverted to obtain  $\mathcal{D}^{95\%}$ . The uncertainty associated to the fit is compute through the covariance matrix  $C$  as<sup>2</sup>

$$\delta \mathcal{D}^{95\%} = \sqrt{\left(\frac{\partial S}{\partial a}\right)^2 C_{aa} + 2 \left(\frac{\partial S}{\partial a}\right) \left(\frac{\partial S}{\partial b}\right) C_{ab} + \left(\frac{\partial S}{\partial b}\right)^2 C_{bb}} \quad (7.16)$$

This procedure is exemplified in Fig. 7.6.

We compute the average wide-band  $\mathcal{D}^{95\%}(f)$  value using Gaussian process regression, as shown in Fig. 7.7. We fit a Gaussian process using to the ensemble of  $\mathcal{D}^{95\%}$  obtained from the injection campaign using `scikit-learn`’s `GaussianProcessRegressor` with an RBF kernel [49]. The uncertainty associated to the fit is computed as the 98% credible region of the deviations with respect to the Gaussian process regression, which corresponds to a 3% relative uncertainty. Equation (7.12) allows us to translate  $\mathcal{D}^{95\%}(f)$  into a corresponding CW amplitude  $h_0^{95\%}(f)$ , shown in Fig. 7.8.

<sup>2</sup>This method is akin to that employed by the `SkyHough` search in [13]. We note that Eq. (19) in [20] is incorrect and should be equivalent to Eq. (7.16) in this document. This is just a typographical error, as the analysis was performed using the correct formulae.




FIGURE 7.6: Example computation of  $\mathcal{D}^{95\%}$  (white star) at a frequency band by fitting a sigmoid function (blue solid line) to a set of efficiencies (blue dots) computed using 200 injections at each sensitivity depth for the **SkyHough** search. Shaded regions represent 1, 2, and 3 sigma envelopes of the sigmoid fit. Error bars are computed as discussed in the main text.




FIGURE 7.7: Wide-band interpolation  $\mathcal{D}^{95\%}(f)$  of the results obtained by the **SkyHough** pipeline. Each dot represents a  $\mathcal{D}^{95\%}$  at a particular frequency band computed using the procedure exemplified in Fig. 7.6. The red solid line represents a non-parametric interpolation using a Gaussian process regression, as discussed in the main text. The shaded region represents a 3% relative error with respect to the interpolation and corresponds to the 98% credible interval.




FIGURE 7.8: CW amplitude  $h_0^{95\%}$  corresponding to the 95% detection efficiency depth along the frequency band analyzed by the `SkyHough` pipeline. Solid line represents the implied  $h_0^{95\%}$  from the wide-band  $\mathcal{D}^{95\%}$  interpolation shown in Fig. 7.7. Shaded region corresponds to the 3% relative error with respect to the interpolation.




FIGURE 7.9: Comparison of broadband search sensitivities obtained by the `FrequencyHough` pipeline (black triangles), the `SkyHough` pipeline (red squares), the `Time-domain  $\mathcal{F}$ -statistic` pipeline (blue circles), and the `SOAP` pipeline (magenta diamonds). Vertical bars mark errors of  $h_0$  obtained in the procedures used by the different search pipelines. Population-averaged upper limits obtained in [50] using the O3a data are marked with dark-green crosses.




FIGURE 7.10: Left panel: detectable ellipticity, given by Eq. (7.17), as a function of the GW frequency for neutron stars with the ‘canonical’ moment of inertia  $I_{zz} = 10^{38} \text{ kg m}^2$  at a distance of 10 kpc, 1 kpc, 100 pc, and 10 pc (from top to bottom). Results for the **FrequencyHough** pipeline are marked in black, **SkyHough** in red and for **Time-domain  $\mathcal{F}$ -statistic** in blue. The right panel shows the relation between the absolute value of the first GW frequency derivative  $\dot{f} = 2\dot{f}_{rot}$  and the GW frequency  $f = 2f_{rot}$  (with  $f_{rot}$  the rotational frequency) of detectable sources as a function of the distance, assuming their spin-down is due solely to the emission of GWs. Constant spin-down ellipticities  $\epsilon^{sd}$ , corresponding to this condition, are denoted by dashed green curves. The magenta horizontal line marks the maximum spin down searched.

## 7.5 Conclusions

In Fig. 7.9 we summarize 95% confidence-level upper limits on strain amplitude  $h_0$  for the **SkyHough** pipeline, compared to the rest of the pipelines used in this search. The upper limits obtained improve on those obtained using the **PowerFlux** method in early O3 LIGO data [18]. These results constitute the most sensitive all-sky search to date for continuous GWs in the range 20-2000 Hz while probing spin-down magnitudes as high as  $1 \times 10^{-8} \text{ Hz/s}$ . Only the O2 Falcon search [17, 16, 51] provides a better sensitivity in the frequency range 20-2000 Hz; however it does so with a dramatically reduced frequency derivative range. In the frequency range of [20, 500] Hz Falcon searches a  $\dot{f}$  range from  $-3 \times 10^{-13} \text{ Hz/s}$  to  $3 \times 10^{-13} \text{ Hz/s}$  and  $\dot{f}$  range up to  $[-7.5 \times 10^{-12}, 3 \times 10^{-12}] \text{ Hz/s}$  for frequencies above 500 Hz. Thus the Falcon search parameter space is smaller than ours by factor of  $\sim 1.8 \times 10^4$  below 500 Hz and factor of  $10^3$  above 500 Hz. A recent search for persistent narrowband gravitational waves using radiometer analysis of combined O1, O2, and O3 LIGO and Virgo data in the frequency range of 20 - 1726 Hz [52] has not revealed any significant signals and has reported upper limits on an *equivalent* strain amplitude in the range of  $(0.030 - 9.6) \times 10^{-24}$ . As briefly discussed in [52], the radiometer search is expected to be significantly less sensitive than the CW searches here discussed for two reasons. First, the former uses frequency bins much larger than the latter (1/32 Hz vs O(mHz)), thus collecting more noise in each bin. Second, it does not take into account the Doppler effect due to the Earth motion, which causes a spread of the signal power over several bins (especially at higher frequencies), thus producing a further sensitivity loss.

We can use the amplitude  $h_0$  given by Eq. (7.2) to calculate star’s ellipticity  $\epsilon$ ,

$$\epsilon = \frac{c^4}{4\pi^2 G} \frac{h_0 d}{I_{zz} f^2} \approx 9.46 \times 10^{-6} \left( \frac{h_0}{10^{-25}} \right) \times \left( \frac{10^{38} \text{ kg m}^2}{I_{zz}} \right) \left( \frac{100 \text{ Hz}}{f} \right)^2 \left( \frac{d}{1 \text{ kpc}} \right). \quad (7.17)$$

Using the above equation the upper limits on the GW strain amplitude  $h_0$  can be converted to upper limits on the ellipticity  $\epsilon$ . The results are plotted in Fig. 7.10 (left panel) for four representative values of the distance  $d$  and they provide astrophysically interesting results. The NSs with ellipticities above a given trace and distance value corresponding to the trace in the left panel of Fig. 7.10 would be detectable by our searches. For instance, at frequency 200 Hz we would be able to detect a CW signal from a NS within a distance of 100 pc if its ellipticity were at least  $3 \times 10^{-7}$ . These levels of ellipticity are below the maximum value of the ellipticity that may be supported by the crust of a NS described by a standard equation of state reported in [53, 54, 55]. However they are above the most recent estimates

in general relativity by [56, 57]. The latter do not, however, exclude larger values of ellipticity when additional physical processes, such as plastic flow in the crust, are taken into account. Our upper limits are starting to probe the range predicted for pulsars by the models of [58], which predict ellipticities up to  $\epsilon \approx 10^{-7} - 10^{-6}$  for younger stars in which the deformation is not supported by crustal rigidity, but by a non-axisymmetric magnetic field at the end of its Hall driven evolution in the crust. Note however that for known pulsars at a distance of a few kpc, such as the Crab, the signal would be at frequencies  $f \lesssim 100$  Hz, so still beyond the reach of our searches.

Another way of representing limits on ellipticity is shown in the right panel of Fig. 7.10. Assuming that the emission of gravitational radiation is the sole energy loss mechanism for a rotating NS, we obtain the so-called spin-down limit  $h_0^{\text{sd}}$  on the amplitude  $h_0$ , see Eqs. (7)–(9) of [59]:

$$h_0^{\text{sd}} = \frac{1}{d} \left( \frac{5}{2} \frac{GI_{zz}}{c^3} \frac{|\dot{f}|}{f} \right)^{1/2} \approx 2.55 \times 10^{-25} \left( \frac{1 \text{ kpc}}{d} \right) \left( \frac{I_{zz}}{10^{38} \text{ kg m}^2} \right)^{1/2} \left( \frac{100 \text{ Hz}}{f} \right)^{1/2} \left( \frac{|\dot{f}|}{10^{-11} \text{ Hz s}^{-1}} \right)^{1/2}. \quad (7.18)$$

Inverting the above equation and replacing the spin-down limit amplitude  $h_0^{\text{sd}}$  with upper limit amplitudes  $h_0^{95\%}$  we have the following relation between the frequency derivative and frequency:

$$|\dot{f}| = \frac{2c^3}{5G} \frac{(h_0^{95\%} d)^2 f}{I_{zz}} \approx 1.54 \times 10^{-10} \left( \frac{h_0^{95\%}}{10^{-24}} \right)^2 \left( \frac{10^{38} \text{ kg m}^2}{I_{zz}} \right) \left( \frac{f}{100 \text{ Hz}} \right) \left( \frac{d}{1 \text{ kpc}} \right)^2. \quad (7.19)$$

In the right panel of Fig. 7.10 we have plotted  $|\dot{f}|$  as a function of frequency  $f$  for several representative values of the distance  $d$  and for a canonical value of the moment of inertia. The NSs with  $|\dot{f}|$  above a given trace and distance value corresponding to the trace in the right panel of Fig. 7.10 would be detectable by our searches.

By equating Eq. (7.2) for the amplitude  $h_0$  and Eq. (7.18) for the spin-down limit, we obtain the following equation for  $\dot{f}$ :

$$|\dot{f}| = \frac{32\pi^4 G}{5c^5} \epsilon^2 I_{zz} f^5 \approx 1.72 \times 10^{-14} \left( \frac{\epsilon}{10^{-6}} \right)^2 \left( \frac{I_{zz}}{10^{38} \text{ kg m}^2} \right) \left( \frac{f}{100 \text{ Hz}} \right)^5. \quad (7.20)$$

The dashed lines in the right panel of Fig. 7.10 are constant ellipticity curves from Eq. (7.20) above. These lines are independent of the distance  $d$ .

# Bibliography

- [1] J. Aasi *et al.*, “Advanced LIGO,” *Class. Quant. Grav.*, vol. 32, p. 074001, Apr. 2015.
- [2] F. Acernese *et al.*, “Advanced Virgo: a second-generation interferometric gravitational wave detector,” *Class. Quant. Grav.*, vol. 32, p. 024001, Jan. 2015.
- [3] B. P. Abbott *et al.*, “Gwtc-1: A gravitational-wave transient catalog of compact binary mergers observed by ligo and virgo during the first and second observing runs,” *Phys. Rev. X*, vol. 9, p. 031040, Sep 2019.
- [4] R. Abbott *et al.*, “GWTC-2: Compact Binary Coalescences Observed by LIGO and Virgo during the First Half of the Third Observing Run,” *Phys. Rev. X*, vol. 11, p. 021053, Apr. 2021.
- [5] R. Abbott *et al.*, “Observation of Gravitational Waves from Two Neutron Star-Black Hole Coalescences,” *Astrophys. J. Lett.*, vol. 915, p. L5, July 2021.
- [6] P. D. Lasky, “Gravitational Waves from Neutron Stars: A Review,” *Publ. Astron. Soc. Austral.*, vol. 32, p. e034, Sept. 2015.
- [7] K. Glampedakis and L. Gualtieri, “Gravitational Waves from Single Neutron Stars: An Advanced Detector Era Survey,” *Astroph. Space Science Library*, vol. 457, p. 673, 2018.
- [8] M. Sieniawska and M. Bejger, “Continuous Gravitational Waves from Neutron Stars: Current Status and Prospects,” *Universe*, vol. 5, p. 217, Oct. 2019.
- [9] M. Isi, M. Pitkin, and A. J. Weinstein, “Probing dynamical gravity with the polarization of continuous gravitational waves,” *Phys. Rev. D*, vol. 96, p. 042001, Aug 2017.
- [10] B. P. Abbott *et al.*, “All-sky search for periodic gravitational waves in the O1 LIGO data,” *Phys. Rev. D*, vol. 96, p. 062002, Sept. 2017.
- [11] B. P. Abbott *et al.*, “Full band all-sky search for periodic gravitational waves in the O1 LIGO data,” *Phys. Rev. D*, vol. 97, p. 102003, May 2018.
- [12] B. P. Abbott *et al.*, “First low-frequency Einstein@Home all-sky search for continuous gravitational waves in Advanced LIGO data,” *Phys. Rev. D*, vol. 96, p. 122004, Dec. 2017.
- [13] B. P. Abbott *et al.*, “All-sky search for continuous gravitational waves from isolated neutron stars using Advanced LIGO O2 data,” *Phys. Rev. D*, vol. 100, p. 024004, July 2019.
- [14] C. Palomba, S. D’Antonio, P. Astone, S. Frasca, G. Intini, I. La Rosa, P. Leaci, S. Mastrogiovanni, A. L. Miller, F. Muciaccia, O. J. Piccinni, L. Rei, and F. Simula, “Direct constraints on the ultralight boson mass from searches of continuous gravitational waves,” *Phys. Rev. Lett.*, vol. 123, p. 171101, Oct 2019.
- [15] B. Steltner, M. A. Papa, H. B. Eggenstein, B. Allen, V. Dergachev, R. Prix, B. Machenschalk, S. Walsh, S. J. Zhu, O. Behnke, and S. Kwang, “Einstein@Home All-sky Search for Continuous Gravitational Waves in LIGO O2 Public Data,” *Astrophys. J.*, vol. 909, p. 79, Mar. 2021.
- [16] V. Dergachev and M. A. Papa, “Results from the First All-Sky Search for Continuous Gravitational Waves from Small-Ellipticity Sources,” *Phys. Rev. Lett.*, vol. 125, p. 171101, Oct. 2020.
- [17] V. Dergachev and M. A. Papa, “Results from high-frequency all-sky search for continuous gravitational waves from small-ellipticity sources,” *Phys. Rev. D*, vol. 103, p. 063019, Mar. 2021.

- [18] R. Abbott *et al.*, “All-sky search for continuous gravitational waves from isolated neutron stars in the early O3 LIGO data,” *Phys. Rev. D*, vol. 104, p. 082004, Oct 2021.
- [19] P. B. Covas and A. M. Sintes, “First all-sky search for continuous gravitational-wave signals from unknown neutron stars in binary systems using Advanced LIGO data,” *Phys. Rev. Lett.*, vol. 124, no. 19, p. 191102, 2020.
- [20] R. Abbott *et al.*, “All-sky search in early O3 LIGO data for continuous gravitational-wave signals from unknown neutron stars in binary systems,” *Phys. Rev. D*, vol. 103, no. 6, p. 064017, 2021.
- [21] K. Wette, L. Dunn, P. Clearwater, and A. Melatos, “Deep exploration for continuous gravitational waves at 171-172 Hz in LIGO second observing run data,” *Phys. Rev. D*, vol. 103, p. 083020, Apr. 2021.
- [22] R. Abbott *et al.*, “All-sky search for gravitational wave emission from scalar boson clouds around spinning black holes in LIGO O3 data.” arXiv:2111.15507 [astro-ph.HE].
- [23] R. Tenorio, D. Keitel, and A. M. Sintes, “Search methods for continuous gravitational-wave signals from unknown sources in the advanced-detector era,” *Universe*, vol. 7, no. 12, 2021.
- [24] P. Astone, A. Colla, S. D’Antonio, S. Frasca, and C. Palomba, “Method for all-sky searches of continuous gravitational wave signals using the frequency-Hough transform,” *Phys. Rev. D*, vol. 90, p. 042002, Aug. 2014.
- [25] B. Krishnan, A. M. Sintes, M. A. Papa, B. F. Schutz, S. Frasca, and C. Palomba, “Hough transform search for continuous gravitational waves,” *Phys. Rev. D*, vol. 70, p. 082001, Oct. 2004.
- [26] P. Jaranowski, A. Królak, and B. F. Schutz, “Data analysis of gravitational-wave signals from spinning neutron stars: The signal and its detection,” *Phys. Rev. D*, vol. 58, p. 063001, Sept. 1998.
- [27] J. Aasi and et al., “Implementation of an f-statistic all-sky search for continuous gravitational waves in virgo VSR1 data,” *Class. Quant. Grav.*, vol. 31, p. 165014, aug 2014.
- [28] J. Bayley, C. Messenger, and G. Woan, “Generalized application of the Viterbi algorithm to searches for continuous gravitational-wave signals,” *Phys. Rev. D*, vol. 100, p. 023006, July 2019.
- [29] L. Sun, E. Goetz, J. S. Kissel, J. Betzwieser, S. Karki, A. Viets, M. Wade, D. Bhattacharjee, V. Bossilkov, P. B. Covas, L. E. H. Datrier, R. Gray, S. Kandhasamy, Y. K. Lecoche, G. Mendell, T. Mistry, E. Payne, R. L. Savage, A. J. Weinstein, S. Aston, A. Buikema, C. Cahillane, J. C. Driggers, S. E. Dwyer, R. Kumar, and A. Urban, “Characterization of systematic error in Advanced LIGO calibration,” *Class. Quant. Grav.*, vol. 37, p. 225008, Nov. 2020.
- [30] L. Sun, E. Goetz, J. S. Kissel, J. Betzwieser, S. Karki, D. Bhattacharjee, P. B. Covas, L. E. H. Datrier, S. Kandhasamy, Y. K. Lecoche, G. Mendell, T. Mistry, E. Payne, R. L. Savage, A. Viets, M. Wade, A. J. Weinstein, S. Aston, C. Cahillane, J. C. Driggers, S. E. Dwyer, and A. Urban, “Characterization of systematic error in Advanced LIGO calibration in the second half of O3.” arXiv:2107.00129 [astro-ph.IM].
- [31] F. Acernese *et al.*, “Calibration of Advanced Virgo and reconstruction of detector strain  $h(t)$  during the Observing Run O3.” arXiv:2107.03294 [gr-qc].
- [32] E. Goetz *et al.*, “O3 lines and combs in found in self-gated c01 data,” Tech. Rep. T2100200-v2, LIGO, 2021.
- [33] A. M. Sintes and B. Krishnan, “Improved hough search for gravitational wave pulsars,” *J. Phys. Conf. Ser.*, vol. 32, pp. 206–211, 2006.
- [34] B. Abbott *et al.*, “All-sky search for periodic gravitational waves in LIGO S4 data,” *Phys. Rev. D*, vol. 77, p. 022001, Jan. 2008.
- [35] B. P. Abbott *et al.*, “Comprehensive all-sky search for periodic gravitational waves in the sixth science run LIGO data,” *Phys. Rev. D*, vol. 94, p. 042002, Aug. 2016.



- [36] P. B. Covas and A. M. Sintes, “New method to search for continuous gravitational waves from unknown neutron stars in binary systems,” *Phys. Rev. D*, vol. 99, no. 12, p. 124019, 2019.
- [37] M. Oliver, D. Keitel, and A. M. Sintes, “Adaptive transient Hough method for long-duration gravitational wave transients,” *Phys. Rev. D*, vol. 99, no. 10, p. 104067, 2019.
- [38] R. Tenorio, D. Keitel, and A. M. Sintes, “Time-frequency track distance for comparing continuous gravitational wave signals,” *Phys. Rev. D*, vol. 103, no. 6, p. 064053, 2021.
- [39] D. Keitel, R. Tenorio, G. Ashton, and R. Prix, “PyFstat: a Python package for continuous gravitational-wave data analysis,” *The Journal of Open Source Software*, vol. 6, p. 3000, Apr. 2021.
- [40] R. Tenorio, D. Keitel, and A. M. Sintes, “Application of a hierarchical MCMC follow-up to Advanced LIGO continuous gravitational-wave candidates,” *Phys. Rev. D*, vol. 104, no. 8, p. 084012, 2021.
- [41] S. J. Zhu, M. Baryakhtar, M. A. Papa, D. Tsuna, N. Kawanaka, and H.-B. Eggenstein, “Characterizing the continuous gravitational-wave signal from boson clouds around Galactic isolated black holes,” *Phys. Rev. D*, vol. 102, no. 6, p. 063020, 2020.
- [42] G. Ashton and R. Prix, “Hierarchical multistage MCMC follow-up of continuous gravitational wave candidates,” *Phys. Rev. D*, vol. 97, p. 103020, May 2018.
- [43] R. Abbott *et al.*, “All-sky search for continuous gravitational waves from isolated neutron stars in the early O3 LIGO data,” *Phys. Rev. D*, vol. 104, no. 8, p. 082004, 2021.
- [44] R. Abbott *et al.*, “Search for continuous gravitational waves from 20 accreting millisecond X-ray pulsars in O3 LIGO data,” *Phys. Rev. D*, vol. 105, p. 022002, 1 2022.
- [45] Prix, Reinhard, “Coherent  $\mathcal{F}$ -statistic on semi-coherent candidate.” <https://dcc.ligo.org/LIGO-T1700236/public>, 2019.
- [46] B. Behnke, M. A. Papa, and R. Prix, “Postprocessing methods used in the search for continuous gravitational- wave signals from the Galactic Center,” *Phys. Rev. D*, vol. 91, p. 064007, Mar. 2015.
- [47] C. Dreissigacker, R. Prix, and K. Wette, “Fast and Accurate Sensitivity Estimation for Continuous-Gravitational-Wave Searches,” *Phys. Rev. D*, vol. 98, no. 8, p. 084058, 2018.
- [48] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, Í. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors, “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python,” *Nature Methods*, vol. 17, pp. 261–272, 2020.
- [49] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, “Scikit-learn: Machine learning in Python,” *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [50] R. Abbott *et al.*, “All-sky Search for Continuous Gravitational Waves from Isolated Neutron Stars in the Early O3 LIGO Data.”
- [51] V. Dergachev and M. Alessandra Papa, “The search for continuous gravitational waves from small-ellipticity sources at low frequencies.” arXiv:2104.09007 [gr-qc].
- [52] R. Abbott *et al.*, “All-sky, all-frequency directional search for persistent gravitational-waves from advanced ligo’s and advanced virgo’s first three observing runs.” arXiv:2110.09834 [gr-qc].
- [53] G. Ushomirsky, C. Cutler, and L. Bildsten, “Deformations of accreting neutron star crusts and gravitational wave emission,” *Month. Not. Roy. Astron. Soc.*, vol. 319, pp. 902–932, 12 2000.

- [54] B. Haskell, D. I. Jones, and N. Andersson, “Mountains on neutron stars: accreted versus non-accreted crusts,” *Month. Not. Roy. Astron. Soc.*, vol. 373, pp. 1423–1439, 11 2006.
- [55] N. K. Johnson-McDaniel and B. J. Owen, “Maximum elastic deformations of relativistic stars,” *Phys. Rev. D*, vol. 88, p. 044004, Aug. 2013.
- [56] F. Gittins and N. Andersson, “Modelling neutron star mountains in relativity,” *Month. Not. Roy. Astron. Soc.*, vol. 507, pp. 116–128, 07 2021.
- [57] F. Gittins, N. Andersson, and D. I. Jones, “Modelling neutron star mountains,” *Month. Not. Roy. Astron. Soc.*, vol. 500, pp. 5570–5582, 11 2021.
- [58] A. G. Suvorov, A. Mastrano, and U. Geppert, “Gravitational radiation from neutron stars deformed by crustal Hall drift,” *Mon. Notices Royal Astron. Soc.*, vol. 459, pp. 3407–3418, July 2016.
- [59] B. P. Abbott *et al.*, “Searches for Gravitational Waves from Known Pulsars at Two Harmonics in 2015-2017 LIGO Data,” *Astrophys. J.*, vol. 879, p. 10, July 2019.

## Chapter 8

# Conclusion and future developments

This thesis presents three new developments to post-process and follow-up the results of a CW search and two searches for CWs from unknown sources in O3 Advanced LIGO data. Taken together, these results demonstrate a new approach for conducting broad parameter-space searches in which candidates are systematically followed-up using increasingly sensitive methods, rather than simply being reported as “interesting outliers” whenever they cannot be ascribed to a known instrumental artifact.

The impact of the new developments in this thesis is manifested by the extensive use of these methods in a significant fraction of the CW searches produced in O3 Advanced LIGO and Advanced Virgo data. Clusterings based on the distance proposed in Chapter 3 have been used in two all-sky searches [1, 2, 3]. The follow-up strategy proposed in Chapter 4 has been used in three all-sky searches [1, 2, 4, 3], as well as a narrow-band search [5] and three directed searches [6, 7, 8]. `distromax`, presented in Chapter 5, has been applied in a narrow-band [9] and directed [10] search. Also, it has been employed to estimate detection thresholds in theoretical sensitivity estimates for the detection of post-glitch emission from known neutron stars [11].

At the time of writing, the Advanced LIGO, Advanced Virgo, and KAGRA detectors are about to start the fourth observing run, O4, which is expected to last for approximately two years, achieving an unprecedented sensitivity that may start to unveil the first signatures of a CW signal [12]. Research in the search for CW signals, however, is far from being complete, as there are still a few open problems to be solved in order to ensure a first detection is achieved and the maximum information is obtained out of it. We shall devote the remainder of this chapter to briefly discuss three of the possible avenues where a significant improvement may be achieved in the forthcoming years.

First, in the most general case, a CW signal is defined essentially by its presence in the detector data for the full duration of an observing run. Even though NSs are considered as the primary source of this sort of signals, other physical systems, such as boson clouds around spinning black holes [13] may produce genuine CW signals. These signals can be searched for either using all-sky searches [14] or targeting specific known black holes [15]; in certain cases, these signals may be even produced by newly formed black holes after the coalescence of a compact binary system [16]. Extending this definition, the search for CWs has proven to be suitable to search for the presence of a galactic dark-matter halo [17, 18, 19, 20], as such a system would produce a monochromatic signal in the detector due to the effect of such a halo on the properties of the detector’s mirrors. Finally, binary systems of compact objects may also be interesting sources of CWs in the appropriate parameter space. A first example are binary systems with low-mass components, such as those composed by planetary-mass primordial black holes [21, 22] or “mini extreme-mass-ratio inspirals” [23], which may allow us to detect exotic compact objects in our galaxy. As we progress into space-borne detectors, such as LISA [24], binary systems’ gravitational-wave emission tends to behave like a CW signal [25], lasting for long periods of time. This is true, for example, for the case of binary-white-dwarf systems, to which CW search techniques have already been applied [26, 27].

The second avenue for improvement is, paradoxically, to increase the amount of prior information used in a blind search. Initial steps in this direction were taken by so-called “spotlight searches” [28, 29], which applied the same modus operandi of an all-sky search to a specific region of the sky. This reduces the number of templates to be analyzed by a search and allows for the use of more sensitive methods as a first stage, increasing the resulting sensitivity. In parallel, a rich body of literature on simulating the galactic population of NSs and their detectability as CW sources has been developed during the past two decades (see [12] and references therein). The analyses therein presented tend to consider electromagnetic emission to dominate the spin frequency evolution of a NS, similarly to what happens in pulsars; there is no strong evidence, however, against the existence of a NS population whose main emission channel

is given by gravitational radiation. Indeed, the properties and detectability of a theoretical galactic population of *gravitars* are studied in [30]. Moreover, as pointed out by [31] the conditions for the existence of a gravitar population are astrophysically plausible. This information, combined with some basic estimations of the sensitivity of CW searches [32], can reduce the standard parameter-space of an all-sky search [3] by an order of magnitude. Furthermore, another reduction of about an order of magnitude in computing cost can be obtained by focusing on the *nearby* population of gravitars [33]. Nevertheless, the dominating factor will be the specific distance to nearby sources: for ellipticities of  $10^{-7}$ , gravitars within 100 pc may be discovered using the current advanced detectors; beyond that, and up to 1 kpc at 1 kHz, gravitars fall into the third generation's territory [34].

Finally, once a CW detection is achieved, the problem of parameter estimation will become a real one. Most of the approaches for the detection of CW signals, as discussed in Chapter 2, focus on estimating the frequency and Doppler modulation parameters, marginalizing out the amplitude parameters. As discussed in [35, 36, 37], the amount of information we can extract from such a result is highly dependent on the availability of a distance measurement to the source. This can be easily achieved for galactic sources with an electromagnetic counterpart, and for nearby sources using gravitational-wave parallax [38].

This brief exposition suggests that, for a first CW signal to be scientifically exploited at its fullest, it would be desirable to come from a nearby source so that there is a smaller dependency on complementary electromagnetic emission. As we progress into a state of routinary detection and follow-up of new CW sources, population studies such as the ones nowadays performed in compact binary coalescences [39] will likely relax such a stringent constraint.

# Bibliography

- [1] R. Abbott *et al.*, “All-sky search in early O3 LIGO data for continuous gravitational-wave signals from unknown neutron stars in binary systems,” *Phys. Rev. D*, vol. 103, no. 6, p. 064017, 2021.
- [2] R. Tenorio, “An all-sky search in early O3 LIGO data for continuous gravitational-wave signals from unknown neutron stars in binary systems,” in *55th Rencontres de Moriond on Gravitation*, 5 2021.
- [3] R. Abbott *et al.*, “All-sky search for continuous gravitational waves from isolated neutron stars using Advanced LIGO and Advanced Virgo O3 data,” *Phys. Rev. D*, vol. 106, no. 10, p. 102008, 2022.
- [4] R. Abbott *et al.*, “All-sky search for continuous gravitational waves from isolated neutron stars in the early O3 LIGO data,” *Phys. Rev. D*, vol. 104, no. 8, p. 082004, 2021.
- [5] R. Abbott *et al.*, “Search for continuous gravitational waves from 20 accreting millisecond x-ray pulsars in O3 LIGO data,” *Phys. Rev. D*, vol. 105, p. 022002, 2022.
- [6] R. Abbott *et al.*, “Search for continuous gravitational wave emission from the Milky Way center in O3 LIGO-Virgo data,” *Phys. Rev. D*, vol. 106, no. 4, p. 042003, 2022.
- [7] R. Abbott *et al.*, “Model-based Cross-correlation Search for Gravitational Waves from the Low-mass X-Ray Binary Scorpius X-1 in LIGO O3 Data,” *Astrophys. J. Lett.*, vol. 941, no. 2, p. L30, 2022.
- [8] J. T. Whelan *et al.*, “Search for Gravitational Waves from Scorpius X-1 in LIGO O3 Data With Corrected Orbital Ephemeris,” *Astrophys. J. Lett.*, vol. 949, no. 117, 2023.
- [9] R. Abbott *et al.*, “Narrowband Searches for Continuous and Long-duration Transient Gravitational Waves from Known Pulsars in the LIGO-Virgo Third Observing Run,” *Astrophys. J.*, vol. 932, no. 2, p. 133, 2022.
- [10] Y. Liu and Y.-C. Zou, “Directed search for continuous gravitational waves from the possible kilonova remnant G4.8+6.2,” *Phys. Rev. D*, vol. 106, no. 12, p. 123024, 2022.
- [11] J. Moragues, L. M. Modafferi, R. Tenorio, and D. Keitel, “Prospects for detecting transient quasi-monochromatic gravitational waves from glitching pulsars with current and future detectors,” *Mon. Not. Roy. Astron. Soc.*, vol. 519, no. 4, pp. 5161–5176, 2023.
- [12] G. Pagliaro, M. A. Papa, J. Ming, J. Lian, D. Tsuna, C. Maraston, and D. Thomas, “Continuous gravitational waves from Galactic neutron stars: demography, detectability and prospects.” arXiv:2303.04714 [gr-qc].
- [13] R. Brito, V. Cardoso, and P. Pani, “Superradiance: New Frontiers in Black Hole Physics,” *Lect. Notes Phys.*, vol. 906, pp. pp.1–237, 2015.
- [14] R. Abbott *et al.*, “All-sky search for gravitational wave emission from scalar boson clouds around spinning black holes in LIGO O3 data,” *Phys. Rev. D*, vol. 105, no. 10, p. 102001, 2022.
- [15] M. Isi, L. Sun, R. Brito, and A. Melatos, “Directed searches for gravitational waves from ultralight bosons,” *Phys. Rev. D*, vol. 99, no. 8, p. 084042, 2019. [Erratum: Phys.Rev.D 102, 049901 (2020)].
- [16] D. Jones, L. Sun, N. Siemonsen, W. E. East, S. M. Scott, and K. Wette, “Methods and prospects for gravitational wave searches targeting ultralight vector boson clouds around known black holes.” arXiv:2305.00401 [gr-qc].

- [17] R. Abbott *et al.*, “Constraints on dark photon dark matter using data from LIGO’s and Virgo’s third observing run,” *Phys. Rev. D*, vol. 105, no. 6, p. 063030, 2022.
- [18] K. Nagano, H. Nakatsuka, S. Morisaki, T. Fujita, Y. Michimura, and I. Obata, “Axion dark matter search using arm cavity transmitted beams of gravitational wave detectors,” *Phys. Rev. D*, vol. 104, no. 6, p. 062008, 2021.
- [19] S. M. Vermeulen *et al.*, “Direct limits for scalar field dark matter from a gravitational-wave detector.” arXiv:2103.03783 [gr-qc].
- [20] A. L. Miller, F. Badaracco, and C. Palomba, “Distinguishing between dark-matter interactions with gravitational-wave detectors,” *Phys. Rev. D*, vol. 105, no. 10, p. 103035, 2022.
- [21] A. L. Miller, S. Clesse, F. De Lillo, G. Bruno, A. Depasse, and A. Tanasijczuk, “Probing planetary-mass primordial black holes with continuous gravitational waves,” *Phys. Dark Univ.*, vol. 32, p. 100836, 2021.
- [22] A. L. Miller, N. Aggarwal, S. Clesse, and F. De Lillo, “Constraints on planetary and asteroid-mass primordial black holes from continuous gravitational-wave searches,” *Phys. Rev. D*, vol. 105, no. 6, p. 062008, 2022.
- [23] H.-K. Guo and A. Miller, “Searching for Mini Extreme Mass Ratio Inspirals with Gravitational-Wave Detectors.” arXiv:2205.10359 [astro-ph.IM].
- [24] P. Amaro-Seoane *et al.*, “Laser interferometer space antenna.” arXiv:1702.00786 [astro-ph.IM].
- [25] P. Amaro-Seoane *et al.*, “Astrophysics with the Laser Interferometer Space Antenna,” *Living Rev. Rel.*, vol. 26, no. 1, p. 2, 2023.
- [26] R. Prix and J. T. Whelan, “F-statistic search for white-dwarf binaries in the first Mock LISA Data Challenge,” *Class. Quant. Grav.*, vol. 24, no. 19, pp. S565–S574, 2007.
- [27] J. T. Whelan, R. Prix, and D. Khurana, “Improved search for galactic white-dwarf binaries in Mock LISA Data Challenge 1B using an F-statistic template bank,” *Class. Quant. Grav.*, vol. 25, p. 184029, 2008.
- [28] J. Aasi *et al.*, “Search of the Orion spur for continuous gravitational waves using a loosely coherent algorithm on data from LIGO interferometers,” *Phys. Rev. D*, vol. 93, no. 4, p. 042006, 2016.
- [29] V. Dergachev, M. A. Papa, B. Steltner, and H.-B. Eggenstein, “Loosely coherent search in LIGO O1 data for continuous gravitational waves from Terzan 5 and the galactic center,” *Phys. Rev. D*, vol. 99, no. 8, p. 084048, 2019.
- [30] C. Palomba, “Simulation of a population of isolated neutron stars evolving through the emission of gravitational waves,” *Mon. Not. Roy. Astron. Soc.*, vol. 359, pp. 1150–1164, 2005.
- [31] B. Knispel and B. Allen, “Blandford’s Argument: The Strongest Continuous Gravitational Wave Signal,” *Phys. Rev. D*, vol. 78, p. 044031, 2008.
- [32] C. Dreissigacker, R. Prix, and K. Wette, “Fast and Accurate Sensitivity Estimation for Continuous-Gravitational-Wave Searches,” *Phys. Rev. D*, vol. 98, no. 8, p. 084058, 2018.
- [33] J. G. Schmidt, M. M. Hohle, and R. Neuhäuser, “Determination of a temporally and spatially resolved supernova rate from ob stars within 5 kpc,” *Astron. Nach.*, vol. 335, no. 9, pp. 935–948, 2014.
- [34] R. Tenorio *et al.*, “in prep.,” 2023.
- [35] M. Sieniawska and D. I. Jones, “Gravitational waves from spinning neutron stars as not-quite-standard sirens,” *Mon. Not. Roy. Astron. Soc.*, vol. 509, no. 4, pp. 5179–5187, 2021.
- [36] D. I. Jones, “Learning from the Frequency Content of Continuous Gravitational Wave Signals.” arXiv:2111.08561 [astro-ph.HE].

- [37] N. Lu, K. Wette, S. M. Scott, and A. Melatos, “Inferring neutron star properties with continuous gravitational waves,” 9 2022.
- [38] M. Sieniawska, D. I. Jones, and A. L. Miller, “Measuring neutron star distances and properties with gravitational-wave parallax,” *Monthly Notices of the Royal Astronomical Society*, vol. 521, pp. 1924–1930, 02 2023.
- [39] R. Abbott *et al.*, “GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo During the Second Part of the Third Observing Run,” 11 2021.





## Appendix A

# Coauthors' authorization



Dra. Alicia M. Sintes Olives, com a coautora dels següents articles

*Time-frequency track distance for comparing continuous gravitational wave signals*

Rodrigo Tenorio, David Keitel, Alicia M. Sintes

Phys. Rev. D 103, 064053 (2021) DOI: 10.1103/PhysRevD.103.064053

*Application of a hierarchical MCMC follow-up to Advanced LIGO continuous gravitational-wave candidates*

Rodrigo Tenorio, David Keitel, Alicia M. Sintes

Phys. Rev. D 104, 084012 (2021) DOI: 10.1103/PhysRevD.104.084012

*Search methods for continuous gravitational-wave signals from unknown sources in the advanced-detector era*

Rodrigo Tenorio, David Keitel, Alicia M. Sintes

Universe 2021, 7(12), 474 DOI: 10.3390/universe7120474

*Empirically estimating the distribution of the loudest candidate from a gravitational-wave search*

Rodrigo Tenorio, Luana M. Modafferi, David Keitel, Alicia M. Sintes

Phys. Rev. D 105, 044029 (2022) DOI: 10.1103/PhysRevD.105.044029

DECLARO:

Que accepto que el Sr. Rodrigo Tenorio Márquez, presenti els articles anteriorment citats com a autor principal i formin part de la seva tesi doctoral. Aquests articles no podran formar part, en cap cas, d'una altra tesi doctoral.

I perquè quedi constància d'això signo aquest document.

Signature,

Palma, 10 de gener, 2023



**Universitat**  
de les Illes Balears

Dr. David Keitel, com a coautor dels següents articles

*Time-frequency track distance for comparing continuous gravitational wave signals*

Rodrigo Tenorio, David Keitel, Alicia M. Sintes

Phys. Rev. D 103, 064053 (2021) DOI: 10.1103/PhysRevD.103.064053

*Application of a hierarchical MCMC follow-up to Advanced LIGO continuous gravitational-wave candidates*

Rodrigo Tenorio, David Keitel, Alicia M. Sintes

Phys. Rev. D 104, 084012 (2021) DOI: 10.1103/PhysRevD.104.084012

*Search methods for continuous gravitational-wave signals from unknown sources in the advanced-detector era*

Rodrigo Tenorio, David Keitel, Alicia M. Sintes

Universe 2021, 7(12), 474 DOI: 10.3390/universe7120474

*Empirically estimating the distribution of the loudest candidate from a gravitational-wave search*

Rodrigo Tenorio, Luana M. Modafferi, David Keitel, Alicia M. Sintes

Phys. Rev. D 105, 044029 (2022) DOI: 10.1103/PhysRevD.105.044029

DECLARO:

Que accepto que el Sr. Rodrigo Tenorio Márquez, presenti els articles anteriorment citats com a autor principal i formin part de la seva tesi doctoral. Aquests articles no podran formar part, en cap cas, d'una altra tesi doctoral.

I perquè quedi constància d'això signo aquest document.

Signature

David Keitel, Palma, 2023/01/09



**Universitat**  
de les Illes Balears

Sra. Luana Michela Modafferi, com a coautora del següent article

*Empirically estimating the distribution of the loudest candidate from a gravitational-wave search*

Rodrigo Tenorio, Luana M. Modafferi, David Keitel, Alicia M. Sintes  
Phys. Rev. D 105, 044029 (2022) DOI: 10.1103/PhysRevD.105.044029

DECLARO:

Que accepto que el Sr. Rodrigo Tenorio Márquez, presenti els articles anteriorment citats com a autor principal i formin part de la seva tesi doctoral. Aquests articles no podran formar part, en cap cas, d'una altra tesi doctoral.

I perquè quedi constància d'això signo aquest document.

Signature

*Luana Michela Modafferi*  
Palma, 9th January 2023