

PyFstat: a Python package for continuous gravitational-wave data analysis

David Keitel¹, Rodrigo Tenorio¹, Gregory Ashton², and Reinhard Prix^{3, 4}

¹ Departament de Física, Institut d'Aplicacions Computacionals i de Codi Comunitari (IAC3), Universitat de les Illes Balears, and Institut d'Estudis Espacials de Catalunya (IEEC), Crta. Valldemossa km 7.5, E-07122 Palma, Spain ² OzGrav, School of Physics & Astronomy, Monash University, Clayton 3800, Victoria, Australia ³ Max-Planck-Institut für Gravitationsphysik (Albert-Einstein-Institut), D-30167 Hannover, Germany ⁴ Leibniz Universität Hannover, D-30167 Hannover, Germany

DOI: [10.21105/joss.03000](https://doi.org/10.21105/joss.03000)

Software

- [Review](#) ↗
- [Repository](#) ↗
- [Archive](#) ↗

Editor: [Daniel S. Katz](#) ↗

Reviewers:

- [@khanx169](#)
- [@RobertRosca](#)

Submitted: 26 January 2021

Published: 06 April 2021

License

Authors of papers retain copyright and release the work under a Creative Commons Attribution 4.0 International License ([CC BY 4.0](#)).

Summary

Gravitational waves in the sensitivity band of ground-based detectors can be emitted by a number of astrophysical sources, including not only binary coalescences, but also individual spinning neutron stars. The most promising signals from such sources, although as of 2020 not yet detected, are the long-lasting, quasi-monochromatic 'Continuous Waves' (CWs). Many search methods have been developed and applied on LIGO ([Aasi et al., 2015](#)) and Virgo ([Acernese et al., 2015](#)) data. See [Prix \(2009\)](#), [Riles \(2017\)](#), and [Sieniawska & Bejger \(2019\)](#) for reviews of the field.

The PyFstat package provides tools to perform a range of CW data analysis tasks. It revolves around the \mathcal{F} -statistic, first introduced by [Jaranowski et al. \(1998\)](#): a matched-filter detection statistic for CW signals described by a set of frequency evolution parameters and maximized over amplitude parameters. This has been one of the standard methods for LIGO-Virgo CW searches for two decades. PyFstat is built on top of established routines in LALSuite ([LIGO Scientific Collaboration, 2018](#)) but through its more modern Python interface it enables a flexible approach to designing new search strategies.

Classes for various search strategies and target signals are contained in three main submodules:

- `core`: The basic wrappers to LALSuite's \mathcal{F} -statistic algorithm. End-users should rarely need to access these directly.
- `grid_based_searches`: Classes to search over regular parameter-space grids.
- `mcmc_based_searches`: Classes to cover promising parameter-space regions through stochastic template placement with the Markov Chain Monte Carlo (MCMC) sampler `ptemcee` ([Vousden et al., 2015](#)).

Besides standard CWs from isolated neutron stars, PyFstat can also be used to search for CWs from sources in binary systems (including the additional orbital parameters), for CWs with a discontinuity at a pulsar glitch, and for CW-like long-duration transient signals, e.g., from *after* a pulsar glitch. Specialized versions of both grid-based and MCMC-based search classes are provided for these scenarios. Both fully-coherent and semi-coherent searches (where the data is split into several segments for efficiency) are covered, and an extension to the \mathcal{F} -statistic that is more robust against single-detector noise artifacts ([Keitel et al., 2014](#)) is also supported. While PyFstat's grid-based searches do not compete with the sophisticated grid setups and semi-coherent algorithms implemented in various LALSuite programs, its main

scientific use cases so far are for the MCMC exploration of interesting parameter-space regions and for the long-duration transient case.

PyFstat was first introduced in [Ashton & Prix \(2018\)](#), which remains the main reference for the MCMC-based analysis implemented in the package. The extension to transient signals, which uses PyCUDA ([Klöckner et al., 2012](#)) for speedup, is discussed in detail in [Keitel & Ashton \(2018\)](#), and the glitch-robust search approaches in [Ashton et al. \(2018\)](#).

Additional helper classes, utility functions, and internals are included for handling the common Short Fourier Transform (SFT) data format for LIGO data, simulating artificial data with noise and signals in them, and plotting results and diagnostics. Most of the underlying LALSuite functionality is accessed through SWIG wrappings ([Wette, 2020](#)) though for some parts, such as the SFT handling, we still (as of the writing of this paper) call stand-alone lalapps executables. Completing the backend migration to pure SWIG usage is planned for the future.

The source of PyFstat is hosted on [GitHub](#). The repository also contains an automated test suite and a set of introductory example scripts. Issues with the software can be submitted through GitHub and pull requests are always welcome. PyFstat can be installed through pip, conda or docker containers. Documentation in html and pdf formats is available from [readthedocs.org](#) and installation instructions can be found there or in the [README](#) file. PyFstat is also listed in the Astrophysics Source Code Library as [ascl:2102.027](#).

Statement of need

The sensitivity of searches for CWs and long-duration transient GWs is generally limited by computational resources, as the required number of matched-filter templates increases steeply for long observation times and wide parameter spaces. The C-based LALSuite library ([LIGO Scientific Collaboration, 2018](#)) contains many sophisticated search methods with a long development history and high level of optimization, but is not very accessible for researchers new to the field or for students; nor is it convenient for rapid development and integration with modern technologies like GPUs or machine learning. Hence, PyFstat serves a dual function of (i) making LALSuite CW functionality more easily accessible through a Python interface, thus facilitating the new user experience and, for developers, the exploratory implementation of novel methods; and (ii) providing a set of production-ready search classes for use cases not yet covered by LALSuite itself, most notably for MCMC-based followup of promising candidates from wide-parameter-space searches.

So far, PyFstat has been used for

- the original proposal of MCMC followup for CW candidates ([Ashton & Prix, 2018](#));
- developing glitch-robust CW search methods ([Ashton et al., 2018](#));
- speeding up long-transient searches with GPUs ([Keitel & Ashton, 2018](#));
- followup of candidates from all-sky searches for CWs from sources in binary systems, see [Covas & Sintes \(2020\)](#) and [Abbott et al. \(2021\)](#);
- studying the impact of neutron star proper motions on CW searches ([Covas, 2020](#)).

Acknowledgements

We acknowledge contributions to the package from Karl Wette, Sylvia Zhu and Dan Foreman-Mackey; as well as helpful suggestions by John T. Whelan, Luca Rei, and the LIGO-Virgo-KAGRA Continuous Wave working group. D.K. and R.T. are supported by European Union FEDER funds; the Spanish Ministerio de Ciencia, Innovación y Universidades and Agencia Estatal de Investigación grants PID2019-106416GB-I00/AEI/10.13039/501100011033,

RED2018-102661-T, RED2018-102573-E, FPA2017-90687-REDC FPU 18/00694, and BEA-GAL 18/00148 (cofinanced by the Universitat de les Illes Balears); the Comunitat Autònoma de les Illes Balears through the Direcció General de Política Universitària i Recerca with funds from the Tourist Stay Tax Law ITS 2017-006 (PRD2018/24) and the Conselleria de Fons Europeus, Universitat i Cultura; the Generalitat Valenciana (PROMETEO/2019/071); and EU COST Actions CA18108, CA17137, CA16214, and CA16104. This paper has been assigned document number LIGO-P2100008.

References

- Aasi, J., Abbott, B. P., Abbott, R., & others. (2015). Advanced LIGO. *Class. Quant. Grav.*, 32, 074001. <https://doi.org/10.1088/0264-9381/32/7/074001>
- Abbott, R., Abbott, T. D., Abraham, S., & others. (2021). All-sky search in early O3 LIGO data for continuous gravitational-wave signals from unknown neutron stars in binary systems. *Phys. Rev. D*, 103(6), 064017. <https://doi.org/10.1103/PhysRevD.103.064017>
- Acernese, F., Agathos, M., Agatsuma, K., & others. (2015). Advanced Virgo: a second-generation interferometric gravitational wave detector. *Class. Quant. Grav.*, 32(2), 024001. <https://doi.org/10.1088/0264-9381/32/2/024001>
- Ashton, G., & Prix, R. (2018). Hierarchical multistage MCMC follow-up of continuous gravitational wave candidates. *Phys. Rev. D*, 97(10), 103020. <https://doi.org/10.1103/PhysRevD.97.103020>
- Ashton, G., Prix, R., & Jones, D. I. (2018). A semicoherent glitch-robust continuous-gravitational-wave search method. *Phys. Rev. D*, 98(6), 063011. <https://doi.org/10.1103/PhysRevD.98.063011>
- Covas, P. B. (2020). Effects of proper motion of neutron stars on continuous gravitational-wave searches. *Mon. Not. Roy. Astron. Soc.*, 500(4), 5167–5176. <https://doi.org/10.1093/mnras/staa3624>
- Covas, P. B., & Sintes, A. M. (2020). First all-sky search for continuous gravitational-wave signals from unknown neutron stars in binary systems using Advanced LIGO data. *Phys. Rev. Lett.*, 124(19), 191102. <https://doi.org/10.1103/PhysRevLett.124.191102>
- Jaranowski, P., Krolak, A., & Schutz, B. F. (1998). Data analysis of gravitational - wave signals from spinning neutron stars. 1. The Signal and its detection. *Phys. Rev. D*, 58, 063001. <https://doi.org/10.1103/PhysRevD.58.063001>
- Keitel, D., & Ashton, G. (2018). Faster search for long gravitational-wave transients: GPU implementation of the transient \mathcal{F} -statistic. *Class. Quant. Grav.*, 35(20), 205003. <https://doi.org/10.1088/1361-6382/aade34>
- Keitel, D., Prix, R., Papa, M. A., Leaci, P., & Siddiqi, M. (2014). Search for continuous gravitational waves: Improving robustness versus instrumental artifacts. *Phys. Rev. D*, 89(6), 064023. <https://doi.org/10.1103/PhysRevD.89.064023>
- Klöckner, A., Pinto, N., Lee, Y., Catanzaro, B., Ivanov, P., & Fasih, A. (2012). PyCUDA and PyOpenCL: A Scripting-Based Approach to GPU Run-Time Code Generation. *Parallel Computing*, 38(3), 157–174. <https://doi.org/10.1016/j.parco.2011.09.001>
- LIGO Scientific Collaboration. (2018). *LIGO Algorithm Library - LALSuite*. free software (GPL). <https://doi.org/10.7935/GT1W-FZ16>
- Prix, R. (2009). Gravitational Waves from Spinning Neutron Stars. In W. Becker (Ed.), *Neutron Stars and Pulsars* (Vol. 357, pp. 651–685). Springer. https://doi.org/10.1007/978-3-540-76965-1_24

- Riles, K. (2017). Recent searches for continuous gravitational waves. *Mod. Phys. Lett. A*, 32(39), 1730035. <https://doi.org/10.1142/S021773231730035X>
- Sieniawska, M., & Bejger, M. (2019). Continuous gravitational waves from neutron stars: current status and prospects. *Universe*, 5(11), 217. <https://doi.org/10.3390/universe5110217>
- Vousden, W. D., Farr, W. M., & Mandel, I. (2015). Dynamic temperature selection for parallel tempering in Markov chain Monte Carlo simulations. *Mon. Not. Roy. Astron. Soc.*, 455(2), 1919–1937. <https://doi.org/10.1093/mnras/stv2422>
- Wette, K. (2020). SWIGLAL: Python and Octave interfaces to the LALSuite gravitational-wave data analysis libraries. *SoftwareX*, 12, 100634. <https://doi.org/10.1016/j.softx.2020.100634>