







Stingray: A Modern Python Library for Spectral Timing

Daniela Huppenkothen¹ , Matteo Bachetti² , Abigail L. Stevens^{3,4} , Simone Migliari^{5,6}, Paul Balm⁶, Omar Hammad⁷, Usman Mahmood Khan⁸, Himanshu Mishra⁹, Haroon Rashid¹⁰, Swapnil Sharma¹¹, Evandro Martinez Ribeiro¹² , and Ricardo Valles Blanco⁶

¹ DIRAC Institute, Department of Astronomy, University of Washington, 3910 15th Avenue NE, Seattle, WA 98195, USA; dhuppenk@uw.edu

² INAF-Osservatorio Astronomico di Cagliari, via della Scienza 5, I-09047 Selargius (CA), Italy

³ Department of Physics & Astronomy, Michigan State University, 567 Wilson Road, East Lansing, MI 48824, USA

⁴ Department of Astronomy, University of Michigan, 1085 South University Avenue, Ann Arbor, MI 48109, USA

⁵ ESAC/ESA, XMM-Newton Science Operations Centre, Camino Bajo del Castillo s/n, Urb. Villafranca del Castillo, SE-28692, Villanueva de la Caada, Madrid, Spain

⁶ Timelab Technologies Ltd., 20-22 Wenlock Road, London N1 7GU, UK

⁷ AinShams University, Cairo, Egypt

⁸ Department of Computer Science, North Carolina State University, Raleigh, NC, USA

⁹ Indian Institute of Technology, Kharagpur West Bengal 721302, India

¹⁰ National University of Sciences and Technology (NUST), Islamabad 44000, Pakistan

¹¹ Indian Institute of Technology Mandi, Mandi, Himachal Pradesh, India

¹² Kapteyn Astronomical Institute, University of Groningen, P.O. Box 800, NL-9700 AV Groningen, The Netherlands

Received 2019 January 22; revised 2019 May 10; accepted 2019 May 28; published 2019 August 9

Abstract

This paper describes the design and implementation of `stingray`, a library in Python built to perform time series analysis and related tasks on astronomical light curves. Its core functionality comprises a range of Fourier analysis techniques commonly used in spectral-timing analysis, as well as extensions for analyzing pulsar data, simulating data sets, and statistical modeling. Its modular build allows for easy extensions and incorporation of its methods into data analysis workflows and pipelines. We aim for the library to be a platform for the implementation of future spectral-timing techniques. We describe the overall vision and framework, core functionality, extensions, and connections to high-level command-line and graphical interfaces. The code is well tested, with a test coverage of currently 95%, and is accompanied by extensive Application Program Interface (API) documentation and a set of step-by-step tutorials.

Key words: methods: data analysis – methods: statistical – X-rays: binaries – X-rays: general

Software reviewed by the [Journal of Open Source Software](#) 

1. Introduction

Variability is one of the key diagnostics in understanding the dynamics, emission processes, and underlying physical mechanisms of astronomical objects. The flux of the majority of sources in the sky, from small asteroids to supermassive black holes, varies on timescales that can range from milliseconds to centuries, depending on the type of source. The detection of periodic variations in the radio flux of certain celestial objects has led to the ground-breaking discovery of pulsars (Hewish et al. 1968). Since this initial detection, the precise measurements of periods, spin-down effects, and intrapulse variations in pulsars across the electromagnetic spectrum from radio to gamma rays has led to new insights on the structure of neutron stars and their magnetic fields (e.g., Lorimer 2008; Abdo et al. 2013; Papitto et al. 2019). Similarly, accurate models of dips in stellar light curves have led to the discovery of thousands of exoplanets (e.g., Charbonneau et al. 2000; Henry et al. 2000; Coughlin et al. 2016). In asteroseismology, the detection of oscillatory modes in the power spectra generated from the light curves of stars, including our Sun, has allowed researchers a rare glimpse into the internal structure of these stars (see, e.g., Di Mauro 2016 for a recent review), and young stellar objects (YSOs) have been shown to be extremely variable (including rapid flaring) across the electromagnetic spectrum from radio to X-rays (see, e.g., Forbrich et al. 2017 for a comprehensive analysis of YSOs in the Orion Nebula Cluster), providing clues about the interaction of the forming stellar object and the circumstellar disk. Methods very similar to those employed in asteroseismology are also used to study the

interior of neutron stars and the dense-matter equation of state through oscillations in magnetar bursts (e.g., Huppenkothen et al. 2013) and giant flares (Israel et al. 2005; Strohmayer & Watts 2005; Watts & Strohmayer 2006). Analogous to studies of oscillations in magnetar giant flares and bursts, Beloborodov et al. (2000) and Guidorzi et al. (2016), among others, have probed the variability of the prompt emission of gamma-ray bursts in the quest to uncover the central engine of these sources, while in solar flares, oscillations may give clues about the nature of the magnetic reconnection that powers these flares (Inglis et al. 2016).

In a range of different astrophysical sources, including stars and compact objects, accretion plays a major role in their evolution and emission properties, giving rise to distinct patterns of brightness variations that allow us to study accretion physics in detail. For example, in white dwarfs, quasi-periodic variations attributed to magnetic gating of the accretion onto the star make it possible to measure the magnetic fields of these stars, even when they are too weak to be measured through other methods (Scaringi et al. 2017). More generally, connections can be found between accretion onto YSOs and onto compact objects including supermassive black holes, linking physics across many different scales in mass and time (Scaringi et al. 2015). Particularly in the study of black holes and neutron stars, the scientific developments of recent decades have brought a growing understanding that time and wavelength are intricately linked. Different spectral components react differently to changes in accretion rate and dynamics, leading to energy-dependent time lags, correlated variability, and higher-order effects (for a review, see Uttley et al. 2014).

This has led to the study of accretion disks, in particular those of active galactic nuclei (AGNs), via reverberation mapping (e.g., Blandford & McKee 1982; Bentz 2016), and probes of the accretion disk geometry using the energy dependence of quasi-periodic oscillations in stellar-mass black holes (e.g., Ingram & van der Klis 2015; Stevens & Uttley 2016). Understanding how the emission at various wavelengths changes with time is crucial for testing and expanding our understanding of general relativity in the strong-gravity limit, the dense-matter equation of state, and other fundamental questions in astrophysics. In X-ray astronomy, there is now a wealth of public data sets of variable objects from missions such as the *Rossi X-ray Timing Explorer* (*RXTE*; Bradt et al. 1993), the X-ray Multi-Mirror Mission (*XMM-Newton*; Jansen et al. 2001), the *Nuclear Spectroscopic Telescope Array* (*NuSTAR*; Harrison et al. 2013), *Astrosat* (Singh et al. 2014), and the *Neutron Star Interior Composition Explorer* (*NICER*; Gendreau et al. 2016). In addition, planned missions such as the *Advanced Telescope for High-Energy Astrophysics* (*Athena*; Barret et al. 2018) and proposed missions like the *Enhanced X-ray Timing Polarimeter* (*eXTP*; Zhang et al. 2016) and the *Spectroscopic Time-Resolving Observatory for Broadband Energy X-rays* (*STROBE-X*; Ray et al. 2018) will produce data sets of unprecedented size and complexity.

Motivated by the ubiquity of (spectral) timing in astronomy and the advent of these new data sets, we present *stingray*, a well-tested, open-source *Python* implementation of a range of core algorithms and methods used in time series analysis and spectral timing across the electromagnetic spectrum. The package has now been employed in a range of studies including radio observations of the galactic black hole X-ray binary Cygnus X-1 (Tetarenko et al. 2019), optical and X-ray observations of accreting pulsars (Brumback et al. 2018; Jaisawal et al. 2018; Kennedy et al. 2018), X-ray observations of accreting low-mass (Beri et al. 2019) and high-mass (Pike et al. 2019) X-ray binaries, ultraluminous X-ray sources (Walton et al. 2018), and statistical investigations of cospectra (Bachetti & Huppenkothen 2018; Huppenkothen & Bachetti 2018).

The paper layout is as follows. In Section 2, we very briefly describe the data sets being used in this paper to showcase the implemented methods. In Section 3, we lay out the overall vision, followed by a description of the general package structure and the general development framework in Section 4. The package’s core functionality is shown in more detail in Section 5, where we introduce basic classes for generating light curves and Fourier spectra of various types. In Sections 6–8, we present the submodules enabling the statistical modeling of Fourier products, simulating light curves from stochastic processes, and pulsar analysis, respectively. Sections 9 and 10 point to connections with a command-line interface and a graphical user interface that are being developed concurrently with *stingray*. Finally, in Section 11 we lay out our future development plans. Note that we intentionally omit code examples and specific implementation details in this manuscript in order to preserve longer-term accuracy. All code to reproduce the figures in this paper is available online,¹³ as is a full suite of up-to-date tutorials.¹⁴

2. Data

Throughout the paper, we use real observations of compact objects to demonstrate the functionality of the software in this package. In the following sections, we give brief introductions into the observations used and the data reduction processes applied before using the resulting event files and light curves with *stingray*.

2.1. GX 339–4

GX 339–4 is a stellar-mass black hole in a low-mass X-ray binary (Hynes et al. 2003). The black hole has a lower mass limit of $\sim 7 M_{\odot}$ (Muñoz-Darias et al. 2008) and possibly a near-maximal spin (Ludlam et al. 2015). The system also likely has a low binary orbit inclination; it has been constrained to $37^{\circ} < i \lesssim 60^{\circ}$ from optical and X-ray observations (Zdziarski et al. 1998; Heida et al. 2017), and spectral modeling by Wang-Ji et al. (2018) estimates $i \approx 40^{\circ}$. We use an observation from the *RXTE* Proportional Counter Array (PCA; Jahoda et al. 1996) in NASA’s High Energy Astrophysics Science Archive Research Center (HEASARC) from the 2010 outburst of GX 339–4 (Yamaoka et al. 2010), with the observation taken from UT 2010 April 22 23:36:52 to UT 2010 April 23 00:01:10 (observation ID 95409-01-15-06). This observation was taken in 64-channel event mode with $122 \mu\text{s}$ time resolution (E_125us_64M_0_1s). The following filtering criteria were used to obtain good time intervals (GTIs): Proportional Counter Unit (PCU) 2 is on, two or more PCUs are on, elevation angle $>10^{\circ}$, and target offset $<0^{\circ}.02$. Time since the South Atlantic Anomaly passage was not filtered on. Applying these filters, we have ~ 1 ks of good data. Since the observation is short, the data were not barycentered¹⁵ before analysis.

2.2. KIC12158940

KIC12158940 is one of 21 AGNs studied in Smith et al. (2018) and was observed for 12 epochs with the *Kepler Space Telescope* (Borucki et al. 2010). Smith et al. (2018) report an average *Kepler* magnitude of 14.85 (originally calculated in Brown et al. 2011) and a black hole mass of $\log M_{\odot} = 8.04$. The power spectrum shows a break at a characteristic timescale of $\tau_{\text{char}} = 31.6$ and a steep power spectral slope of $\Gamma = -3.3$.

We downloaded the preproduced light curves for all 12 epochs from the Barbara A. Mikulski Archive for Space Telescopes (MAST). We note that Smith et al. (2018) caution that because *Kepler* was designed for observing stars, these preproduced AGN light curves contain instrumental artifacts (e.g., too-small extraction apertures and rolling band noise) that render these light curves too imprecise to derive scientific conclusions from them. Because the purpose of this work is to showcase *stingray*’s capabilities on different types of data rather than deriving physical properties, we continue with the preproduced MAST light curves.

We did spot checks using a larger aperture to extract photometric fluxes from the full-frame images using the *Python* package *Lightkurve* (Barentsen et al. 2019) and find no discernible differences due to the smaller apertures used by the *Kepler* team. For astrophysical investigations, however, we urge the reader interested in analyzing AGN light curves

¹³ <https://github.com/StingraySoftware/stingraypaper>

¹⁴ <https://github.com/StingraySoftware/notebooks>

¹⁵ “Barycentering” the data applies a spacecraft clock correction to correct the photon arrival times to the solar system barycenter. This is commonly done with the FTOOL *barycorr* from HEASoft.

observed with *Kepler* to follow the methods laid out in Smith et al. (2018) to produce debiased light curves before using *stingray*. We split the *Kepler* light curves for KIC12158940 along data gaps to produce contiguous segments, and we rebinned all segments to a resolution of 0.075 day in order to generate evenly sampled light curves.

2.3. Hercules X-1

Hercules X-1 (Her X-1) is a well-known persistent X-ray binary pulsar with a period of $P = 1.23$ s (Tananbaum et al. 1972) in a binary system with an $\sim 2.2 M_{\odot}$ stellar companion HZ Herculis (Bahcall & Bahcall 1972; Davidsen et al. 1972; Forman et al. 1972; Reynolds et al. 1997; Leahy & Abdallah 2014) with an orbital period of $P_{\text{orb}} = 1.7$ days and superorbital variations on a ~ 35 day timescale (Giacconi et al. 1973; Scott & Leahy 1999; Igna & Leahy 2011). The companion’s type varies between late-type A and early-type B with orbital phase (Anderson et al. 1994; Cheng et al. 1995). For this work, we considered two of the several observations of Her X-1 with *NuSTAR*. The first observation was taken from UT 2012 September 19 to UT 2012 September 20 and was one among several used by Fürst et al. (2013) to characterize the cyclotron resonance scattering features (CRSFs) in the spectrum of the source. The second was taken from UT 2016 August 20 to UT 2016 August 21 and was used by Staubert et al. (2017) to detect an inversion of the decay of the CRSF. We used observation IDs 30002006002 and 10202002002 from HEASARC and barycentered (see footnote 15) the data with the latest (as of 2018 November 27) *NuSTAR* clock correction file. For our analysis, we considered photons from 3 to 79 keV at most $50''$ from the nominal position of the source, extracted from the two identical Focal Plane Modules A and B (FPMA and FPMB, respectively) on board the spacecraft. We used a total of 32.67 ks of good data in the first observation and 36.56 ks in the second, only selecting intervals longer than 10 s.

3. Vision and General Package Framework

Despite decades of research, the field of spectral timing in high-energy astrophysics is fragmented in terms of software; there is no commonly accepted, up-to-date framework for the core data-analysis tasks involved in (spectral) timing. Code is often siloed within groups, making it difficult to reproduce scientific results. Additionally, the scarcity of fully open-source tools constitutes a significant barrier to entry for researchers new to the field, since it effectively requires anyone not part of collaborations with an existing private code base to write their own software from scratch. The NASA library *xronos* is, to our knowledge, the only widely used open-source library in this field and has several shortcomings. In particular, it performs only a few of the most basic tasks, and it has not been maintained since 2004. Other open-source projects use languages that either require an expensive license (e.g., IDL) or have a limited scope (e.g., S-Lang). This dearth of software for spectral timing motivated the development of *stingray*,¹⁶ a library built entirely in the widely used Python language and based on *astropy* functionality. *stingray* aims to make many of the core Fourier analysis tools used in timing and spectral-timing analysis available to a large range of

researchers while providing a common platform for new methods and tools as they enter the field.

It includes the most relevant functionality in its core package, while extending that functionality in its subpackages in several ways, allowing for easy modeling of light curves and power spectra, simulation of synthetic data sets, and pulsar timing.

Its core idea is to provide time series analysis methods in an accessible, unit-tested way, built as a series of object-oriented modules. In practice, data analysis requirements are varied and depend on the type of data, the wavelength the observation was taken at, and the object being observed. With this in mind, *stingray* does not aim to provide full-stack data analysis workflows; rather, it provides the core building blocks for users to build such workflows themselves, based on the specific data analysis requirements of their source and observation. The modularity of its classes allows for easy incorporation of existing *stingray* functionality into larger data-analysis workflows and pipelines, while being easily extensible for cases that the library currently does not cover.

stingray separates out core functionality from several more specialized tasks based on those core classes and functions. Constructs related to data products as well as Fourier transforms of the data (e.g., power spectra, cross spectra, time lags, and other spectral timing products) are considered core functionality, as are some utility functions and classes, for example related to GTI calculations.

This core functionality is extended in various ways in currently three subpackages. The *modeling* subpackage (see also Section 6) provides a framework for modeling light curves and Fourier spectra with parametric functions. Based on this framework, it allows users to search for (quasi-)periodic oscillations in light curves with stochastic variability and provides convenience functions to aid standard tasks like fitting Lorentzian functions to power spectra.

The subpackage *simulator* (Section 7) provides important functionality to allow efficient simulation of time series from a range of stochastic processes. This includes simulation of light curves from power spectral models, as well as the use of transfer functions to introduce time lags and higher-order effects.

Finally, the subpackage *pulsar* implements a range of methods particularly useful for period searches in pulsars.

stingray is designed to be used as a standalone package and is also at the core of two other software packages currently under development: *HENDRICS* and *DAVE*. *HENDRICS* (Bachetti 2015; see also Section 9) provides prebuilt data-analysis workflows using *stingray* core functionality. These workflows are accessible from the command line and are provided for some common data types and data analysis tasks. *DAVE* (see also Section 10) provides a graphical user interface on top of *stingray* to allow for easy interactive exploratory data analysis.

As of v0.1, the core functionality of *stingray* depends exclusively on *numpy* (van der Walt et al. 2011), *scipy* (Jones et al. 2001), and *astropy* (The Astropy Collaboration et al. 2018), with optional plotting functionality supplied by *matplotlib* (Hunter 2007). The *modeling* subpackage optionally uses sampling methods supplied by *emcee* (Foreman-Mackey et al. 2013), some functionality implemented in *statsmodels*, and plotting using *corner* (Foreman-Mackey 2016). The *pulsar* subpackage optionally

¹⁶ *stingray* was named partly in homage to the popular 1960s children’s TV series, from which *stingray*’s motto derives: Anything can happen in the next half hour (including spectral timing made easy)!

allows for just-in-time compilation using `numba` (Lam et al. 2015) for computational efficiency, and for advanced pulsar timing models using `PINT`.¹⁷

This paper describes `stingray` v0.1, released on 2018 February 12. As with most open-source packages, `stingray` is under continuous development and welcomes contributions from the community, including suggestions for new sub-packages to be implemented.

4. Development and Integration Environment

`stingray` is developed entirely in Python 3, with backwards compatibility to Python 2.7 where possible through the integration package `six`. Development is version-controlled through `git` and officially hosted on GitHub through the organization *StingraySoftware*,¹⁸ where several interconnected repositories related to `stingray` live, including the core library, extension packages `HENDRICS` and `DAVE`, the suite of tutorials, the website, and this manuscript. All patches and code are submitted via pull requests to the `stingray` repository and checked by a maintainer for correctness of algorithms, adherence to standards of code, documentation, and tests. As an *Astropy* Affiliated Package,¹⁹ we follow the coding standards as well as community guidelines (including the Code of Conduct) set out by the *Astropy* community. All code within the `stingray` core library is subject to extensive unit testing, with compatibility across platforms as well as different versions of Python and required packages controlled through Continuous Integration services *Travis* (Unix platforms) and *AppVeyor* (Windows). Test coverage is checked using *Coveralls*. All user-facing functions and classes within `stingray` must have documentation in the form of docstrings, compiled and built along with the main documentation pages using *sphinx* and hosted on `readthedocs`.²⁰ Tutorials are provided in the form of executable *Jupyter* notebooks in a separate repository, (see footnote 14), which can either be run interactively using *Binder* (Jupyter et al. 2018) or viewed as part of the documentation.

5. Core Functionality

`stingray` imports its core functions and classes from the top-level package. These classes define the basic data structures such as light curves and cross as well as power spectra that are used in much of the higher-level functionality provided in the subpackages. Additionally, it incorporates a number of utilities for dealing with GTIs as well as input and output of data sets.

5.1. The *Lightcurve* Class

We expect `stingray` to be used largely on data sets of two forms: (1) event data (recordings of arrival times of individual photons) or (2) binned light curves (measurements of brightness in units of flux, magnitude, or counts as a function of time).

The majority of methods in `stingray` use binned light curves, which we thus currently consider the default format. The `Lightcurve` class defines a basic data structure to store binned light curves. Its attributes include arrays describing time

bins and associated (flux or counts) measurements, the number of data points in the light curve, the time resolution, and the total duration of the light curve. For unevenly sampled light curves, the time resolution Δt will be defined as the median difference between time bin midpoints. Users can pass uncertainties for measurements directly, or pass a string defining the statistical distribution of the data points for automatic calculation. By default, a Poisson distribution is assumed, appropriate for binned event data.

There are two ways to generate a `Lightcurve` object. In the standard case, the instrument has recorded a binned time series of N pairs of time stamps and count (rate) or flux values, $\{t_k, c_k\}_{k=1}^N$. In this case, one can simply instantiate a `Lightcurve` object with the keywords `time` and `counts` (and optionally set `use_counts = False` when the input is in units of counts per second). In cases where the native data format is events (e.g., photon arrival times), it is possible to use the static method `Lightcurve.make_lightcurve`, passing the array of events as well as a time resolution Δt to create a new light curve from the events.

Various operations are implemented for class `Lightcurve`. Custom behavior of the `+` and `-` operators allows straightforward addition and subtraction of light curves from one another. Assuming the light curves have the same time bins, the `+` and `-` operators will add or subtract the flux or counts measurements, respectively, and return a new `Lightcurve` object with the results. Other common operations implemented include time-shifting the light curve by a constant factor, joining two light curves into a single object, truncating a light curve at a certain time bin, and input/output operations to read or write objects from or to disk in various formats (HDF5, FITS, and ASCII are currently supported). For light curves that do not have consecutive time bins, there is a sorting operation, as well as the option to sort the light curve by the ascending or descending flux or counts.

We provide support for GTIs in many methods and implement rebinning the light curve to a new time resolution larger than the native resolution of the data (interpolation to a finer resolution is currently not supported). In Figure 1 (left panel), we show an example observation of GX 339–4 as taken with *RXTE*, and in Figure 2 (left panel) a *Kepler* observation of KIC12158940. `stingray` implements basic methods for plotting (useful for a quick look at the data).

5.2. The *Events* Class

At short wavelengths, data are largely recorded as *photon events*, where arrival times at the detector are recorded for each photon independently, along with a number of other properties of the event (for example, an energy channel in which the photon was recorded in, which can be transformed to a rough estimate of the energy of the original photon arriving at the detector).

Even for a single instrument, multiple types of data can often be recorded, resulting in a plethora of data formats and internal schemas for how data are stored within the binary files distributed to the community. `stingray` implements a basic `EventList` class that acts as a container for event data, but does not aim to encompass all data types of all current (and future) instruments. Instead, it aims to abstract away from instrument-specific idiosyncrasies as much as possible and remain mission agnostic. In its basic form, it takes arrays with time stamps and optionally corresponding photon energies as

¹⁷ <https://github.com/nanograv/PINT>

¹⁸ <https://github.com/StingraySoftware/>

¹⁹ <http://www.astropy.org/affiliated/>

²⁰ <https://stingray.readthedocs.io/en/latest/>




Figure 1. Left panel: ~ 1 ks *RXTE* observation of the black hole X-ray binary GX 339–4. Details of the observation can be found in Section 2.1. In gray, we show the light curve produced by binning the events into 0.02 s bins. The blue line corresponds to the rebinned light curve at $dt = 1.0$ s. Right panel: the power spectrum calculated from the light curve in the left panel (gray), as well as a version of the same power spectrum that has been linearly rebinned (blue) and logarithmically rebinned (orange) in frequency.

input and implements a set of basic methods. Similarly to `Lightcurve`, it provides basic input/output (I/O) functionality in the form of `read` and `write` methods, as well as a method to join event lists, which can be particularly useful when data are recorded in several independent detectors, as is common for several current and future X-ray missions. The `to_lc` method provides straightforward connection to create a `Lightcurve` directly out of an `EventList` object. In return, it is possible to create an `EventList` out of a `Lightcurve` object using the `from_lc`. The latter will create N_i events, each with a time stamp equation to the time bin t_i , where N_i is the number of counts in bin i (event lists are, by their very definition, only a useful data product if the light curve used to simulate comes from photon counting data in the first place). It is possible to simulate more physically meaningful photon events from a given light curve and energy spectrum using the `simulate_times` and `simulate_energies` methods (from the `simulator` package, Section 7), which employ a combination of interpolation and rejection sampling to accurately draw events from the given light curve and spectrum.

5.3. Cross Spectra and Power Spectra

The cross spectrum and the power spectrum²¹ are closely related (for a pedagogical introduction to Fourier analysis, see van der Klis 1989; see also Uttley et al. 2014 for a recent review of spectral timing techniques). Computing the cross spectrum requires two evenly sampled time series $y_1 = \{y_{1,i}\}_{i=1}^N$ and $y_2 = \{y_{2,i}\}_{i=1}^N$ taken simultaneously at exactly the same time intervals $\{t_i\}_{i=1}^N$. Under this assumption, one may then

compute the discrete Fourier transform of each time series, \mathcal{F}_1 and \mathcal{F}_2 , independently, and multiply \mathcal{F}_1 with \mathcal{F}_2^* , that is, the Fourier transform of y_1 with the *complex conjugate* of the Fourier transform of y_2 .

Because the power spectrum is defined as the square of the real part of the Fourier amplitudes of a single evenly sampled time series, it can be formulated as the special case of the cross spectrum where $y_1 = y_2$. In `stingray`, we implement a class `Crossspectrum`, which takes two `Lightcurve` objects as input and internally calculates the complex cross spectrum in one of a number of common normalizations (see below). Because many of the internal calculations are the same, the class `Powerspectrum` is implemented as a subclass of `Crossspectrum`, but takes only a single `Lightcurve` object instead of two.

There are several popular normalizations for the real part of the cross spectrum as well as the power spectrum implemented in `stingray`: the *Leahy normalization* (Leahy et al. 1983a) is defined such that for simple white noise, the power spectrum will follow a χ^2 distribution with two degrees of freedom around a mean value of 2, and the cospectrum—the real part of the cross spectrum—will follow a Laplace distribution centered on zero with a scale parameter of 1 (Huppenkothen & Bachetti 2018). It is particularly useful for period searches, because the white noise level is well understood and always the same (but be aware that detector effects like dead time can distort the power spectrum in practice; Bachetti et al. 2015). For light curves with complex variability patterns, and especially for understanding how these patterns contribute to the overall variance observed, the *fractional rms-squared normalization* (Belloni & Hasinger 1990; Miyamoto et al. 1992) or the *absolute rms-squared normalization* (Uttley & McHardy 2001) may be more appropriate choices.

The classes `Crossspectrum` and `Powerspectrum` share most of the implemented methods, except where otherwise noted. Both classes include methods to rebin cross and power spectra. Linear rebinning is implemented analogously to the method in class `Lightcurve`. Additionally, logarithmic binning is implemented in the method `rebin_log` in such a

²¹ In the signal processing literature, generally a distinction is made between the power spectrum, which describes the process at the source generating variable time series, and the periodogram, which denotes a realization of said power spectrum, that is, the time series we actually observe, which is an *estimator* of the underlying process. While the products generated by `stingray` are generally derived from data, and therefore periodograms, the astronomy literature usually denotes them by the term power spectrum. We follow this convention here, as we do within the software package itself.




Figure 2. Left panel: light curve of all 12 epochs of the *Kepler* observation of the AGN KIC12158940, rebinned to 0.075 day. Details of the observation can be found in Section 2.2. While there are error bars on the *Kepler* data points, they are so small on this scale as to be virtually invisible. The right panel shows the corresponding averaged power spectrum, using a total of 50 segments, each with a length of 10 days.

way that the bin width at a given frequency increases by a fraction of the previous bin width:

$$d\nu_{i+1} = d\nu_i(1 + f),$$

where f is some constant factor by which the frequency resolution increases, often $f = 0.01$.

Classical period searches are often formulated as outlier detection problems from an expected statistical distribution. Assuming the signal is sufficiently coherent such that all of the signal power is concentrated in one bin, one may calculate the chance probability that an observed power in the spectrum was generated by statistical fluctuations alone. For the white noise case, the equations to accurately calculate a p value of rejecting the hypothesis that a given outlier in the power spectrum was generated by noise are defined in Groth (1975), and they can be calculated for one or multiple powers in a `Powerspectrum` object using the `classical_significances` method, which enables computation of a (trial-corrected) p value for a given power in the presence of white noise. Note that the cross spectrum does not follow the same distribution (Huppenkothen & Bachetti 2018), and the recently derived statistical distributions for this case will be implemented in a future version of `stingray`.

In many practical applications, users may wish to average power or cross spectra from multiple light-curve segments in order to suppress statistical noise. This can be done with the appropriate classes `AveragedPowerspectrum` and `AveragedCrossspectrum`, which take a `Lightcurve` object or list of `Lightcurve` objects as an input and will compute averaged Fourier products by dividing the light curve into N segments of a given size τ_{seg} . The Fourier spectra (either cross spectra or power spectra) are averaged together. Both are subclasses of `Crossspectrum`, and either inherit or override many of the methods relevant for those classes as well. Examples of the kinds of products produced by the classes and methods introduced above are given in Figures 1 and 2 (right panels).

For averaged cross spectra, it is possible to calculate the time lag between variability in two simultaneous light curves, for

example, if the two light curves cover different energy bands (Vaughan et al. 1994). The time lag τ_j is defined as

$$\tau_j = \frac{\phi_j}{2\pi\nu_j}$$

for a phase angle ϕ_j derived from the imaginary component of the complex cross spectrum and a midbin frequency ν_j . Similarly, it is possible to calculate the coherence from the cross spectrum (Vaughan & Nowak 1997; Nowak et al. 1999), defined as

$$c_j = \frac{C_{xyj}}{C_{xj}C_{yj}}. \quad (1)$$

Here, C_{xyj} corresponds to the real part of the unnormalized cross spectrum, and C_{xj} and C_{yj} correspond to the analogous squared amplitudes of the power spectrum for each individual light curve. The errors on τ_j and c_j are also computed in `stingray`.

For long observations with quasi-periodic oscillations (QPOs), spectrograms, more commonly known in the astronomy literature as dynamic power spectra, can be a useful way to track changes in the QPO centroid frequency over time. We have implemented `DynamicalPowerspectrum` as a subclass to `AveragedPowerspectrum` to provide this functionality. Like `AveragedPowerspectrum`, this class takes a `Lightcurve` object and a segment size as input, but instead of averaging the power spectra of each individual segment, it will create a matrix of time bins (one bin for each segment) as columns and Fourier frequencies as rows. Rebinning both along the time and frequency axes is possible. Moreover, the method `trace_maximum` automatically finds the frequency with the highest power in each segment in a given range of frequencies and traces this maximum over time. An example using data from the source GX 339-4 is shown in Figure 3.

Closely related to the cross spectrum and power spectrum are the cross correlation and the autocorrelation, implemented in classes `CrossCorrelation` and `AutoCorrelation`. As their respective Fourier spectra equivalents, they take either one




Figure 3. Example of a dynamic power spectrum generated from the GX 339 light curve shown in Figure 1. We generated 63 light-curve segments of 16 s length with a 0.02 s time resolution and Fourier-transformed each to generate a power spectrum. The dynamic power spectrum here plots each power spectrum as a vertical slice as a function of time, with the color indicating the fractional rms-squared-normalized power in each bin (yellow are large powers; purple, small). The dynamic power spectrum was clipped to around the range of the QPO and smoothed using bicubic interpolation to improve clarity. The QPO is clearly visible as a yellow streak and seems not to be present during the entire observation (consistent with Belloni et al. 2005). In red, we show the frequency with the highest power found in each segment (excluding frequencies below 3 Hz to exclude the low-frequency red noise), using the `trace_maximum` method.

(autocorrelation) or two (cross correlation) `Lightcurve` objects as input and compute the correlation between the two light curves or of the single light curve with itself, along with the time lags for which the correlation was produced and the time lag at which the maximum correlation is measured.

It is useful to note that all classes in this section are compatible with GTIs. The classes `Powerspectrum` and `Crossspectrum` will generate warnings if the observations contain gaps; their averaged versions will take GTIs correctly into account by producing power spectra only from light-curve segments for which data are available.

6. The modeling Subpackage

Modeling data sets with parametric (often physically motivated) models that map an independent variable (e.g., time or frequency) to one or more dependent variables (e.g., flux, counts, or Fourier powers) is a common task in astronomy. Constructing a universal modeling framework is a highly nontrivial task, and excellent packages exist for general-purpose model building (e.g., STAN, Carpenter et al. 2017). Thus, `stingray`'s modeling interface restricts itself to models of commonly used spectral-timing products, in particular (averaged) power spectra. While it makes heavy use of the `astropy.modeling.FittableModel` definitions, it uses custom definitions for fitting algorithms motivated by the statistical properties of spectral-timing products, which deviate significantly from other data types commonly found in astronomy and thus cannot easily be modeled with standard approaches defined in `astropy`.

The modeling subpackage logically separates out statistical models—likelihoods and posteriors—from the fitting functionality, such that different likelihoods and posteriors can be straightforwardly dropped in and out depending on the data set and problem at hand. In line with the overall philosophy of `stingray`, the modeling subpackage is designed to be

modular and easily extensible to specific problems a user might try to solve, while many typical tasks one might do with Fourier products are already built in. It makes use of the `scipy.optimize` interface for optimization, as well as the package `emcee` for Markov chain Monte Carlo (MCMC) sampling.

6.1. Statistical Models

All statistical models are implemented as a subclass of an Abstract Base Class `Likelihood` in module `stingray.posterior`. In its most basic form, each subclass of `Likelihood` takes data in some form (most commonly two arrays, one with the independent and one with the dependent variable) as well as an object of type `astropy.modeling.FittableModel`. The likelihood computes model values for each data point in the array of independent variables and statistically compares these model values with the data points stored in the dependent variable, assuming the particular statistical distribution of the likelihood definition. The result is a single scalar, which can then be, for example, used in an optimization algorithm in order to find a maximum likelihood (ML) solution.

For all likelihoods in `stingray`, an equivalent subclass of `stingray.modeling.Posterior` is available, which uses the `Likelihood` definitions to compute posterior probability densities for the parameters of a model given data. All subclasses of `Posterior` also require definition of a `logprior` method, which calculates the value of the prior probability density of the parameters. Because priors are strongly problem-dependent, they cannot be hard-coded into `stingray`. Even for relatively straightforward problems such as modeling quasi-periodic oscillations of X-ray binaries, the physical properties and their effect on the data can differ strongly from source to source, indicating that a prior set for XTE J1550–564 may not be appropriate for, for example, GRS 1915+105. Separating out the likelihood and posterior in distinct classes makes it possible to allow the use of the likelihood for maximum likelihood estimation, while requiring priors for estimating the Bayesian posterior probability through, for example, MCMC simulations.

`Loglikelihood` and `Posterior` subclass definitions currently exist within `stingray` for different statistical models useful in the context of astronomical data. `GaussianLogLikelihood` and `GaussianPosterior` implement statistical models for data with normally distributed uncertainties. `GaussianLogLikelihood` will compute what astronomers generally call χ^2 , because the likelihood calculated by this statistical model generally follows a χ^2 distribution with $N - P$ degrees of freedom (where N is the number of data points and P the number of free parameters). Note, however, that this is not the same as the χ^2 likelihood defined below!

`PoissonLogLikelihood` and `PoissonPosterior` calculate the likelihood and posterior for Poisson-distributed data, respectively. This likelihood is equivalent to what in astronomy is often called the Cash statistic (Cash 1979) and is the appropriate likelihood to use for count- or event-type data often found in X-ray astronomy time series and spectra.

`PSDLogLikelihood` and `PSDPosterior` implement the statistical model appropriate for modeling (averaged) power spectra, a χ^2 distribution. We broke with the rule of naming likelihoods and posteriors after the statistical distribution they

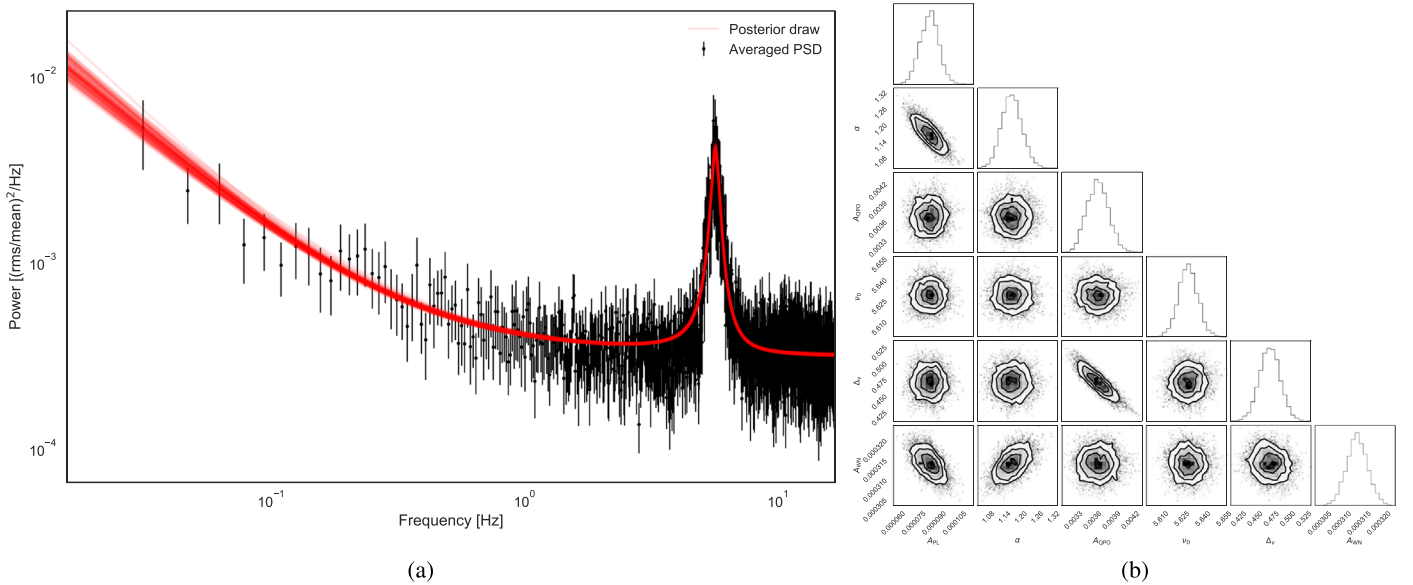


Figure 4. Left panel: in black, a power spectrum averaged out of 15 light-curve segments of 64 s each of the GX 339–4 observation, along with draws from the posterior distribution of the power-law model plus the Lorentzian QPO model and constant used to represent the data (red). Right panel: corner plot showing the marginal posterior distributions (diagonal) of the six parameters of the model: the amplitude of the power law A_{PL} , the power-law index α , the amplitude of the Lorentzian A_{QPO} , the QPO centroid frequency ν_{QPO} , the width of the QPO $\Delta\nu$, and the amplitude of the white noise, A_{WN} . The right-hand figure was produced using the package `corner` (Foreman-Mackey 2016).

implement in this case, because, as mentioned above, astronomers tend to call the likelihood for normally distributed data χ^2 , and this naming helps avoid any confusion. These two classes implement a χ^2 distribution for Fourier spectra generated with the `Powerspectrum` class, and a χ^2_{MK} distribution for power spectra generated with the `AveragedPowerspectrum` class, where M is the number of averaged segments and K is the number of averaged neighboring frequency bins. Please note that as laid out in Huppenkothen & Bachetti (2018), these distributions are not appropriate for use on (averaged) cross spectra. The appropriate distributions for these products are under development for the next version of `stingray`.

Other statistical models can be easily implemented by subclassing the `LogLikelihood` and `Posterior Abstract Base Classes` and using the existing classes as a template.

6.2. General Parameter Estimation and Model Comparison Functionality

`stingray` implements utility functions in order to reduce some of the overhead required for standard parameter estimation and model comparison tasks. In particular, the `parameterestimation` module implements classes and functions to aid users in fitting models to data and estimating the probability distributions of parameters.

The class `ParameterEstimation` provides the basis for more sophisticated, specialized implementations for particular data types. Its core methods are `fit` and `sample`. The former takes an instance of a `LogLikelihood` or `Posterior` subclass and uses minimization algorithms implemented in `scipy.optimize` to find the ML or maximum a posteriori (MAP) solution. The `sample` method uses the affine-invariant MCMC sampler implemented in `emcee` (Foreman-Mackey et al. 2013) to generate samples from a posterior distribution passed as an instance of a subclass of `Posterior`. Note that *you should never pass a `LogLikelihood` instance into the `sample` method*, because sampling from a likelihood is statistically

invalid. In addition to these core methods, higher-level functionality implemented in this class includes calculating the likelihood ratio test (LRT) for two different models M_1 and M_2 via the `compute_lrt` method (note the statistical assumptions of the LRT and where they fail, e.g., Protassov et al. 2002). In addition, the `calibrate_lrt` method allows calibrating the p value for rejecting the model M_1 via simulations of M_1 , using either an MCMC sample (for Bayesian inference and posterior predictive p values) or the covariance matrix derived from the optimization (both Bayesian and ML approaches).

`stingray` also implements two classes that summarize results of the optimization and sampling procedures in concise, useful ways. The `fit` method returns an instance of class `OptimizationResults`. This contains the most important outputs from the optimizer, but will also behind the scenes calculate a number of useful quantities, including the covariance between parameters (or a numerical approximation for some minimization algorithms), the Akaike and Bayesian information criteria (AIC: Akaike 1974; BIC: Schwarz 1978), and various summary statistics.

Similarly, an instance of class `SamplingResults` is returned by the `sample` method, which returns the posterior samples calculated by the MCMC sampler, as well as computes a number of helpful quantities using the MCMC chains. It calculates useful diagnostics including the acceptance fraction, the autocorrelation length, and the Rubin–Gelman statistic (Gelman & Rubin 1992) to indicate convergence, and infers means, standard deviations, and user-defined credible intervals for each parameter. An example of posterior inference, a power spectrum with a QPO and a broadband red noise component, is shown in Figure 4, including posterior probability distributions for the model parameters.

6.3. Special Functionality for Fourier Products

The subclass `PSDParEst` implements a number of additional methods particularly useful for modeling power




Figure 5. Left: power spectral shape generated using a compound `astropy.modeling.models` object of a power law and a Lorentzian. Right: corresponding light curve generated by the `simulator` subpackage with a time resolution of 0.05 s, a total duration of 15 ks, a mean count rate of 40 counts s^{-1} , and a fractional rms amplitude of 0.2. In blue, we show a binned version of the same light curve.

spectra. One particularly common task is to search for periodic signals (e.g., from pulsars) in a power spectrum, which reduces to finding outliers around an assumed power spectral shape (assuming the signal is *strictly* periodic, and thus all power is approximately concentrated in one bin). In the presence of other variability, the probability of observing a certain power P_j at a frequency ν_j under the assumption that no periodic signal is present depends on the shape and parameters of the underlying power spectral model assumed to have generated the data. As Vaughan (2010) shows, there is an inherent uncertainty in our inference of the parameters of this power spectral model, which must be taken into account via simulations. PSDParEst implements a method `calibrate_highest_outlier`, which finds the k highest outliers (where k is a user-defined number) and calculates the posterior predictive p value that said outliers cannot be explained by noise alone. It makes heavy use of the method `simulate_highest_outlier`, which uses the `sample` method to derive an MCMC sample and then simulate fake power spectra from that model for a range of plausible parameter values in order to include our model uncertainty in the posterior predictive p value. For details of the overall procedure, see Vaughan (2010).

As of this version, the `stingray.modeling` subpackage has no functionality to model higher-order Fourier products. For spectral timing in particular, this would involve being able to read and apply instrument responses to models, as well as being able to interface with the library of spectral models associated with the X-ray spectral fitting tool XSPEC (Arnaud 1996). Providing this functionality is planned for a future release of `stingray`.

7. The `simulator` Subpackage

The `simulator` subpackage contains a number of methods to generate simulated light curves out of known power spectral shapes and generate event lists from light curves.

7.1. Simulating Light Curves from Input Power Spectra

The basic `Simulator` object uses the algorithm from Timmer & Koenig (1995) to generate light curves out of a spectral shape. The spectral shape can be input as a spectral power-law index, `astropy.modeling.models` objects,

and a user-given array of powers. In Figure 5, we present a light curve as generated by a given power spectral shape. The output is a `Lightcurve` object that can be used like real data sets, including all functionality related to GTIs, spectral-timing products, and modeling.

7.2. Use Transfer Functions on Light Curves

Most astrophysical signals we receive in our instruments are a mixture of different input signals. Often, a signal emitted in one region can be reflected and re-emitted from another region or filtered in different ways. Spectral-timing studies can decompose the signals and try to understand how the signal is transformed by these phenomena between the emission region and the observer (see Uttley et al. 2014 for a review). This transformation can be encoded in an impulse response function, which describes the response of the system to a delta-function impulse. This is the Fourier transform of another well-known quantity in signal processing, the transfer function (see Girod et al. 2001). The `Simulator` object is capable of generating multiple light curves starting from an initial light curve and multiple input responses, mimicking observations in different energy bands.

7.3. Simulating Event Lists from Light Curves

The `simulator.base.simulate_times` method is able to simulate event lists from input light curves. It implements the acceptance–rejection method:

1. Generate a light curve (and smooth out any Poisson noise if generated through the Timmer & Koenig 1995 method) over the whole observation; normalize it so that the maximum is 1.
2. Generate an event, with uniform probability over the observing time.
3. Associate to this event a uniform random number \mathcal{P} between 0 and 1.
4. If \mathcal{P} is lower than the normalized light curve at the event time, *accept* the event, otherwise *reject* it.

In `stingray`, we use arrays of events for better performance, using the functionality contained in the `numpy` library.

8. The pulse Subpackage

The subpackage `pulse` contains the basic operations to perform the search and characterization of pulsed signals for use, for example, in searches of X-ray pulsars.

8.1. Epoch Folding

Among the basic algorithms used in pulsar astronomy, one cannot overstate the importance of epoch folding (EF). The algorithm consists of cutting the signal at every pulse period and summing all subintervals in phase. An alternative way of seeing it, more useful for photon data, is as a *histogram of pulse phases*.

If the period is exactly correct and assuming a stable pulsation, the signal-to-noise ratio will get better approximately with the square root of the number of summed subintervals. This is the method used to obtain practically all pulse profiles shown in the literature, as most pulsar signals are orders of magnitude below the noise level.

The `pulse.pulsar` submodule contains the functionality to calculate the phase given a simple pulse ephemeris consisting of any number of pulse frequency derivatives, or using a number of methods for the orbit of the pulsar (using the optional dependency `PINT`). Moreover, the module also includes a mechanism to calculate the exposure of single bins in the pulse profile. This is particularly useful for very-long-period pulsars where the pulsed period is comparable to the length of the GTIs. The different exposure of pulse bins caused by the absence of signals during GTIs is taken into account in the calculation of the final pulse profile by the folding algorithm, if the user asks for it.

8.2. Epoch Folding Searches and Z_n^2 Searches

During a search for pulsations, the first step is usually calculating a power spectrum through a fast Fourier transform (FFT). However, often pulsations do not leave a clear signature above the noise level in the power spectrum, because they are weak or they fall close to bin edges, where the sensitivity is reduced.²² Even when the signature is clear, the frequency resolution of the power spectrum is often inadequate to measure precisely the pulse frequency. Therefore, an additional statistical analysis is needed.

`stingray` implements two statistical methods for pulsar searches, which can be applied to event lists or light curves (that are treated as event lists with “weights”).

The epoch folding search (EFS) method consists of executing the folding at many trial frequencies around the candidate frequency. Once the folding is performed, the following statistics are calculated on the profile:

$$\mathcal{S} = \sum_i \frac{(P_i - \bar{P})^2}{\sigma^2} \quad (2)$$

where P_i are the bins of the profile, \bar{P} is the mean level of the profile, and σ is the standard deviation. Here, \mathcal{S} is the summed squared error of the actual pulsed profile with respect to a *flat* model and follows a χ^2 distribution.

²² This is due to the convolution of the signal with the observing window, which produces a sinc-like response inside the bins of the FFT; periodic signals with the same amplitude are detected with a lower Fourier amplitude if they fall far from the center of the spectral bin (van der Klis 1989).

If there is no pulsation, \mathcal{S} will assume a random value distributed around the number of degrees of freedom $n - 1$ (where n is the number of bins in the profile) with a well-defined statistical distribution (χ_{n-1}^2) that allows an easy calculation of detection limits. When observing a peak of a given height is very unlikely under the null hypothesis (meaning that the probability to obtain this peak by noise is below a certain ϵ), this peak is considered a pulse candidate. If the frequency resolution is sufficiently high, close to the correct frequency, as described by Leahy et al. (1983b) and Leahy (1987), the peak in the epoch folding periodogram has the shape of a sinc² function whose width is driven by the length T of the observation (FWHM $\Delta\nu \sim 0.9/T$).

The epoch folding statistic, however, can give the same value for a pulse profile at the correct frequency and, for example, a harmonic that produces a deviation from a Poisson distribution. A more effective method is the Z_n^2 statistics (Buccheri et al. 1983), which is conceptually similar to EF but has high values when the signal is well described by a small number of sinusoidal harmonics:

$$Z_n^2 = \frac{2}{N} \sum_{k=1}^n \left[\left(\sum_{j=1}^N \cos k\phi_j \right)^2 + \left(\sum_{j=1}^N \sin k\phi_j \right)^2 \right], \quad (3)$$

where N is the number of photons, n is the number of harmonics, and ϕ_j are the phases corresponding to the event arrival times t_j ($\phi_j = \nu t_j$, where ν is the pulse frequency).

The Z_n^2 statistics defined in this way, far from the pulsed profile, follow a χ_n^2 distribution, where n is the number of harmonics this time. This allows us, again, to easily calculate thresholds based on the probability of obtaining a given Z_n^2 by pure noise.

The standard Z_n^2 search calculates the phase of each photon and calculates the sinusoidal functions above for each photon. This is very computationally expensive if the number of photons is high. Therefore, in `stingray`, the search is performed by binning the pulse profile first and using the phases of the folded profile in the formula above, multiplying the squared sinusoids of the phases of the pulse profile by a weight corresponding to the number of photons at each phase.

$$Z_n^2 \approx \frac{2}{\sum_j w_j} \sum_{k=1}^n \left[\left(\sum_{j=1}^m w_j \cos k\phi_j \right)^2 + \left(\sum_{j=1}^m w_j \sin k\phi_j \right)^2 \right]. \quad (4)$$

Since the sinusoids are only executed on a small number of bins, while the epoch folding procedure just consists of a very fast histogram-like operation, the speedup of this new formula is obvious. Care must be put into the choice of the number of bins, in order to maintain a good approximation even when the number of harmonics is high. We recommend in the documentation to use a number of bins at least 10 times larger than the number of harmonics.

8.3. Characterization of Pulsar Behavior

As seen in Section 8.2, the Z_n^2 or EF periodograms of a perfectly stable pulsation have the shape of a sinc² function. `stingray` has functionality to fit these periodograms with a

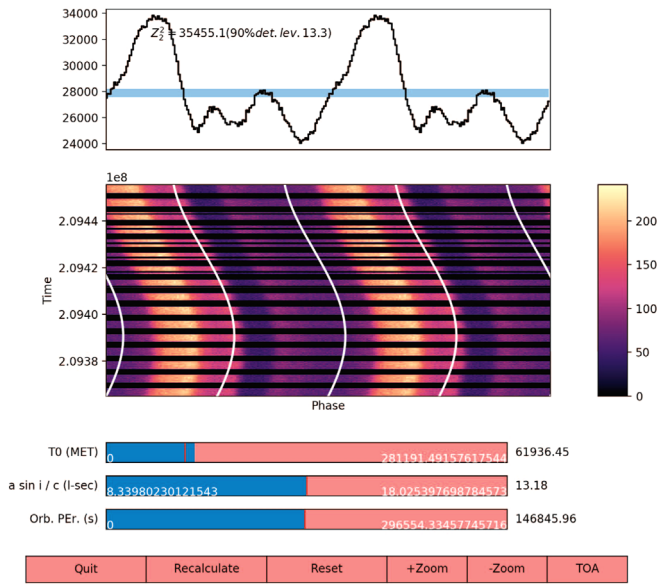


Figure 6. Phaseogram showing the variation of the pulse phase corresponding to an imperfect orbital solution (in this case the time at the ascending node T_0) in a *NuSTAR* observation of Her X-1, executed with *stingray* and plotted in a convenient, interactive interface with HENDRICS. The time-of-arrival button allows the user to calculate the time of arrival for use with Tempo2, PINT, or similar programs.

sinc^2 function or alternatively a Gaussian model, and find the mean frequency with high precision.²³

A significant deviation from the expected shape from these models can happen if the pulsation is not stable. Calculating the *phaseogram* (Figure 6) is an option to investigate how the pulse phase varies in time. The phaseogram in this context consists of a 2D histogram of the phase and arrival times of the pulses. If the pulsation is stable and the pulse frequency was determined with precision, the phaseogram shows vertical stripes corresponding to perfectly aligned pulses. If the frequency is not as precise, the stripes become more and more diagonal. If the pulse has a detectable frequency derivative, these stripes bend with a parabolic shape. If the orbital solution is imperfect, the stripes show specific periodic features.²⁴

A very precise way to determine the exact pulse ephemeris is out of the scope of *stingray*. Nonetheless, *stingray* has a mechanism to calculate the pulse arrival times (or times of arrival, TOAs) to be analyzed with more specialized software like Tempo, Tempo2, or PINT. We use the same `ffitit` algorithm used for radio pulsars (Taylor 1992), which calculates the cross correlation between a template profile and the folded profile in the Fourier domain. This is implemented in the `get_TOA` function in *stingray*. `pulse.pulsar`.

The functionality to plot the phaseogram, interactively change the timing parameters (either pulse parameters or orbital parameters) and adjust the solution, and calculate the TOAs for use with external programs is conveniently accessible in HENDRICS (see Section 9) and Figure 6 and in DAVE (Section 10).

²³ When using the Gaussian model, the width of the impulse is similar to the FWHM $\Delta\nu \sim 0.9/T$ of the sinc^2 function (Section 8.2). It does *not* represent an error bar to the frequency measurement.

²⁴ See for example https://github.com/matteobachetti/timing-lectures/blob/master/no-binder/Timing_residuals.ipynb.

9. HENDRICS: A Command-line Interface for *stingray*

The HENDRICS package²⁵—formerly called MaLTPyNT (Bachetti 2015)—builds upon *stingray* by providing a suite of easy-to-execute command-line scripts whose primary use is providing an accurate quick-look (spectral) timing analysis of X-ray observations, useful for a range of use cases, including exploratory data analysis and quality assessment of larger data-analysis pipelines. While its initial development proceeded independently from *stingray*, much of its core functionality since version 3.0 is based on the classes and methods *stingray* provides, and some key functionality has been shifted to *stingray* where appropriate.

The key feature distinguishing it from established command-line interfaces such as FTOOLS is the accurate treatment of gaps in the data (for example due to the Earth’s occultation or the South Atlantic Anomaly), as well as its treatment of dead time for certain detectors like *NuSTAR*. Where *stingray* aims to provide flexible building blocks for designing sophisticated spectral-timing analysis workflows, HENDRICS provides an end-to-end solution for common tasks such as power and cross spectra, time lags, pulsar searches, and color-color as well as color-intensity diagrams, at the cost of losing some flexibility during the creation of those products. Like *stingray*, HENDRICS is an *astropy*-affiliated package and aims to build upon and be compatible with functionality provided as part of the *astropy* ecosystem. HENDRICS supports a range of output data formats, including netCDF4 and ASCII formats, which can then be read into other astronomical data analysis systems such as XSPEC (Arnaud 1996) or ISIS (Houck & Denicola 2000).

HENDRICS is in release version 4 as of 2018 February 12 and under active development, utilizing the same continuous integration, testing, and code review standards as *stingray*.

10. DAVE: Exploratory Data Analysis in a Graphical User Interface

DAVE²⁶—the Data Analysis for Variable Events package—is a graphical user interface built on top of *stingray* in order to provide users with interactive capabilities for exploratory data analysis of variable time series. Much of the core functionality within *stingray* is available in DAVE as well: creation of power spectra, cross spectra, dynamical power spectra, and spectral-timing products such as time lags and coherence measurements. In addition, it implements interactive filtering of light curves with respect to energy channels or energies (if a response matrix file is loaded), time ranges, and count rates. Users may compare light curves and power spectra from different energy ranges and may create auxiliary products such as color-color and color-intensity diagrams that further aid the exploration of the data. An example of the interface is shown in Figure 7. The full interface and its capabilities will be described in a future publication.

11. Future Development Plans

Near- and medium-term plans for *stingray* development are largely aimed at extending current functionality related to Fourier spectra and continuing work toward comprehensive spectral-timing capabilities. While open-source reference

²⁵ <https://github.com/stingraySoftware/hendrics>

²⁶ <https://github.com/stingraySoftware/dave>

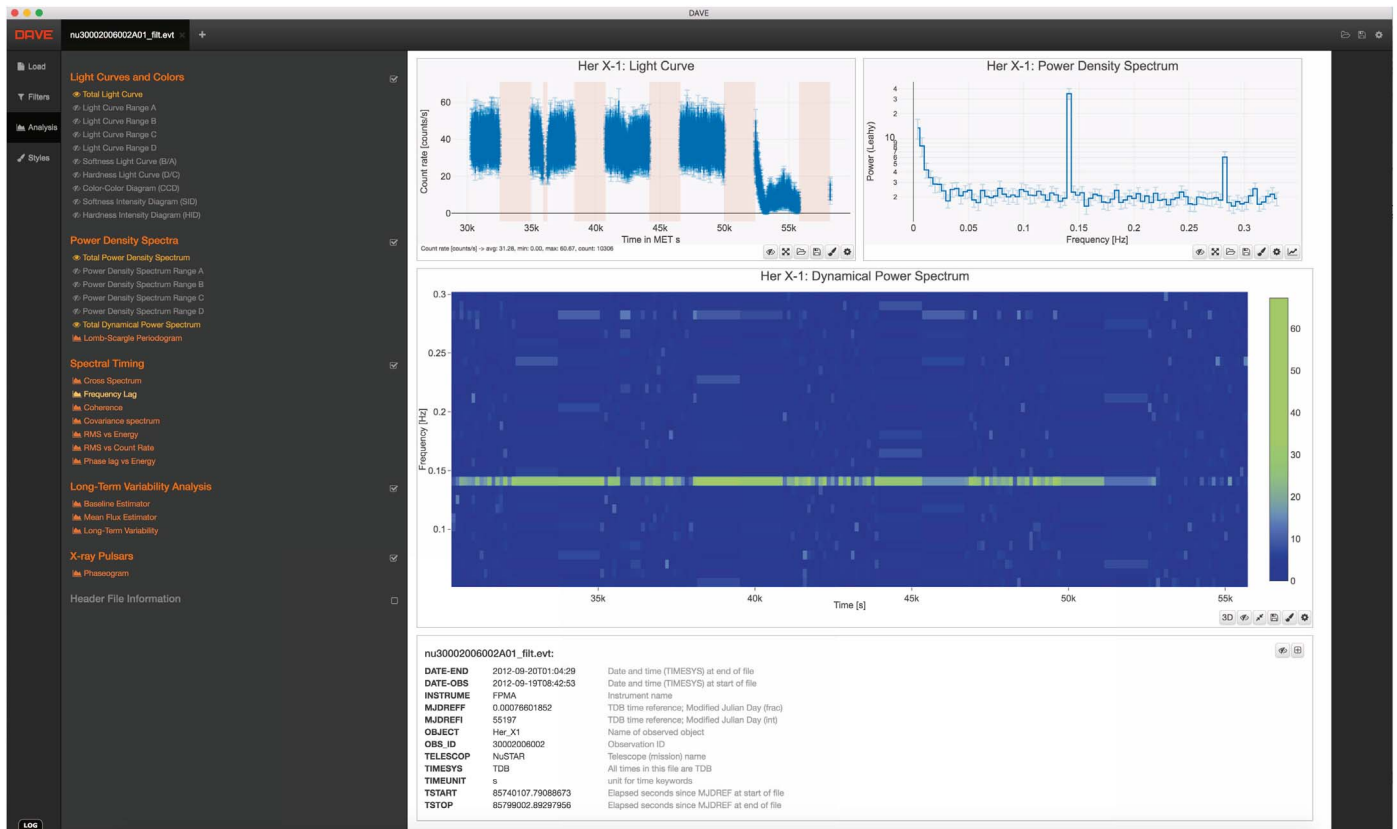


Figure 7. Example of the DAVE graphical user interface for the Her X-1 pulsar data observed with *NuSTAR*. In the top left, we show the last 30 ks of the pulsar light curve with the GTIs clearly marked. In the top right, we plot the averaged power spectrum generated from 109 segments of 256 s duration with a binned time resolution of 1.5 s. In the middle, we present the dynamic power spectrum generated from the same 256 s segments that generated the top-right averaged power spectrum. Below, header metadata are shown for reference. On the left, the menu presents a range of options of figures to plot and compare, including spectral-timing capabilities. All figures are interactive, including panning and zooming, including interactive choices of data selection.

implementations of higher-order Fourier products such as bispectra, biphase, and bicoherence exist (Maccarone & Coppi 2002; Maccarone & Schnittman 2005; Maccarone 2013), they require additional extensions to be useful for X-ray spectral timing. New key features in the next version of *stingray*, based on an existing reference implementation of covariance spectra (Wilkinson & Uttley 2009), will include lag-energy spectra (Vaughan et al. 1994), rms-energy spectra (Revnivtsev et al. 1999), and excess variance spectra (Vaughan et al. 2003). In addition, while at the moment there is rudimentary functionality to build spectral-timing products, it is currently not possible to seamlessly work with these products using *stingray*, because *stingray* currently has no native capability for energy-spectral modeling. Instead, they would have to be exported (e.g., saved to disk) and then loaded into another software package, significantly disrupting workflows and pipeline development. In order to streamline this process, we aim to connect *stingray* with existing packages for modeling X-ray spectra. Here, it will be necessary to connect *stingray* with the extensive suite of physical models implemented in XSPEC, as well as existing spectral fitting codes implemented in *Python*, most notably the open-source package *Sherpa* (Burke et al. 2018).

Data rates from current and future X-ray instruments are increasing at a precipitous rate, pushing memory and processing requirements for even simple tasks like fast Fourier transforms of data observed, for example, with *NICER* and

Astrosat into a regime that is difficult with standard desktop computing architectures. Therefore, the other strong emphasis for the second version of *stingray* will be code and algorithm optimization. Where possible, we will replace existing implementations with high-performance equivalents that take advantage of recent developments in computing (such as GPU-enabled computations and multicore batch processing) and optimize and streamline existing code to minimize computational overhead and memory usage of the classes and functions implemented within *stingray*.

Typical X-ray timing observations—long, deep stares at single objects with high time resolution—are particularly well suited for Fourier-based methods, owing to their relatively regular time sampling and observation duration much longer than the physical timescales of interest encoded in the light curves. While some optical telescopes, most notably *Kepler* and the *Transiting Exoplanet Survey Satellite (TESS)* (Ricker et al. 2015), employ similar modes of observation that make the current suite of methods in *stingray* transferable to data from these instruments, many current and future survey instruments such as the *Zwicky Transient Facility (ZTF)* (Bellm et al. 2019; Graham et al. 2019) or the *Large Synoptic Survey Telescope (LSST)* (Ivezic et al. 2019) will provide the community with long-baseline light curves that are irregularly sampled. A range of methods have been developed for time series analysis of these light curves in the time domain, such as *CARMA* (see, e.g., Kelly et al. 2014; Foreman-Mackey et al. 2017) and

ARIMA (e.g., Feigelson et al. 2018) processes. Developing spectral timing methods for irregularly sampled light curves is a major future challenge for the field and a high-priority, long-term goal for *stingray*.

We thank Astro Hack Week for providing the venue that started this project and the Lorentz Center workshop “The X-ray Spectral-Timing Revolution” (2016 February) that started the collaboration, as well as the Google Summer of Code Program for funding five coauthors (O.M., U.M.K., H.M., H.R., S.S.) who implemented a large fraction of the various library components over three summers. We thank Jim Davenport for providing help with the *Kepler* data. We are grateful to Matteo Guanazzi and Erik Kuulkers for their advice and suggestions for implementing DAVE, as well as the anonymous referee and the AAS statistics editor for their helpful comments on the manuscript. D.H. acknowledges support from the DIRAC Institute in the Department of Astronomy at the University of Washington. The DIRAC Institute is supported through generous gifts from the Charles and Lisa Simonyi Fund for Arts and Sciences and the Washington Research Foundation. M.B. is supported in part by the Italian Space Agency through agreement ASI-INAF n.2017-12-H.0 and ASI-INFN agreement n.2017-13-H.0. A.L.S. is supported by an NSF Astronomy and Astrophysics Postdoctoral Fellowship under award AST-1801792. E.M.R. acknowledges support from Conselho Nacional de Desenvolvimento Científico e Tecnológico (CNPq—Brazil).

ORCID iDs

Daniela Huppenkothen  <https://orcid.org/0000-0002-1169-7486>

Matteo Bachetti  <https://orcid.org/0000-0002-4576-9337>

Abigail L. Stevens  <https://orcid.org/0000-0002-5041-3079>

Evandro Martinez Ribeiro  <https://orcid.org/0000-0002-2985-6369>

References

- Abdo, A. A., Ajello, M., Allafort, A., et al. 2013, *ApJS*, **208**, 17
- Akaike, H. 1974, *ITAC*, **19**, 716
- Anderson, S. F., Wachter, S., Margon, B., et al. 1994, *ApJ*, **436**, 319
- Arnaud, K. A. 1996, in ASP Conf. Ser. 101, *Astronomical Data Analysis Software and Systems V*, ed. G. H. Jacoby & J. Barnes (San Francisco, CA: ASP), 17
- Bachetti, M. 2015, *MaLTPyNT: Quick Look Timing Analysis for NuSTAR Data*, *Astrophysics Source Code Library*, ascl:1502.021
- Bachetti, M., & Huppenkothen, D. 2018, *ApJL*, **853**, L21
- Bachetti, M., Harrison, F. A., Cook, R., et al. 2015, *ApJ*, **800**, 109
- Bahcall, J. N., & Bahcall, N. A. 1972, *ApJL*, **178**, L1
- Barentsen, G., Hedges, C., Vincius, Z., et al. 2019, *KeplerGO/lightkurve: Lightkurve v1.0b30*, Zenodo, doi:10.5281/zenodo.2611871
- Barret, D., Lam Trong, T., den Herder, J.-W., et al. 2018, *Proc. SPIE*, **10699**, 106991G
- Bellm, E. C., Kulkarni, S. R., Graham, M. J., et al. 2019, *PASP*, **131**, 018002
- Belloni, T., & Hasinger, G. 1990, *A&A*, **227**, L33
- Belloni, T., Homan, J., Casella, P., et al. 2005, *A&A*, **440**, 207
- Beloborodov, A. M., Stern, B. E., & Svensson, R. 2000, *ApJ*, **535**, 158
- Bentz, M. C. 2016, in *AGN Reverberation Mapping*, ed. H. M. J. Boffin et al. (Cham: Springer International), 249
- Beri, A., Tetarenko, B. E., Bahramian, A., et al. 2019, *MNRAS*, **485**, 3064
- Blandford, R. D., & McKee, C. F. 1982, *ApJ*, **255**, 419
- Borucki, W. J., Koch, D., Basri, G., et al. 2010, *Sci*, **327**, 977
- Bradt, H. V., Rothschild, R. E., & Swank, J. H. 1993, *A&AS*, **97**, 355
- Brown, T. M., Latham, D. W., Everett, M. E., & Esquerdo, G. A. 2011, *AJ*, **142**, 112
- Brumback, M. C., Hickox, R. C., Bachetti, M., et al. 2018, *ApJL*, **861**, L7
- Buccheri, R., Bennett, K., Bignami, G. F., et al. 1983, *A&A*, **128**, 245
- Burke, D., Laurino, O., dtnguyen2, et al. 2018, *sherpa/sherpa: Sherpa 4.10.1*, Zenodo, doi:10.5281/zenodo.1463962
- Carpenter, B., Gelman, A., Hoffman, M., et al. 2017, *J. Stat. Software*, **76**, 1
- Cash, W. 1979, *ApJ*, **228**, 939
- Charbonneau, D., Brown, T. M., Latham, D. W., & Mayor, M. 2000, *ApJL*, **529**, L45
- Cheng, F. H., Vrtilik, S. D., & Raymond, J. C. 1995, *ApJ*, **452**, 825
- Coughlin, J. L., Mullally, F., Thompson, S. E., et al. 2016, *ApJS*, **224**, 12
- Davidson, A., Henry, J. P., Middleditch, J., & Smith, H. E. 1972, *ApJL*, **177**, L97
- Di Mauro, M. P. 2016, *Frontier Research in Astrophysics II (Trieste: SISSA)*, 29
- Feigelson, E. D., Babu, G. J., & Caceres, G. A. 2018, *FrP*, **6**, 80
- Forbrich, J., Reid, M. J., Menten, K. M., et al. 2017, *ApJ*, **844**, 109
- Foreman-Mackey, D. 2016, *JOSS*, **1**, 24
- Foreman-Mackey, D., Agol, E., Ambikasaran, S., & Angus, R. 2017, *AJ*, **154**, 220
- Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, **125**, 306
- Forman, W., Jones, C. A., & Liller, W. 1972, *ApJL*, **177**, L103
- Fürst, F., Grefenstette, B. W., Stauber, R., et al. 2013, *ApJ*, **779**, 69
- Gelman, A., & Rubin, D. B. 1992, *StaSc*, **7**, 457
- Gendreau, K. C., Arzoumanian, Z., Adkins, P. W., et al. 2016, *Proc. SPIE*, **9905**, 99051H
- Giacconi, R., Rabsky, H., Kellogg, E., et al. 1973, *ApJ*, **184**, 227
- Girod, B., Rubenstein, R., & Stenger, A. 2001, *Signals and Systems* (New York: Wiley)
- Graham, M. J., Kulkarni, S. R., Bellm, E. C., et al. 2019, *PASP*, **131**, 078001
- Groth, E. J. 1975, *ApJS*, **29**, 285
- Guidorzi, C., Dichiaro, S., & Amati, L. 2016, *A&A*, **589**, A98
- Harrison, F. A., Craig, W. W., Christensen, F. E., et al. 2013, *ApJ*, **770**, 103
- Heida, M., Jonker, P. G., Torres, M. A. P., & Chiavassa, A. 2017, *ApJ*, **846**, 132
- Henry, G. W., Marcy, G. W., Butler, R. P., & Vogt, S. S. 2000, *ApJL*, **529**, L41
- Hewish, A., Bell, S. J., Pilkington, J. D. H., Scott, P. F., & Collins, R. A. 1968, *Natur*, **217**, 709
- Houck, J. C., & Denicola, L. A. 2000, in ASP Conf. Ser. 216, *Astronomical Data Analysis Software and Systems IX*, ed. N. Manset, C. Veillet, & D. Crabtree (San Francisco, CA: ASP), 591
- Hunter, J. D. 2007, *CSE*, **9**, 90
- Huppenkothen, D., & Bachetti, M. 2018, *ApJS*, **236**, 13
- Huppenkothen, D., Watts, A. L., Uttley, P., et al. 2013, *ApJ*, **768**, 87
- Hynes, R. I., Steeghs, D., Casares, J., Charles, P. A., & O’Brien, K. 2003, *ApJL*, **583**, L95
- Ignia, C. D., & Leahy, D. A. 2011, *MNRAS*, **418**, 2283
- Inglis, A. R., Ireland, J., Dennis, B. R., Hayes, L., & Gallagher, P. 2016, *ApJ*, **833**, 284
- Ingram, A., & van der Klis, M. 2015, *MNRAS*, **446**, 3516
- Israel, G. L., Belloni, T., Stella, L., et al. 2005, *ApJL*, **628**, L53
- Ivezic, Ž., Kahn, S. M., Tyson, J. A., et al. 2019, *ApJ*, **873**, 111
- Jahoda, K., Swank, J. H., Giles, A. B., et al. 1996, *Proc. SPIE*, **2808**, 59
- Jaiswal, G. K., Naik, S., Gupta, S., Chenevez, J., & Epili, P. 2018, *MNRAS*, **478**, 448
- Jansen, F., Lumb, D., Altieri, B., et al. 2001, *A&A*, **365**, L1
- Jones, E., Oliphant, T., Peterson, P., et al. 2001, *SciPy: Open Source Scientific Tools for Python*, <http://www.scipy.org/>
- Jupyter, P., Bussonnier, M., Forde, J., et al. 2018, in *Proc. 17th Python in Science Conf.*, ed. F. Akici, 113, 10.1007/978-3-319-39739-9_13
- Kelly, B. C., Becker, A. C., Sobolewska, M., Siemiginowska, A., & Uttley, P. 2014, *ApJ*, **788**, 33
- Kennedy, M. R., Clark, C. J., Voisin, G., & Breton, R. P. 2018, *MNRAS*, **477**, 1120
- Lam, S. K., Pitrou, A., & Seibert, S. 2015, in *Proc. Second Workshop on the LLVM Compiler Infrastructure in HPC, LLVM’15* (New York: ACM), 7
- Leahy, D. A. 1987, *A&A*, **180**, 275
- Leahy, D. A., & Abdallah, M. H. 2014, *ApJ*, **793**, 79
- Leahy, D. A., Darbro, W., Elsner, R. F., et al. 1983a, *ApJ*, **266**, 160
- Leahy, D. A., Elsner, R. F., & Weisskopf, M. C. 1983b, *ApJ*, **272**, 256
- Lorimer, D. R. 2008, *LRR*, **11**, 8
- Ludlam, R. M., Miller, J. M., & Cackett, E. M. 2015, *ApJ*, **806**, 262
- Maccarone, T. J. 2013, *MNRAS*, **435**, 3547
- Maccarone, T. J., & Coppi, P. S. 2002, *MNRAS*, **336**, 817
- Maccarone, T. J., & Schnittman, J. D. 2005, *MNRAS*, **357**, 12

- Miyamoto, S., Kitamoto, S., Iga, S., Negoro, H., & Terada, K. 1992, *ApJL*, **391**, L21
- Muñoz-Darias, T., Casares, J., & Martínez-Pais, I. G. 2008, *MNRAS*, **385**, 2205
- Nowak, M. A., Dove, J. B., Vaughan, B. A., Wilms, J., & Begelman, M. C. 1999, *NuPhS*, **69**, 302
- Papitto, A., Ambrosino, F., Stella, L., et al. 2019, arXiv:1904.10433
- Pike, S. N., Harrison, F. A., Bachetti, M., et al. 2019, *ApJ*, **875**, 144
- Protassov, R., van Dyk, D. A., Connors, A., Kashyap, V. L., & Siemiginowska, A. 2002, *ApJ*, **571**, 545
- Ray, P. S., Arzoumanian, Z., Brandt, S., et al. 2018, *Proc. SPIE*, **10699**, 1069919
- Revnivtsev, M., Gilfanov, M., & Churazov, E. 1999, *A&A*, **347**, L23
- Reynolds, A. P., Quaintrell, H., Still, M. D., et al. 1997, *MNRAS*, **288**, 43
- Ricker, G. R., Winn, J. N., Vanderspek, R., et al. 2015, *JATIS*, **1**, 014003
- Scaringi, S., Maccarone, T. J., D'Angelo, C., Knigge, C., & Groot, P. J. 2017, *Natur*, **552**, 210
- Scaringi, S., Maccarone, T. J., Körding, E., et al. 2015, *SciA*, **1**, e1500686
- Schwarz, G. 1978, *AnSta*, **6**, 461
- Scott, D. M., & Leahy, D. A. 1999, *ApJ*, **510**, 974
- Singh, K. P., Tandon, S. N., Agrawal, P. C., et al. 2014, *Proc. SPIE*, **9144**, 91441S
- Smith, K. L., Mushotzky, R. F., Boyd, P. T., et al. 2018, *ApJ*, **857**, 141
- Staubert, R., Klochkov, D., Fürst, F., et al. 2017, *A&A*, **606**, L13
- Stevens, A. L., & Uttley, P. 2016, *MNRAS*, **460**, 2796
- Strohmayer, T. E., & Watts, A. L. 2005, *ApJL*, **632**, L111
- Tananbaum, H., Gursky, H., Kellogg, E. M., et al. 1972, *ApJL*, **174**, L143
- Taylor, J. H. 1992, *RSPTA*, **341**, 117
- Tetarenko, A. J., Casella, P., Miller-Jones, J. C. A., et al. 2019, *MNRAS*, **484**, 2987
- The Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, **156**, 123
- Timmer, J., & Koenig, M. 1995, *A&A*, **300**, 707
- Uttley, P., Cackett, E. M., Fabian, A. C., Kara, E., & Wilkins, D. R. 2014, *A&ARv*, **22**, 72
- Uttley, P., & McHardy, I. M. 2001, *MNRAS*, **323**, L26
- van der Klis, M. 1989, in *Timing Neutron Stars*, NATO Advanced Science Institutes (ASI) Ser. C, Vol. 262, ed. H. Ögelman & E. P. J. van den Heuvel (New York: Kluwer Academic), 27
- van der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, *CSE*, **13**, 22
- Vaughan, B., van der Klis, M., Lewin, W. H. G., et al. 1994, *ApJ*, **421**, 738
- Vaughan, B. A., & Nowak, M. A. 1997, *ApJL*, **474**, L43
- Vaughan, S. 2010, *MNRAS*, **402**, 307
- Vaughan, S., Edelson, R., Warwick, R. S., & Uttley, P. 2003, *MNRAS*, **345**, 1271
- Walton, D. J., Bachetti, M., Fürst, F., et al. 2018, *ApJL*, **857**, L3
- Wang-Ji, J., García, J. A., Steiner, J. F., et al. 2018, *ApJ*, **855**, 61
- Watts, A. L., & Strohmayer, T. E. 2006, *ApJL*, **637**, L117
- Wilkinson, T., & Uttley, P. 2009, *MNRAS*, **397**, 666
- Yamaoka, K., Sugizaki, M., Nakahira, S., et al. 2010, *ATel*, **2380**
- Zdziarski, A. A., Poutanen, J., Mikolajewska, J., et al. 1998, *MNRAS*, **301**, 435
- Zhang, S. N., Feroci, M., Santangelo, A., et al. 2016, *Proc. SPIE*, **9905**, 99051Q