

Self-lensing flares from black hole binaries. V. Systematic searches in LSST

Kevin Park^{1,2,3,*} Zoltán Haiman^{1,4,2,3} Chengcheng Xin^{4,5} Tzuken Shen⁵ Ashley Villar^{5,6} and Jordy Davelaar^{7,†}

¹*Institute of Science and Technology Austria (ISTA), Am Campus 1, Klosterneuburg 3400, Austria*

²*Department of Physics, Columbia University, 550 W 120th Street New York, New York 10027, USA*

³*Astrophysics Laboratory, Columbia University, 550 W 120th Street New York, New York 10027, USA*

⁴*Department of Astronomy, Columbia University, New York, New York 10027, USA*

⁵*Center for Astrophysics—Harvard and Smithsonian, 60 Garden Street,
Cambridge, Massachusetts 02138, USA*

⁶*The NSF AI Institute for Artificial Intelligence and Fundamental Interactions, 77 Massachusetts Avenue,
26-505 Cambridge, Massachusetts 02139-4307, USA*

⁷*Department of Astrophysical Sciences, Peyton Hall, Princeton University,
Princeton, New Jersey 08544, USA*



(Received 9 December 2025; accepted 4 February 2026; published 25 February 2026)

The Vera C. Rubin Observatory has now seen first light, and over a 10 year duration, Legacy Survey of Space and Time (LSST) is projected to catalog tens of millions of quasars, many of which are expected to be associated with subparsec supermassive black hole binaries (SMBHBs). Out of these SMBHBs, up to thousands of relatively massive binary-quasars are expected to exhibit gravitational self-lensing flares (SLFs) that last for at least 20–30 days. We assess the effectiveness of the Lomb-Scargle (LS) periodogram and matched filters (MFs) as methods for systematic searches for these binaries, using toy-models of hydrodynamical, Doppler, and self-lensing variability from equal-mass, eccentric SMBHBs. We inject SLFs into random realizations of damped random walk (DRW) lightcurves, representing stochastic quasar variability, and compute the LS periodogram with and without the SLF. We find that periodograms of SLF + DRW light-curves do not have maximum peak heights that could not arise from DRW-only periodograms. On the other hand, the matched filter signal-to-noise ratio (SNR) can distinguish SLFs from noise even with LSST-like cadences and DRW noise. Furthermore, we develop a three-step procedure with matched filters, which can also recover injected binary parameters from these light-curves. We expect this method to be computationally efficient enough to be applicable to millions of quasar light-curves in LSST.

DOI: [10.1103/p328-62sl](https://doi.org/10.1103/p328-62sl)

I. INTRODUCTION

Supermassive black holes (SMBHs) with masses between 10^6 – $9M_\odot$ have been observed in the cores of most nearby galaxies [1]. If galaxies primarily grow through mergers, then the two central SMBHs should settle at the center of a merged galaxy driven by dynamical interactions with their environment. Although the fraction of systems that harden to close separations is still an open question, gravitationally bound SMBH binaries (SMBHBs) should be common in galactic nuclei [2]. Hydrodynamical simulations of the electromagnetic (EM) emission from a SMBHB accreting from circumbinary gas show that

this emission can be bright and also periodic long before merger [3–6].

There are dozens of observations of dual active galactic nuclei (AGN)—two active SMBHs in a single galaxy separated at kpc scales [7]. However, there are no directly imaged sub-pc SMBHBs due to their close separation.

Currently, one of the most prominent methods to find sub-pc SMBHBs is to find apparent periodicity in quasar/AGN light-curves [8,9]. Systematic searches for periodicity in AGN light-curves have been performed for large time-domain optical surveys including the Catalina Real-time Transient Survey (CRTS) [10], the Zwicky Transient Facility (ZTF) [11], the Palomar Transient Factory (PTF) [12], and GAIA [13], identifying hundreds of SMBHB candidates (but see also [14–16]).

In addition, recently, pulsar timing arrays (PTAs) have detected a stochastic gravitational wave background (GWB) in the nHz bands, which is consistent with a cosmological population of coalescing SMBHBs [17–20]. Targeted continuous wave (CW) searches with PTAs have further identified one AGN (dubbed Rohan) which shows modest

*Contact author: ksp2136@columbia.edu

†NASA Hubble Fellowship Program, Einstein Fellow.

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International license](https://creativecommons.org/licenses/by/4.0/). Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

Bayes-factor support for a CW signal and has sustained quasiperiodicity in its optical light-curve for ~ 20 years [21]. Gravitational waves from SMBHBs will also be directly detectable by the Laser Interferometer Space Antenna (LISA) [22], planned for launch in the 2030s. Thus, at the current stage, compiling a catalog of periodic quasars can later enable combining EM and GW signals from the same sources [23], which would allow for new tests of general relativity, high energy physics, and cosmology [24]. The Vera C. Rubin Observatory’s Legacy Survey of Space and Time (LSST), which has now seen first light in 2025, is expected to contain 20–100 million quasars [25] over a 10 year survey duration and can contribute significantly to this catalog of periodic AGN.

Most systematic searches for SMBHB candidates in the time-domain have relied on Lomb-Scargle (LS) periodograms or similar sinusoid-fitting methods, assuming that the periodicity of the emission roughly follows the orbital period of the SMBHB and is sinusoidal. However, it is known that quasars exhibit stochastic red-noise variability which is well described by a “damped random walk” (DRW) model [26] or a “damped harmonic oscillator” model [27]. This quasar red-noise variability can mimic sinusoidlike periodicities for a few periods [28,29], meaning many SMBHB candidates, without additional evidence, can exhibit spurious periodicity by chance.

A SMBHB signature that would be more difficult to mimic by red-noise is a “self-lensing flare” (SLF). General relativistic ray-tracing models [30,31] (hereafter Papers I, II) show that self-lensing flares occur when binaries are inclined nearly edge-on and the two SMBHBs are aligned to the line of sight within an Einstein radius. Then the emission from the “source” SMBH passing behind is gravitationally lensed and magnified by the “lens” SMBH in front, leading to a symmetric flare in the light-curve occurring twice per orbit, assuming both BHs are actively accreting. [32] (hereafter paper III) shows using hydrodynamical simulations that these flares exist in light-curves even in the presence of circumbinary gas and [33] (hereafter paper IV) estimate that up to thousands of self-lensing flares at least 20–30 days long, with many more narrower flares, can be expected in LSST. This paper (paper V) considers the question of how to recover these flares from noisy and sparse light-curves.

There are currently a handful of self-lensing binary candidates reported. An AGN in the Kepler catalog, dubbed Spikey, also had a light-curve which was well-fit with a relativistic Doppler boost + SLF model from an eccentric binary. The hypothesized self-lensing flare occurred as a symmetric $\sim 10\%$ magnification at the expected orbital phase—when the Doppler modulation is at a minimum [34]. Radio observations of Spikey have shown evidence for a wobbling jet [35], which would be consistent with a SMBHB with one BH having an active and precessing jet, but a second flare was not found. Meanwhile, follow-up observations of Spikey with the *Chandra Space Telescope*

at a predicted flare (2020 March-May) [36] were inconclusive, as the intrinsic x-ray variability was larger than the amplitude of the self-lensing flare and the relativistic Doppler boost.

In addition, [37] report a symmetric flare in the AGN NGC 1566, which lasted for roughly 155 days and is well-fit by a microlensing curve with a lens mass of $5 \times 10^5 M_\odot$. The primary SMBH mass was estimated to be $5 \times 10^6 M_\odot$ from optical broad emission lines and the SMBHB has an inferred orbital period of roughly 4000 years.

While there are currently only a few uncertain candidates, [38] have also showed that several hundred self-lensing flares lasting 30 days or longer are expected in LSST [39] using the Illustris simulation [40]. This is consistent with paper IV, which was based on the observed quasar luminosity function (QLF) under the assumption that most AGN are SMBH binaries. Notably, this assumption is supported by the fact that the population of quasars from the QLF yields a gravitational wave background consistent with the recent PTA detections [41]. Thus, a systematic method tailored to finding periodic flares in large datasets such as those expected soon from LSST AGN is warranted.

In addition, for a fraction of self-lensing binaries, the angular separation at alignment can be comparable to the size of the source BH’s shadow ($\sqrt{27}GM_{\text{bh}}/c^2$), which leads to a “dip” in luminosity near the peak of a flare when the BH shadow is lensed. Papers I, II and [42] show that the height and width of the dips can vary depending on the SMBHB parameters and that these dips, along with the self-lensing flares, also exist in light-curves even in the presence of circumbinary gas (paper III). Paper IV finds that dozens of binaries will be inclined enough to exhibit dips in LSST. Finally, in edge-on ultracompact binaries, the binary’s orbital period will evolve due to the emission of gravitational waves and thus a series of “chirping” self-lensing flares could be observable in an even smaller subset of sources [43,44].

In this paper, we develop statistical and computational tools for the search for self-lensing flares in LSST. First, we use the PYTHON package BINLITE [45] for toy-models of light-curves of eccentric, equal-mass SMBHBs. BINLITE incorporates accretion variability from hydrodynamical simulations of binaries and circumbinary disks, relativistic Doppler modulation, and self-lensing flares to generate synthetic light-curves. We also incorporate stochastic quasar DRW variability and downsample our light-curves at LSST-like cadences. In our initial attempt, we use standard LS periodograms, the most commonly used method for SMBHB searches, to recover the periodicity of the self-lensing flares. We show that even in ideal conditions (no noise, ideal sampling), the LS periodogram peak heights of periodic self-lensing flares cannot be robustly distinguished from that of DRW. This result is consistent with the fact that LS periodograms are known to

struggle with detecting nonsinusoidal (sawtooth shape) periodicity [46]. Alternatively, we use matched filters to identify periodicities and to recover binary parameters from self-lensing flares even in the presence of DRW noise and LSST-like observation gaps. In order to avoid searching a wide parameter space, we develop a three-step procedure to zoom-in and sample the vicinity of the true injected binary parameters.

The rest of this paper is organized as follows. In Sec. II, we describe our methodology, including a summary of BINLITE, simulations of DRW variability, and simulations of LSST noise and sampling. In Sec. III, we present our attempt to use LS periodograms as well as the parameter recovery capability using matched filters. In Sec. IV, we discuss the current limitations and future directions for matched filter approaches using hydrodynamical simulations. In Sec. V, we summarize our main conclusions.

II. METHODS

In this section we outline the methods of generating realistic mock light-curves. We start from our models for binary AGN periodicity and then describe our implementation of DRW noise. Finally, we describe our adopted LSST instrumental noise and irregular sampling. Both DRW and LSST white noise are assumed to be additive, so our final light-curve is

$$d(t) = s_{\text{binlite}}(t) + n_{\text{DRW}}(t) + n_{\text{LSST}}(t), \quad (1)$$

where the BINLITE model is denoted s_{binlite} to emphasize that it is the signal we want to extract from the DRW and LSST noise.

A. Model binary light-curves

To simulate self-lensing flares and accretion variability, we use the PYTHON package BINLITE from [45] (hereafter B24). We provide a brief overview of their methods below, but refer the reader to B24 for more details.

B24 consider a binary characterized by total mass $M = M_1 + M_2$ where the masses of the two components are equal: $q = M_2/M_1 = 1$. The binary orbit has eccentricity $e \in (0, 0.8)$, semimajor axis a , and orbital period T_{orb} . A locally isothermal circumbinary disk (CBD) accretes onto the binary at a user-specified fixed rate of $\dot{M}_{\text{CBD}} = \langle \dot{M}_1(t) + \dot{M}_2(t) \rangle$, and $Q \equiv \langle \dot{M}_2(t) \rangle / \langle \dot{M}_1(t) \rangle$ is defined as the ratio of time-averaged accretion rates between the accretion onto each component calculated from a hydrodynamical simulation. Note that the averaging is performed over dozens of orbits.

An important feature of BINLITE is that despite the two black holes having the same mass, their accretion rates, when averaged over dozens of orbits, differ ($Q \neq 1$). This accounts for the behavior observed in hydrodynamical simulations: the two black holes are surrounded by a

nonaxisymmetric circumbinary disk, which typically precesses on time-scales corresponding to hundreds of orbits. As a result, one black hole out-accretes the other, with their roles switching back-and-forth on this precession time-scale. A typical observation time (i.e., several years to a decade) is much shorter than this timescale, except for rare, ultracompact binaries with periods of at most several days. In practice, this implies that during an LSST observation, an equal-mass binary can be observed with a range of possible long-term average accretion-rate ratios Q . Furthermore, the probability distribution of this average Q depends on eccentricity. In our work, we include the dependence of Q on e , as provided by BINLITE. This assumes, for simplicity, that $Q = Q(e)$ is deterministic, i.e., yields a single Q , roughly corresponding to the value at which the binary spends most of its time. Future analyses should incorporate the full probabilistic nature of Q and allow for a range of possible Q values for each eccentricity.

The above fixes the mean accretion rate ratio. On top of this, the variation of the accretion rates on short (orbital) timescales are included as

$$p(t) = \frac{\dot{M}_1(t)}{\langle \dot{M}_1(t) + \dot{M}_2(t) \rangle}, \quad s(t) = \frac{\dot{M}_2(t)}{\langle \dot{M}_1(t) + \dot{M}_2(t) \rangle}. \quad (2)$$

Here, $\dot{M}_1(t)$ and $\dot{M}_2(t)$ are Fourier reconstructions of accretion rates in 2D hydrodynamical simulations and thus are deterministic and periodic. This reconstruction is valid as a close approximation to the original accretion rates as the power of the Fourier components is concentrated at integer multiples of the orbital frequency, as shown in the periodograms of B24's Fig 1. A comparison, with close match, between the simulation and the Fourier reconstructions can be found in B24's Fig 2.

Next, each BH component of the binary is assumed to have a minidisk characterized by the thin-disk temperature profile $T(r) \propto (M_i \dot{M}_{i,\text{avg}} / r_i^3)^{1/4}$, where $M_i, \dot{M}_{i,\text{avg}} = \langle \dot{M}_i \rangle$ and r_i are the masses, time-averaged accretion rates and the distance to each BH, respectively. The CBD has the same temperature distribution but is assumed to be centered around a BH with total mass of $M = M_1 + M_2$ located at the center of mass of the system. Each minidisk is assumed to extend from its innermost stable circular orbit (ISCO) $r_{\text{ISCO},i}(M_i) = 6GM_i/c^2$ to the tidal truncation radius $r_{\text{tidal},i} = 0.27a$, where $a = (GMT^2/4\pi^2)^{1/3}$. The CBD is assumed to extend from $2a$ to $100a$. Finally, assuming the two minidisks and the CBD are optically thick and their spectra can be approximated as blackbodies, the average fluxes from each disk are identically calculated as:

$$F_{\nu,i} = \frac{2\pi \cos I}{d^2} \int B_\nu[T_i(r_i)] r_i dr_i, \quad (3)$$

for a binary located at luminosity distance $d = d(z)$, where z is the cosmological redshift, B_ν is the Planck function and




FIG. 1. Left panel: normalized accretion rates that were used to generate the two light-curves in the right panel. The dashed and dash-dotted curves indicate the accretion from the primary and the secondary, respectively. Note that in this example, the ratio of the time-averaged accretion rates $\langle \dot{M}_2(t) \rangle / \langle \dot{M}_1(t) \rangle$ is different for $e = 0.15$ vs $e = 0.45$, meaning that despite $M_1 = M_2$ for BINLITE, the relative brightness of the two disks $F_{\nu,1}/F_{\nu,2}$ can be different for different eccentricities (See Sec. II A for discussion). Right panel: Examples of BINLITE-simulated optical (r -band) light-curves from equal-mass binaries with total mass $M = 10^9 M_\odot$, orbital period $T_{orb} = 2$ yr at redshift $z = 1$, accreting at 30% of its Eddington limit and inclined at $I = 88^\circ$. We visualize the effects of varying the orbital eccentricity for two values: $e = 0.15$ and $e = 0.45$. The dashed curves indicate the light-curves assuming no self-lensing magnification or Doppler boost effects (i.e., $\mathcal{M}_1 = \mathcal{M}_2 = 1$ in Eq. (4)). We also indicate the amplitude of two periodic binary signatures: self-lensing (A_{SLF}) and accretion variability (A_{hydro}) for reference.

I is the binary inclination where we use the convention that 90° is edge-on and 0° is face-on. The factor of $\cos I$ indicates that we assume isotropic emission and that only the perpendicular component of the flux from the disk will

be observable. The binary orbital plane and the two minidisks are assumed to be coplanar, given that the alignment of the accretion planes of SMBHs with the orbital axis is expected to be efficient [47].




FIG. 2. Histograms of the maximum peak powers in 10,000 DRW (orange) realizations and in corresponding SLF + DRW (blue, green, red) light-curves. The height of the self-lensing flare A_{SLF} was varied manually so that it was 1, 2.5, 5 (blue, green, red respectively) times larger than the DRW amplitude A_{DRW} . The maximum peak height value among the 10,000 pure DRW cases is marked by a vertical orange line. Out of the SLF-included light-curves, none have peak heights greater than this maximum. Counterintuitively, as A_{SLF}/A_{DRW} is increased, the LS power decreases (see Fig. 3 for a visual explanation of this effect).

Finally, the total observed flux from the system is

$$F_\nu = F_{\nu,1}\mathcal{M}_1(t)p(t) + F_{\nu,2}\mathcal{M}_2(t)s(t) + F_{\nu,\text{CBD}}, \quad (4)$$

where $F_{\nu,1}, F_{\nu,2}, F_{\nu,\text{CBD}}$ are the average fluxes from the primary, secondary, and CBD respectively and $\mathcal{M}_i(t)$ is the total modulation from self-lensing magnification and Doppler boost, and we have explicitly noted the time dependence. The self-lensing magnification of the flux is calculated assuming the point source approximation

$$\mathcal{M}(u) = \frac{u^2 + 2}{u\sqrt{u^2 + 4}}, \quad (5)$$

where u is the angular separation of the binary in units of the Einstein radius. All of the features introduced above in this section, with the final product being a total observed flux time series, can be obtained by generating a `BinaryAlphaDisk` object and using the `periodic_flux_series_from_bad` function. After obtaining the flux series (given in cgs units) from `BINLITE` we convert to AB magnitudes. In Sec. II C we take this flux series F_ν and simulate it with LSST-like photometric noise and observation gaps before converting to magnitudes. The statistical analysis using LS periodograms and matched filters was done in magnitudes. We used magnitudes to match the survey data products, though the analysis could equivalently be performed in flux. In particular, for weak flares, i.e., $A_{\text{SLF}}/A_{\text{DRW}}$, we believe working in flux units could improve performance.

B. Intrinsic quasar variability

Meanwhile, we model the intrinsic optical variability of quasars as a damped random walk (DRW) [26], which has a typical characteristic timescale τ_{DRW} on the order of a few hundred days and can mimic periodic signals for a typical survey length, giving the appearance of a binary [28]. The DRW is characterized by an exponential covariance matrix:

$$S_{ij} = A_{\text{DRW}}^2 \exp[-|t_i - t_j|/\tau_{\text{DRW}}]. \quad (6)$$

Here A_{DRW} (typically denoted as σ_{DRW} in other works, but denoted differently in this work for direct comparison to A_{SLF}) is the long-term ($\Delta t \gg \tau_{\text{DRW}}$) DRW variance in magnitude. τ_{DRW} is a characteristic observer-frame damping timescale and t_i, t_j are two observation times of the light-curve. We used the PYTHON package `astroML`'s `generate_damped_RW` function to generate mock DRW light-curves.

C. LSST simulation

The simulated SMBH binary light-curves are then downsampled on the cadence of the Rubin Observatory LSST `baseline_v3.3_10 yrs` OpSim run [48]. LSST-like instrumental noise is calculated using the semianalytic model outlined in the publicly released `LSE-40`, LSST

Photon Rates and SNR Calculations, v1.2 document [49], of which we give a brief overview.

Given a source flux $F_\nu(\lambda)$, the flux transmitted through the atmosphere to the telescope pupil is given by

$$F_\nu^{\text{pupil}}(\lambda) = F_\nu(\lambda)S^{\text{atm}}(\lambda), \quad (7)$$

where S^{atm} is the probability that a photon with wavelength λ will reach the telescope. Then given F_ν^{pupil} , the number C_b of recorded analog-to-digital units (ADUs) in a given band b is given by

$$C_b = C \int_0^\infty F_\nu^{\text{pupil}}(\lambda)S_b^{\text{sys}}(\lambda)/\lambda d\lambda \\ \approx CF_\nu(\lambda_{\text{eff},b}) \int_0^\infty S^{\text{atm}}(\lambda)S_b^{\text{sys}}(\lambda)/\lambda d\lambda, \quad (8)$$

where the effective wavelength in each band b is calculated as $\lambda_{\text{eff},b} = \int \lambda S_b^{\text{sys}}(\lambda) d\lambda / \int S_b^{\text{sys}}(\lambda) d\lambda$, $S_b^{\text{sys}}(\lambda)$ is the probability that a photon at the telescope pupil will be converted into a ADU count, and the conversion from flux to ADU count is given by $C = A\Delta t_{\text{exp}}/gh$, where $A = \pi(321.15)^2 \text{ cm}^2$ is the mirror area, $\Delta t_{\text{exp}} = 30\text{s}$ is the average exposure time, $h = 6.626 \times 10^{-27} \text{ erg} \times \text{s}$ is Planck's constant, and $g = 2.2$ is the number of photoelectrons per ADU count. We incorporate the sky background flux density $B_\nu[\text{Jy/arcsec}^2]$ provided in the OpSim data by converting the flux density to a ADU count per pixel

$$B_b = CB_\nu(\lambda_{\text{eff},b}) \int_0^\infty S^{\text{atm}}(\lambda)S_b^{\text{sys}}(\lambda)/\lambda d\lambda \cdot \left(\frac{0.2 \text{ arcsec}}{\text{pixelScale}}\right)^2. \quad (9)$$

In LSST, pointlike sources will be optimally collected over $n_{\text{eff}} = 2.266(\text{FWHM}/\text{pixelScale})^2$ pixels, where we adopt the simulated data in OpSim for atmospheric seeing (FWHM).

Finally, we take the instrumental noise per pixel (dark current, read noise from sensors and electronics) to be $\sigma_{\text{instr}}^2 = 12.7$ electrons per pixel.

Combining all sources of noise, we can calculate the signal-to-noise ratio (SNR) of each data point in our simulation as

$$\text{SNR} = \frac{C_b}{\sqrt{\frac{C_b}{g} + (\frac{B_b}{g} + \sigma_{\text{instr}}^2)n_{\text{eff}}}}. \quad (10)$$

We can then calculate the LSST noise in magnitude units as

$$n_{\text{LSST}}(\text{mag}) = \left| \frac{dm}{dF_\nu} \right| \Delta F_\nu = \frac{2.5}{\ln 10} \frac{\Delta F_\nu}{F_\nu} \sim \frac{1.09}{\text{SNR}}. \quad (11)$$

To summarize, we incorporate LSST-like noise $n_{\text{LSST}}(t)$ and LSST-like cadences from the `baseline_v3.3_10 yrs` OpSim run in our mock light-curves.

III. RESULTS

In order to detect the periodicity in the mock light-curves, we initially attempt using the most common technique in the literature, i.e. LS periodograms. We inject self-lensing flares into a DRW and see if a LS periodogram will produce a peak that is located near the true binary periodicity and if the peak height is high enough that it cannot be mimicked by pure DRWs.

A. LS periodograms: SLF + DRW vs DRW

We begin by introducing a fiducial binary, whose parameters are summarized in Table I. We chose a relatively long orbital period and large total mass, so that the lensing duration will be long but multiple orbits will be present in the light-curve. We also chose an eccentricity which is expected from hydrodynamical simulations [50]. We restrict our analysis to equal-mass binaries, consistent with the simulations of B24. We additionally select a relatively nearby and bright system (in terms of redshift and Eddington ratio) and adopt typical DRW parameters [26]. We simulate the photometric variability of this binary over 10 years to mimic a light-curve expected in LSST.

To analyze the periodicity of self-lensing flares relative to DRW, we follow a simplified version of the analysis used in [12], using a LS periodogram to compare periodogram peak heights between an injected signal with a DRW noise added relative to a pure DRW noise light-curve. We only consider the periodic variability in magnitude units caused by self-lensing, given that LS periodograms are known to struggle with nonsinusoidal signals [46], and manually add in a damped random walk. We neglect any hydrodynamical variability and noise/gaps from LSST. The injected binary,

TABLE I. Binary parameters, DRW parameters, and light-curve settings of the fiducial binary. The resulting self-lensing flares have height of $A_{\text{SLF}} \sim 0.5$ mag.

Fiducial Binary	
Binary Parameter	Injected
T_{orb} (years)	1
e	0.4
I	88°
ω	90°
$q = M_2/M_1$	1
$\log(M/M_\odot)$	9
z	0.5
f_{edd}	0.2
<i>DRW Parameters</i>	
A_{DRW} (mag)	0.1
τ_{DRW} (days)	300
<i>Light-curve Settings</i>	
Total Duration	10 years
Filter	LSST r-band

DRW parameters, and the light-curve settings are shown in Table I.

We identically downsample a 10-year SLF + DRW light-curve and DRW-only light-curve at 10,000 random times, using the same DRW realization, and compute the LS periodograms for each scenario. We repeat this process for the same SLF, but with 10,000 DRW realizations. For each pair (SLF + DRW vs DRW) of LS periodograms, we recorded the frequency and height of the highest peak. We present these distributions in Fig. 2, where we also show the results for different $A_{\text{SLF}}/A_{\text{DRW}}$ ratios, varied by dividing through the fiducial $A_{\text{SLF}} \sim 0.5$ mag manually with a constant factor of 2.0 or 5.0, and keeping $A_{\text{DRW}} = 0.1$ mag constant.

Finally, we checked one criterion from the periodograms, whether the periodogram power of the highest peak was greater than 10,000 DRW realizations, shown in a vertical orange line in Fig. 2. This criterion is generous as it only guarantees false alarm probabilities of 10^{-4} while there are millions of quasars expected in LSST. However, we find that even this criterion is not satisfied, i.e. this analysis is not a reliable method of detecting flares even for this idealized light-curve. In all scenarios the number of SLF-included light-curves satisfying this criterion is 0 out of 10,000, which can also be seen in Fig. 2 by comparing the orange (DRW-only) distribution vs the red, green, blue SLF-included distributions. In Fig. 2, increasing the $A_{\text{SLF}}/A_{\text{DRW}}$ ratio seems to counterintuitively shift the peak-height distributions to lower heights.

To understand these trends, we show a typical light-curve realization in Fig. 3. In the first (leftmost) periodogram, there is an aliasing peak of height ~ 0.25 caused by the DRW realization, which is higher than the injected SLF's tallest peak height of ~ 0.13 in the second periodogram. Note that this DRW periodogram peak height is typical, as can be seen from the orange DRW distribution from Fig. 2. This aliasing peak at $f_{\text{false}} \sim 4 \times 10^{-4}$ [days $^{-1}$] renders the true $f_{\text{orb}}, 2f_{\text{orb}}$ peaks undetectable. For larger $A_{\text{SLF}}/A_{\text{DRW}}$, the false peaks from DRW progressively decrease, and the true $f_{\text{orb}}, 2f_{\text{orb}}$ peaks increase, but do not emerge as the highest peak even for $A_{\text{SLF}}/A_{\text{DRW}} = 5$ (as shown in the rightmost periodogram).

To isolate sampling effects, we also show the Fourier transform of the SLF-only light-curve (second column of Fig 3) in Fig. 4. The Fourier transform is the LS periodogram given regularly spaced and well-sampled data and even in this case the maximum peak height is only ~ 0.13 . Thus we conclude that a naive application of the LS periodogram fails to detect SLFs [46]. This conclusion is very similar to the recent work by [46], who found that sawtooth-shaped hydrodynamical periodicity is difficult for LS periodograms to pick up. In both cases, this failure is caused by the nonsinusoidal pulse shape, which results in the power moving away from the single orbital frequency, and spread over many Fourier modes. In principle, these nonsinusoidal pulse shapes could be modeled an




FIG. 3. light-curves and their LS periodograms representative of the trends in Fig. 2. Each column varies $A_{\text{SLF}}/A_{\text{DRW}}$, where the first column shows pure noise, the second column shows pure SLF signal, and the final three columns show scenarios where the signal to noise ratio $A_{\text{SLF}}/A_{\text{DRW}}$ is increased from 1, 2.5, 5. The same DRW noise realization is kept throughout the plots. The top row shows the injected signal of each light-curve and the light-curves are shown in the middle row. The bottom row shows LS periodograms of the light-curves, where the two vertical red lines show the injected orbital frequency $f_{\text{orb}} = 1/T_{\text{orb}}$ and $2f_{\text{orb}}$ in units of 1/day. As $A_{\text{SLF}}/A_{\text{DRW}}$ increases, the DRW peaks decrease and the SLF peaks increase, but the SLF peaks never become the highest peaks.

incorporated into a LS analysis, but we instead turn to matched filters as an alternative approach.

B. Matched filters

The problem of recovering SLFs involves identifying not just a single prominent Fourier peak, but a collection of

Fourier peaks at the orbital frequency and its harmonics (Fig. 4). We propose matched filters as an alternative to LS periodograms for this problem, with the additional advantage that we can infer binary parameters, unlike with LS periodograms or in many other period-detection algorithms.




FIG. 4. Fourier transform of the injected self-lensing flares in the second column of Fig. 3. There was no noise included and the flares were sampled at a regular and dense cadence (opposed to random cadence in Fig. 3) for this illustration.

1. Matched filters: Approach

With matched filters, we aim to analyze the light-curve of the fiducial binary, shown in Fig. 5 where the binary-template has been degraded by adding DRW and LSST noise, and sampled at LSST-like observation times. The realistic mock light-curve of the fiducial binary is shown with black points in Fig. 5, and the faint blue light-curve shows the BINLITE light-curve used to generate the mock light-curve. We also assume that from other measurements, we have found the redshift z and total mass $M = M_1 + M_2$




FIG. 5. Illustration of the mock light-curve of the fiducial binary with realistic LSST data quality. The faint blue light-curve represents the BINLITE binary-template used to generate the signal $s(\hat{\theta})$ and the black data points $d(t) = s(\hat{\theta}) + n(t)$ represent the downsampled binary-template at LSST-like observation times with LSST-like noise and a DRW added. An estimate for the flare threshold is drawn with a dashed blue line, calculated via a simple analytical approximation for the lensing duration in Eq. (15).

TABLE II. The priors of the four free parameters in this matched filter exercise and the true injected values that we attempt to recover via matched filters. Any parameters not mentioned in this table ($q, M, z, f_{\text{edd}}, A_{\text{DRW}}, \tau_{\text{DRW}}$) are assumed to be known and are taken directly from Table I. $U[a, b]$ denotes a uniform distribution in the range $[a, b]$.

Fiducial Binary		
Binary Parameter	Injected	Prior
T_{orb} (years)	1	$U[0.5, 5]$
e	0.4	$[0.1, 0.2, 0.3, \dots, 0.7, 0.79]$
I	88°	$\cos^{-1}(U[0, 1])$
ω	90°	$U[0^\circ, 360^\circ]$

of the quasar to infinite precision. In addition, having a 10-year light-curve at the end of LSST will determine a mean luminosity, which then determines an Eddington ratio for a given mass M —which again we will assume we know precisely. Under these assumptions, the binary is characterized by the four free parameters of orbital period T_{orb} , eccentricity e , binary inclination I , and the argument of periaapsis ω (which is defined so that $\omega = 0^\circ$ means the line of sight is parallel to the semiminor axis and $\omega = 90^\circ$ means the line of sight is parallel to the semimajor axis). We assume that these parameters are unknown, and have priors summarized in Table II.

The matched filter is a method used to detect weak signals embedded in noise by maximizing the signal-to-noise ratio

$$\rho(\theta) \equiv \frac{s(\theta) * d}{\sqrt{s(\theta) * s(\theta)}}, \quad (12)$$

which is the normalized inner product between a binary-template time series $s(\theta)$ (i.e. $s_{\text{binlite}}(\theta)$) where θ are the template binary parameters, $\hat{\theta}$ are the true binary parameters that create the binary signal and observed time series $d(t) = s(\hat{\theta}) + n(t)$, where $n(t) = n_{\text{DRW}}(t) + n_{\text{LSST}}(t)$ with covariance $C_{ij} = S_{ij} + N_{\text{LSST}}$, where C_{ij} is a square matrix whose length is the number of light-curve data points, S_{ij} is the DRW covariance from Eq. (6) and N_{LSST} is a diagonal matrix with (uncorrelated) LSST noise. The binary-template $s(\theta)$ is interpolated and sampled at the observation times of the mock light-curve $d(t)$. We also maximize the SNR up to a time shift, so the numerator of Eq. (12) is the maximum SNR out of all possible shifts Δ of the binary-template light-curve

$$s(\theta) * d = \max_{\Delta} \left(\sum_t s[\Delta + t]^T C^{-1} d[t] \right). \quad (13)$$

The denominator is calculated as

$$\sqrt{s(\theta) * s(\theta)} = \sqrt{s(\theta)^T C^{-1} s(\theta)}, \quad (14)$$

which prevents bias towards binary-templates with poor fits but large amplitudes. In addition, we subtract the mean from the mock data $d(t)$ and the templates $s(\theta)$ before these calculations for the same purpose of preventing biases towards templates with large amplitudes. In practice, for a given template with known orbital period T'_{orb} we shift a given light-curve from 0 to T'_{orb} in 1.0 day increments to find the maximum SNR.

We note that due to large LSST observation gaps, the Fourier transform of the data $\tilde{d}$ cannot be computed without some degree of error, especially without some interpolation scheme. Thus we opt to compute the cross-correlation Eq. (12) in the time domain and not the frequency domain—although this method would be computationally more efficient through the use of fast Fourier transforms.

To efficiently match a large number of binary-templates to the mock light-curve, we precompute BINLITE binary-templates on the parameter space (the “template bank” method in gravitational wave analysis) and for a given light-curve, we calculate the SNR ρ for all the templates and look at the templates which have the highest SNR ρ . This has the advantage that these binary-templates are reusable for any number of light-curves that one wishes to analyze, as opposed to a method like Markov chain

Monte Carlo (MCMC), which requires template generation or interpolation during a search. However, this introduces a dependence on the spacing of the grid of binary-templates in parameter space, where a sparse grid will lead to missing the SNR-maximizing template, but an overly dense grid will become exponentially difficult to store and process.

Next, from the priors described in Table II, we choose 10000 points according to a Sobol sequence [51] using SCIPY’s SOBOLE package and calculate the r-band BINLITE light-curve for each point. A DRW realization $n_{\text{DRW}}(t)$ with DRW amplitude A_{DRW} of 0.1 mag and observed damping timescale τ_{DRW} of 300 days was added to the binary signal and the signal + noise light-curve was sampled at LSST-like observation cadences and gaps, and LSST-like noise $n_{\text{LSST}}(t)$ was also applied. For our fiducial binary at ~ 20 mag, the 1σ of LSST white noise is typically 0.01–0.05 mag. In Fig. 5, we show the BINLITE-simulated light-curve in blue and the noise-degraded LSST light-curve with black points. The mean of the light-curve was subtracted off to clearly visualize the various variability (hydrodynamical, self-lensing, DRW).

2. Matched filters: Results

We show 1-D slices of the distribution of SNR $\rho(\theta)$ in the parameter space in Fig. 6, i.e., $\rho(e, T_{\text{orb}} = 1 \text{ yr})$,




FIG. 6. Matched filter SNR, maximized over a phase shift Δ , of the fiducial binary relative to binary-templates in which 3 out of 4 binary parameters are fixed. The true binary parameters $e = 0.4$, $I = 88^\circ$, $T_{\text{orb}} = 1 \text{ yr}$, $w = 90^\circ$ are marked with red vertical lines. Top left panel: eccentricity dependence, i.e., $\rho(e, T_{\text{orb}} = 1 \text{ yr}, I = 88^\circ, w = 90^\circ)$. Top right panel: orbital period dependence—there are significant false peaks at $0.5T_{\text{orb}}$, $2.0T_{\text{orb}}$, $3.0T_{\text{orb}}$ etc. Bottom left panel: binary inclination dependence. Bottom right panel: argument of periaapsis dependence—there is only a mild dependence on ω , with SNRs of 67 and above being achieved as long as the other orbital parameters are fixed. The true injected parameters have the highest overall SNR of 71.

$I = 88^\circ, w = 90^\circ$) vs e , $\rho(I, T_{\text{orb}} = 1 \text{ yr}, e = 0.4, w = 90^\circ)$ vs I , etc. We again mark the values of the injected binary parameters $\hat{\theta}$ with red vertical lines. For T_{orb} , we find narrow SNR peaks at $0.5T_{\text{orb}}$ and $2T_{\text{orb}}$ and other SNR peaks $1.5T_{\text{orb}}$ and $2.5T_{\text{orb}}$. It is notable that these peaks are extremely narrow, of order few % of T_{orb} . This can be attributed to the fact that the maximum SNR will be achieved when all the flares are “matched” in a given binary template, and a binary-template with an orbital period that is too far away from T_{orb} or one of its multiples will necessarily miss some of the 10 lensing flares of the fiducial binary.

For the argument of periapsis ω , if all the orbital parameters are fixed, changing ω will impact the timing of the flares. However, a shift of the template can account for this, leading to relatively small variations ~ 3 in the SNR, caused by small deviations in the fit to the mock light-curve, and the fact that we used shift spacings of $\Delta = 1.0$ day.

For the orbital eccentricity, we find that there is only a mild loss in the SNR even for a template with a much different eccentricity. This is because for a template with the correct orbital period, inclination and argument of periapsis, the flares are already occurring at the correct times and (roughly) correct heights. We note that we use a discrete prior $e \in [0.1, 0.2, \dots]$ because of how the accretion rates from the B24 hydrodynamical simulations are sampled as a function of eccentricity due to their sweep technique they employ. We discuss this issue in Sec. IV.

Finally, for the orbital inclination, which primarily determines the angular separation between the two BHs at closest alignment, shows a significant loss of SNR when $I \lesssim 80^\circ$, which will not produce flares. In the other direction, we also find that there is a loss of SNR when $I = 90^\circ$, which should produce flares that are higher compared to $I = 88^\circ$.

With an understanding of the expected SNR distribution in parameter space, and the SNR ($\rho_{\text{max}} = 71$) at the true parameters for the fiducial binary, we perform a three-step analysis procedure to recover the injected binary-template:

- (i) Step #1: Draw 10,000 binary-templates from the wide prior described in Table II, and compute the matched-filter SNR for each.
- (ii) Step #2: Delete the self-lensing flares from the data and repeat Step #1 to accurately determine the orbital period. Steps #1 and #2 together reduce the search space in T_{orb} and I .
- (iii) Step #3: Draw 10,000 binary-templates from the narrowed prior in T_{orb} and I and compute the SNRs.

3. Step No. 1: Course grid

We present the results of the SNRs of the 10,000 binary-templates in Fig. 7. The SNRs should be compared to the null hypothesis of assuming there is no binary in the




FIG. 7. Distribution of the SNR ρ when matching the mock light-curve $d(t)$ to 10^4 sparsely sampled binary-templates and 10^5 DRW realizations. The maximum SNR of the DRW realizations, 25, is noted with a dashed orange vertical line and the SNR at the injected values, 71 (see Fig. 6), is noted with a dashed blue vertical line. We find that 536 templates out of the sparsely sampled 10^4 binary-templates have SNRs greater than the maximum SNR from DRW templates.

light-curve. Thus we also generate 10^5 pure DRW noise realizations with identical parameters ($A_{\text{DRW}} = 0.1$, $\tau_{\text{DRW}} = 300\text{d}$) to the ones used to generate the mock light-curve and compute their matched filter SNRs relative to the mock data. We find that DRWs can mimic binary signals up to an SNR of 25. This maximum value is shown with an orange dashed vertical line. We also find that 536 out of 10,000 binary-templates have SNRs that cannot be reached by these noise realizations. We show 5 binary-templates (after optimizing their time shift) with the highest SNR out of these 536 templates in Fig. 8. Although the templates match to several of the flares, none of them match to all of the flares present in the data. This can also be seen numerically, where the best-fit templates have SNRs of 47-54 while a best-fit close to the correct model should have a SNR close to 71 (see Fig. 6). The five binary-templates shown in Fig. 8 have binary parameters and optimal template shifts summarized in Table III. We find that for these five examples, the orbital periods are within 1% of either 1, 1.5, 2 years and all binary-templates had inclinations of at least 80 degrees.

We also show the parameter and SNR distribution of these binary-templates in a 2D histogram in Fig. 9, again noting that while many binary-templates have higher SNR than the largest among 10^5 DRW realizations (orange horizontal) none reach the SNR at injected values (marked with blue horizontal). We also note that there is a preference for higher eccentricities, which we attribute to a numerical artifact of the B24 hydrodynamical simulations. We discuss this further in Sec. IV.

We also find constraints on the binary inclination, with 86% of these 536 high-SNR templates having values above greater than 85° . However, the highest SNRs from this do not reach the true injected value, and we conclude that the




FIG. 8. The five binary-templates with highest SNR from Step #1, Fig. 7, overlaid on the mock fiducial light-curve. In order of higher SNR to lower SNR, the binary-templates have orbital periods of $T_{\text{orb}} = 1.00, 0.50, 1.50, 2.00, 1.50$ years.

initial search of sampling 10,000 binary-templates from a relatively sparse prior (Table II) is insufficient to converge on the correct model because of the steep SNR dependence on T_{orb} .

4. Step No. 2: Excised SLFs

A better constraint on periodicity is necessary as it is the parameter that leads to the largest changes in the SNR (see Fig. 6). From Step No. 1, although we were able to constrain the injected orbital period to 1, 1.5 or 2.0 years, we need a more robust constraint. We take

advantage of a second source of periodicity besides the flares: the hydrodynamical variability in the baseline. In order to analyze the baseline periodicity, we clip the flares from the light-curve, using the approximation that a lensing event lasts for the time it takes the lens to cross 4 Einstein radii of the source at alignment

$$t_{\text{lensing}} = 2 \times \left(\frac{2R_{E,\text{source}}}{v_{\text{orb}}} \right) \approx 2 \times \left(\frac{R_{E,\text{source}}}{\pi a} T_{\text{orb}} \right), \quad (15)$$

where $R_{E,\text{source}}$ is the Einstein radius of the source and we assume a circular orbit to approximate the orbital velocity v_{orb} . Then we remove the fraction $t_{\text{lensing}}/T_{\text{orb}} = 2R_{E,\text{source}}/\pi a$ of the brightest points from the light-curves, assuming that these correspond to the flares.

While Fig. 9 shows that $T_{\text{orb}} = 0.5, 1.0$ year has the highest SNR we purposely use $\tilde{T}_{\text{orb}} = 2.0$ yr as an incorrect Ansatz to calculate the semimajor axis a in Eq. (15). As shown in Fig. 5, this threshold still effectively deletes most of the flares.

We then repeat the analysis of Step No. 1, matching the same 10,000 templates to the mock light-curve with flares deleted. We present the SNR, orbital period distribution results of Step No. 2 in Fig. 10. We find that this is effective in identifying the injected periodicity, with 4 binary-templates at $T_{\text{orb}} \sim 1.00$ year and 1 binary-template at $T_{\text{orb}} \sim 2.00$ year (see Table IV) having SNRs

TABLE III. Top-ranked binary-templates from Step No. 1 (Fig. 7). First row shows the injected parameters for the fiducial binary and each row below lists the ranking (by SNR), SNR value, and corresponding binary parameters, including the optimal template shift Δ_{opt} which maximizes the SNR.

Rank	SNR	T_{orb} (yr)	$I(^{\circ})$	e	$\omega(^{\circ})$	Δ_{opt} (d)
Injected	70.6	1.00	88.00	0.40	90.00	0
1	54.3	1.00	80.96	0.70	84.54	2
2	54.2	0.50	84.15	0.10	40.07	23
3	48.0	1.50	88.01	0.40	70.25	376
4	47.8	2.00	87.27	0.79	333.67	353
5	47.5	1.50	87.56	0.70	230.37	376




FIG. 9. Step No. 1, 2D histogram of 536 out of 10,000 binary-templates with parameters sampled from a wide prior, binned by their binary parameters (horizontal axis) and matched filter SNRs (vertical axis). The color of each (parameter, SNR) bin indicates the fraction of the selected 536 binary templates that fall into that bin. These binary-templates have SNRs higher than the maximum SNR achieved with 10^5 DRW realizations from Fig. 7, marked with an orange dashed line. While our goal was to sample the SNR peak at the injected parameters (~ 71 , marked with blue dashed line), we find that most binary-templates have SNRs of 25-40 when matched to the mock data. Injected binary parameters are marked with red lines.




FIG. 10. Step No. 2, 2D histogram of 10,000 binary-templates with parameters sampled from a wide prior, binned by their binary parameters (horizontal axis) and matched filter SNRs (vertical axis). The color of each (parameter, SNR) bin indicates the fraction of templates that fall into that bin. The horizontal orange dashed line indicates the maximum SNR from 10^5 DRW realizations. Injected orbital period is marked with a vertical red line. We find 4 binary-templates with orbital period ~ 1 yr and 1 binary-template with orbital period ~ 2 yr which have SNRs higher than DRWs.

greater than SNRs from DRW realizations (these correspond to the 5 binary-templates in Fig. 10 which have SNR greater than the orange dashed line representing the maximum SNR from DRWs). Thus we can deduce, together from the information from Fig. 9 that there are periodically occurring flares from an edge-on $I \gtrsim 85^\circ$ binary, and from the T_{orb} panel of Step No. 2 that the

TABLE IV. Top-ranked binary-templates from Step No. 2. These templates are also the only five binary-templates with SNRs greater than the maximum from DRWs (orange dashed line in Fig. 10). First row shows the injected parameters for the fiducial binary and each row below lists the ranking (by SNR), SNR value, recovered orbital period, and the optimal template shift Δ_{opt} which maximizes the SNR.

Rank	SNR	T_{orb} (yr)	Δ_{opt} (d)
1	8.8	1.0042	38
2	8.4	1.0019	4
3	8.2	1.0005	0
4	8.1	0.9976	76
5	7.9	2.0009	5

TABLE V. The narrowed prior for the fiducial binary inferred in Steps No. 1, No. 2. Priors of parameters not mentioned in this table are kept the same as in Table I and Table II.

Parameter	Fiducial $\hat{\theta}$	Narrowed prior
T_{orb} (years)	1	$U[0.95, 1.05]$
$I(^{\circ})$	88	$U[85, 90]$

true orbital period is ~ 1 yr. Using these new constraints, we set a narrowed prior of $T_{\text{orb}} \in [0.95, 1.05]$ years around the orbital period and the binary inclination $I = [85^{\circ}, 90^{\circ}]$, which are summarized in Table V.

5. Step No. 3: Fine grid

Finally, we Sobol-sample 10,000 binary-templates from the narrowed prior in I, T_{orb} of Table V. Priors of parameters not mentioned in Table V are kept the same as in Table I and Table II. We present the parameter and SNR

distribution of 5,006 binary-templates which exceed the DRW SNR threshold of 25 in Step No. 1 in Fig. 11. As a result we find that the SNRs of various binary-templates nearly reach the SNR at the injected values. We also plot the three highest-SNR binary-templates in Fig. 12, which now match all of the flares. The binary parameters of these binary-templates are listed in Table VI. We also show the resulting parameter and SNR distribution in a corner plot (Fig. 13). The marginalized posterior contours are drawn using relative log-likelihood weights of $\ln \mathcal{L} - \ln \mathcal{L}_{\text{max}} = -\frac{1}{2}(\rho_{\text{max}}^2 - \rho^2)$, which follows from the log-likelihood $\ln \mathcal{L}(\theta) = \frac{1}{2}\rho^2(\theta) - \frac{1}{2}d^T C^{-1}d$ (assuming Gaussian-distributed noise) and the second term on the right hand side is independent of the binary parameters θ . [52]. Here, $\rho_{\text{max}} = 69.8$ is the maximum SNR found among the 10,000 binary-templates. We recover the injected orbital period $T_{\text{orb}} = 1.00^{+0.03}_{-0.04}$ yr and the binary inclination $I = 88.12^{+1.35^{\circ}}_{-2.00^{\circ}}$, where the values indicate the medians and the 16th, 84th percentile intervals of the marginalized




FIG. 11. Step No. 3, 2D histogram of 5006 out of 10,000 binary-templates with parameters sampled from a narrow prior, binned by their binary parameters (horizontal axis) and matched filter SNRs (vertical axis). The color of each (parameter, SNR) bin indicates the fraction of the selected 5006 binary templates that fall into that bin. These binary-templates have SNRs higher than the maximum SNR achieved with 10^5 DRW realizations from Fig. 7, marked with an orange dashed line. Injected binary parameters are marked with red lines. The SNR at injected values, marked by a blue dashed line, is now achieved as a result of our three-step procedure.




FIG. 12. Final results (Step No. 3): the three highest-SNR templates. The template SNRs are close to the SNR at the injected value 71 (see also Table VI) and the templates match all of the self-lensing flares present in the mock data.

posteriors of each parameter. We again find a preference for high eccentricities in Fig. 13. There are also two clear peaks at $\omega = 90^\circ, 270^\circ$ but the distribution is overall uniform, as can also be seen from the top-ranked binary-templates in Table VI. This is because, for a fixed eccentricity, a change in the argument of periapsis primarily changes the timing of the flare within an orbit, but this change can be accounted for by the time shifts for each binary-template (see, for example, the ω and Δ_{opt} values of the high-SNR binary-templates in Table VI). Since the secondary flare was poorly sampled in this fiducial binary, it does not contribute to the SNR, and thus any value of ω can lead to a good fit with the mock light-curve as long as the other parameters are close to the injected parameters.

While the best-fit values of T_{orb} , I , and ω in Table VI are consistent with the posterior shown in Fig. 13, the eccentricity behaves differently. Although the injected eccentricity $e = 0.4$ yields the highest SNR at its peak (see Fig. 6), the peak is extremely narrow in (I, ω) space. In contrast, templates with higher eccentricities ($e = 0.7, 0.79$) produce

slightly lower SNRs but over a much broader region of parameter space. When marginalizing over (I, ω) , this broader “plateau” contributes a larger total likelihood volume than the narrow peak at $e = 0.4$, leading to an apparent preference for higher eccentricity in the corner plot. A visualization of this plateau in (I, ω) space is provided in Fig. 16 in the Appendix.

IV. DISCUSSION

A. Limitations of mock binary light-curves

Our analysis and results in Sec. III are based on the toy models for periodic binary light-curves in BINLITE. While this is suitable a first approximation, it leads to several caveats and limitations.

The BINLITE light-curves are toy models which are built from blackbody emission from minidisks with equilibrium temperature profiles $T(r) \propto (M\dot{M}/r^3)^{1/4}$ and accretion rates measured in equal-mass, 2D Newtonian, hydrodynamical simulations with a locally isothermal equation of state. In reality, for quasars, the connection between accretion rates and luminosity is more complicated and not deterministic, so that light-curves should be directly extracted from more realistic simulations. For the construction of such light-curves, it will be necessary to incorporate the effects of radiation pressure, magnetic fields, general relativity, and the self-gravity of the gas into three-dimensional simulations. BINLITE also needs to be extended to include hydrodynamical simulations of unequal-mass binaries [50]—note that Paper IV and [38] both find that mass ratios down to $q \sim 0.1$ will be relevant for self-lensing.

We also had to use a discrete and sparse prior in e , because of a more subtle issue with BINLITE that results in a rapid oscillation of accretion rate ratio

TABLE VI. Top-ranked binary-templates from Step #3 (Fig. 12). Like in Table III, the first row shows the injected parameters for the fiducial binary and each row below lists the ranking (by SNR), SNR value, corresponding binary parameters, and the optimal template shift Δ_{opt} which maximizes the SNR.

Rank	SNR	T_{orb} (yr)	I ($^\circ$)	e	ω ($^\circ$)	Δ_{opt} (d)
Injected	70.6	1.00	88.00	0.40	90.00	0
1	69.8	1.00	88.97	0.20	45.81	37
2	69.2	1.00	89.48	0.20	304.05	133
3	68.7	1.00	87.05	0.50	42.03	15




FIG. 13. Corner plots of the results from Step #3 with contours drawn using relative log-likelihood weights of $\ln \mathcal{L} - \ln \mathcal{L}_{\text{max}} = -\frac{1}{2}(\rho_{\text{max}}^2 - \rho^2)$. We recover the injected orbital period and binary inclination accurately. We find that the posterior distribution for the argument of periaapsis also peaks at the injected orientation but we identify a bias towards higher eccentricities. See Sec. IV for further discussion.

$Q \equiv \langle \dot{M}_2(t) \rangle / \langle \dot{M}_1(t) \rangle$ as a function of e . This is an artifact from the BINLITE accretion rates, which are generated from a single hydrodynamical simulation sweep in which the eccentricity is varied continuously from $e = 0$ to $e = 0.8$. Each eccentricity value therefore corresponds to a particular time in their sweep, from which BINLITE extracts its accretion rates. Along this sweep in eccentricity, the accretion rates $\dot{M}_1, \dot{M}_2$ onto the two minidisks fluctuate,

causing the ratio $\dot{M}_1/\dot{M}_2$ to fluctuate between values above and below unity as well. We demonstrate the accretion rate switching in Fig. 14 (and in the Appendix), where “neighboring” eccentricities produce different flare profiles and SNRs due to frequent switches in which minidisk is brighter. This leads to a matched filter SNR profile that oscillates as a function of eccentricity whenever the accretion switching occurs.




FIG. 14. Matched filter SNR for the fiducial binary for a continuous prior in eccentricity. The oscillatory trend as a function of e is artificial and arises from how BINLINE samples the accretion switching between the two BH components. See Sec. IV for further discussion.

As a result, lensing flare profiles look different as well, especially when the source and lens roles are reversed (for example, $\omega = 90^\circ, 270^\circ$). This introduces a dependency on the argument of periastris. For example, for certain eccentricities (e.g., $e = 0.79$), the accretion ratio is closer to unity. Then, the flare profiles and the fit to the mock light-curve of the fiducial binary are similar when lens/source roles are swapped. Meanwhile for $e = 0.4$, the accretion ratio is unequal, so when lens/source roles are switched ($\omega = 90^\circ \rightarrow 270^\circ$) this leads to a much smaller flare because the dimmer disk becomes the source and the brighter disk becomes the lens. When marginalizing over ω in the corner plot (Fig. 13), we find that this causes an apparent bias towards higher eccentricities. We visualize this effect further in the Appendix.

B24 also use the point-source approximation for self-lensing. In reality, most self-lensing binaries in LSST are expected to have masses of $\log(M/M_\odot) = 7\text{--}9$, mass ratios of $q = 0.1\text{--}1$, and orbital periods of a few years [33,38]. According to [53], these binaries will have accretion disc sizes that are comparable to the Einstein radius of the lenses, making the point-source approximation invalid. When the source is treated as a resolved disk and not a point, the effect is that the “Paczynski curve” (magnification vs time) is smeared out (dimmed and widened). When considering such resolved accretion disks, their physical properties may become relevant. For example, for radiatively inefficient accretion flows (RIAF), the radius of the mini-disk primarily responsible for the observed-frame optical emission could be significantly different from that in a thin-disk model. However, paper IV shows that the widening of the flares is more important than the dimming for detection rates. Thus incorporating the finite-source calculation of self-lensing in BINLITE is crucial as well and we expect that it could improve our results by making the flares less narrow and harder to miss. Beyond BINLITE, a future matched filter approach should also incorporate uncertainties on the binary total mass, Eddington ratio and redshift, which were assumed to




FIG. 15. Distribution of 1D posteriors when applying Steps No. 1–3 on the fiducial binary but with two other noise realizations. Black distributions are identical to the 1D distributions on the diagonal of the corner plot Fig. 13 and blue, green distributions correspond to the two other noise realizations.

be known. Developing mock binary light-curves incorporating various uncertainties could potentially be addressed by the concept of “fuzzy templates” [54].

Finally, our analysis here was applied to a single band (r -band, chosen arbitrarily), whereas LSST will have data in 6 optical bands. Since SLFs are roughly achromatic between optical bands even in the finite-source limit, data from multiple bands will improve detections by increasing the number of data points and also the confidence of a detection. For these multiwavelength matched filters (MMFs), correlations of quasar noise between optical bands will be required.

B. Different DRW realizations

Another issue is that our analysis above was applied to a single mock binary, with a single noise realization. To assess the robustness of our matched filter methods to the

noise realization, we repeated Steps No. 1, No. 2, No. 3 on the same fiducial binary parameters but for two more different realizations of the quasar DRW variability. Figure 15 shows the resulting 1D marginalized posterior distributions for the four inferred parameters. The black distributions are identical to the posteriors obtained in Fig. 13, while the blue and green curves show the distributions for the two additional DRW realizations. The three sets of distributions overlap closely, showing that our procedure is not sensitive to the specific DRW realization.

C. Beyond the fiducial binary

We next discuss the application of our matched filter procedure to binaries with parameters different from those of the fiducial binary. Given that flares dominate the matched filter SNR, the performance of our method depends sensitively on the sampling of the self-lensing flares. Binaries with longer-duration flares will have more data points covering each individual flare, which increases their detectability in our pipeline and improves periodicity recovery. Conversely, if some flares are missed due to observational gaps of the light-curve or because of a short flare durations, the detectability will degrade.

Our method also partially relies (in Step No. 2) on the hydrodynamical and Doppler variability in the baseline of the light-curve. In the case that the flares are poorly sampled, and quasar DRW noise dominates over baseline variability (e.g., small $A_{\text{hydro}}/A_{\text{DRW}}$ or $A_{\text{DB}}/A_{\text{DRW}}$), it will be difficult to recover the periodicity in Step No. 2 and sample binary-templates in the vicinity of the injected parameters in Step No. 3. However, we also found that Step No. 1 identified multiple (though not all) flares in the mock light-curve (see Fig. 8) with SNRs above the maximum of 10^5 DRWs, despite sampling binary-templates from a wide prior (Table II). This indicates that our pipeline can still flag such noisy light-curves, although the subsequent parameter recovery is less accurate when only a subset of flares is detected.

In our fiducial binary's mock light-curve (Fig. 5), the secondary self-lensing flares were much smaller compared to the primary and were also often missed due to sampling. In self-lensing binaries with brighter secondary flares that are adequately sampled and not drowned out by noise, we expect improved parameter recovery in (e, ω) . For example, for fixed $\omega, T_{\text{orb}}, I$, a higher eccentricity leads to a shorter time interval between the flares followed by a longer waiting time for the next pair of flares to occur. On the other hand, for fixed e, T_{orb}, I , flare widths can also be used as an additional constraint. For example, the width of a self-lensing flare is shorter if $\omega \sim 90^\circ$ when the secondary is at periapsis and moving quickly at closest alignment.

We further tested whether our weak eccentricity dependence (Fig. 6, top-left panel) is specific to the fiducial orientation $\omega = 90^\circ$, where the line of sight is parallel to the semimajor axis and the flares always occur at periapsis and apoapsis for any eccentricity. By repeating the analysis at a

different argument of periapsis $\omega = 40^\circ$, we found a similar weak dependence of SNR on eccentricity. Again, this insensitivity arises because only one of the two flares per orbit contributes significantly to the SNR. Investigating the detectability and parameter estimation of binaries with a pair of bright and well-sampled self-lensing flares per orbit will be the subject of future work.

D. Computational cost

We consider computational costs necessary to search the 10-year catalog of LSST, which is projected to contain $N_{\text{QSO}} = 20\text{--}100$ million quasars [25]. The number of CPU hours required for such a search will come from the amount of time to prepare the binary-templates $N_{\text{templates}}$, which can be shared for the analysis of all quasar light-curves, and the amount of time to compute matched filter SNRs $N_{\text{QSO}} \times N_{\text{templates}}$ times. Although we plan to include other parameters like the mass ratio q or the total mass M of the binary in future searches, we assume that $N_{\text{templates}} = 10^4$ Sobol-sampled points will be more than sufficient to explore our parameter space.

On average, we estimate that preparing 1 BINLITE binary-template requires $t_{\text{pre}} \sim 0.4$ seconds and calculating 1 matched filter SNR requires $t_{\text{SNR}} \sim 0.2$ seconds. Then the total CPU time required to search N_{QSO} light-curves is

$$\begin{aligned} t_{\text{CPU}} &= N_{\text{templates}} \times t_{\text{pre}} + N_{\text{QSO}} \times N_{\text{templates}} \times t_{\text{SNR}} \\ &\simeq 22 \text{ million} \left(\frac{N_{\text{QSO}}}{2 \times 10^7} \right) \left(\frac{N_{\text{templates}}}{10^4} \right) \text{ CPU hours,} \end{aligned} \quad (16)$$

where the template-preparation time (first term) is negligible compared to the total SNR calculation time (second term). t_{SNR} consists of the time to perform interpolation and dot-product operations using SCIPY [55] and we plan to optimize these calculations further, which would proportionally reduce the computation costs.

A simpler approach could be to use computationally inexpensive methods like second- or third-order structure functions as an initial filter to reduce the number of candidates requiring our relatively expensive matched filter approach. In particular, the third-order structure function is known to effectively capture temporally asymmetric flares and other variability [56]. Beyond matched filters, we plan to consider and test a variety of methods in preparation of the LSST data releases.

V. SUMMARY AND CONCLUSIONS

It is estimated that there will be up to thousands of edge-on binaries exhibiting self-lensing flares lasting 30-days or longer and having orbital periods of 5 years or less (at least 2 orbits in 10 years). Using toy models of SMBHB photometric variability, we tested the efficacy of two

statistical tools, LS periodograms and matched filters, to recover injected self-lensing flares in noisy light-curves. The conclusions of our study can be summarized as follows:

- (i) None of the self-lensing flares have LS periodogram peaks exceeding the height of flares among 10,000 DRWs, even in ideal scenarios. A standard LS periodogram is unable to effectively recover periodicity from self-lensing flares.
- (ii) The matched filter SNR effectively separates self-lensing flares from DRW noise, yielding SNRs much greater than the maximum value found in 10^5 DRW realizations.
- (iii) For our fiducial binary, a three-step matched filter procedure can identify periodicities as well as recover the injected periodicity and binary inclination.
- (iv) We find that constraining eccentricity is difficult because the secondary flare was not adequately sampled in our mock light-curve, leading to degeneracies among e , ω , and an overall phase shift in the light-curve. We expect that bright and well-sampled secondary flares can break these degeneracies and help constrain (e, ω) .

We demonstrate that matched filters can potentially find self-lensing flares in LSST with reasonable computational cost. In the future, this approach will be refined by widening our prior to include the mass ratio as a free parameter, and more realistic light-curves that include 3D, general-relativistic effects, etc. Moreover, our single-band analysis should be extended to simultaneously analyzing all 6 bands of LSST.

Software: PYTHON [57,58], NUMPY [59], MATPLOTLIB [60], SCIPY [55], astroML [61], BINLITE [45], CORNER [62]

ACKNOWLEDGMENTS

We acknowledge support by NASA Grants No. 80NSSC22K0822 and No. 80NSSC24K0440 (Z. H.). V. A. V. acknowledges support through the David and Lucile Packard Foundation, National Science Foundation under No. PHY-2019786 (The NSF AI Institute for Artificial Intelligence and Fundamental Interactions, [63]), No. AST-2433718, No. AST-2407922 and No. AST-2406110, as well as an Aramont Fellowship for Emerging Science Research. J. D. acknowledges support by a joint Columbia/Flatiron Postdoctoral Fellowship. Research at the Flatiron Institute is supported by the Simons Foundation. This research was supported in part by the National Science Foundation under Grant No. NSF PHY-1748958 to the Kavli Institute for Theoretical Physics (KITP). This research has made use of NASA's Astrophysics Data System. The authors acknowledge the Texas Advanced Computing Center (TACC) at The University of Texas at Austin for providing computational resources. V. A. V. acknowledges support through the David and Lucile Packard Foundation, National Science Foundation under No. AST-2433718, No. AST-2407922 and No. AST-2406110, as well as an Aramont Fellowship for Emerging Science Research. This work is supported by the National Science Foundation under Cooperative Agreement No. PHY-2019786 (The NSF AI Institute for Artificial Intelligence and Fundamental Interactions, [63]).

DATA AVAILABILITY

The data are not publicly available. The data are available from the authors upon reasonable request.

APPENDIX: SNR SURFACES




FIG. 16. First row: SNR surfaces for $e = 0.4$ (injected value) and $e = 0.79$ of the fiducial binary in (I, ω) space, where the brightest color indicates the SNR at the injected values (71). The two blue, orange points on each panel indicate points that characterize the e plateau described in Sec. III. For $e = 0.4$, at the blue point $(\omega, I) = (90^\circ, 85.5^\circ)$, the SNR is 67.9, whereas at the orange point $(\omega, I) = (270^\circ, 85.5^\circ)$, the SNR is 46.2. However for $e = 0.79$, at the blue point $(\omega, I) = (90^\circ, 85.5^\circ)$, the SNR is 65.5, whereas at the orange point $(\omega, I) = (270^\circ, 85.5^\circ)$, the SNR is 65.2. There is a clear difference in SNR for $e = 0.4$ vs $e = 0.79$ when the lensing flare is viewed from the opposite side. Second row: Binary-template light-curves generated from the blue and orange points in the (I, ω) parameter space in the first row. There is a large difference in the flare heights for $e = 0.4$ for $\omega = 90^\circ, 270^\circ$ whereas for $e = 0.79$ the difference is much smaller. Third row: The accretion rates of the primary (dashed) and secondary (solid) minidisks, determined by the BINLITE simulation sweep. These are plotted in gray to emphasize that they depend only on eccentricity and not on (I, ω) . The accretion-rate contrast between the two minidisks is large for $e = 0.4$ and much smaller for $e = 0.79$, explaining the asymmetric versus symmetric flare behavior seen in the first two rows.

- [1] J. Kormendy and L. C. Ho, *Annu. Rev. Astron. Astrophys.* **51**, 511 (2013).
- [2] M. C. Begelman, R. D. Blandford, and M. J. Rees, *Nature (London)* **287**, 307 (1980).
- [3] B. D. Farris, P. Duffell, A. I. MacFadyen, and Z. Haiman, *Mon. Not. R. Astron. Soc.: Lett.* **446**, L36 (2015).
- [4] Y. Tang, Z. Haiman, and A. MacFadyen, *Mon. Not. R. Astron. Soc.* **476**, 2249 (2018).
- [5] L. Major Krauth, J. Davelaar, Z. Haiman, J. R. Westernacher-Schneider, J. Zrake, and A. MacFadyen, *Mon. Not. R. Astron. Soc.* **526**, 5441 (2023).
- [6] A. J. Dittmann, G. Ryan, and M. C. Miller, *Astrophys. J. Lett.* **949**, L30 (2023).
- [7] A. De Rosa *et al.*, *New Astron. Rev.* **86**, 101525 (2019).
- [8] D. J. D’Orazio and M. Charisi, [arXiv:2310.16896](https://arxiv.org/abs/2310.16896).
- [9] T. Bogdanovic, M. C. Miller, and L. Blecha, *Living Rev. Relativity* **25**, 3 (2022).
- [10] M. J. Graham, S. G. Djorgovski, D. Stern, A. J. Drake, A. A. Mahabal, C. Donalek, E. Glikman, S. Larson, and E. Christensen, *Mon. Not. R. Astron. Soc.* **453**, 1562 (2015).
- [11] Y.-J. Chen, S. Zhai, J.-R. Liu, W.-J. Guo, Y.-C. Peng, Y.-R. Li, Y.-Y. Songsheng, P. Du, C. Hu, and J.-M. Wang, *Mon. Not. R. Astron. Soc.* **527**, 12154 (2024).
- [12] M. Charisi, I. Bartos, Z. Haiman, A. M. Price-Whelan, M. J. Graham, E. C. Bellm, R. R. Laher, and S. Márka, *Mon. Not. R. Astron. Soc.* **463**, 2145 (2016).
- [13] P. Huijse, J. Davelaar, J. D. Ridder, N. Jannsen, and C. Aerts, [arXiv:2505.16884](https://arxiv.org/abs/2505.16884).
- [14] J. Robnik, A. E. Bayer, M. Charisi, Z. Haiman, A. Lin, and U. Seljak, *Mon. Not. R. Astron. Soc.* **534**, 1609 (2024).
- [15] A. Sesana, Z. Haiman, B. Kocsis, and L. Z. Kelley, *Astrophys. J.* **856**, 42 (2018).
- [16] K. El-Badry, D. W. Hogg, and H.-W. Rix, [arXiv:2509.10601](https://arxiv.org/abs/2509.10601).
- [17] G. Agazie, A. Anumalapudi, A. M. Archibald *et al.* (The Nanograv Collaboration), *Astrophys. J. Lett.* **951**, L8 (2023).
- [18] J. Antoniadis, P. Arumugam, S. Arumugam *et al.* (The EPTA and InPTA Collaborations), *Astron. Astrophys.* **678**, A50 (2023).
- [19] D. J. Reardon, A. Zic, R. M. Shannon, G. B. Hobbs, M. Bailes, V. Di Marco, A. Kapur, A. F. Rogers, E. Thrane *et al.*, *Astrophys. J. Lett.* **951**, L6 (2023).
- [20] H. Xu *et al.*, *Res. Astron. Astrophys.* **23**, 075024 (2023).
- [21] E. A. Nikita Agarwal, *Astrophys. J. Lett.* **998**, L11 (2026).
- [22] P. Amaro-Seoane *et al.*, [arXiv:1702.00786](https://arxiv.org/abs/1702.00786).
- [23] C. Xin and Z. Haiman, *Mon. Not. R. Astron. Soc.* **533**, 3164 (2024).
- [24] J. Baker *et al.*, [arXiv:1903.04417](https://arxiv.org/abs/1903.04417).
- [25] C. Xin and Z. Haiman, *Mon. Not. R. Astron. Soc.* **506**, 2408 (2021).
- [26] C. L. MacLeod, Ž. Ivezić, C. S. Kochanek, S. Kozłowski, B. Kelly, E. Bullock, A. Kimball, B. Sesar, D. Westman, K. Brooks, R. Gibson, A. C. Becker, and W. H. de Vries, *Astrophys. J.* **721**, 1014 (2010).
- [27] W. Yu, G. T. Richards, J. J. Ruan, M. S. Vogeley, F. E. Bauer, and M. J. Graham, [arXiv:2508.12076](https://arxiv.org/abs/2508.12076).
- [28] S. Vaughan, P. Uttley, A. G. Markowitz, D. Huppenkothen, M. J. Middleton, W. N. Alston, J. D. Scargle, and W. M. Farr, *Mon. Not. R. Astron. Soc.* **461**, 3145 (2016).
- [29] K. El-Badry, D. W. Hogg, and H.-W. Rix, [arXiv:2509.10601](https://arxiv.org/abs/2509.10601).
- [30] J. Davelaar and Z. Haiman, *Phys. Rev. D* **105**, 103010 (2022).
- [31] J. Davelaar and Z. Haiman, *Phys. Rev. Lett.* **128**, 191101 (2022).
- [32] L. M. Krauth, J. Davelaar, Z. Haiman, J. R. Westernacher-Schneider, J. Zrake, and A. MacFadyen, *Phys. Rev. D* **109**, 103014 (2024).
- [33] K. Park, C. Xin, J. Davelaar, and Z. Haiman, *Phys. Rev. D* **111**, 063011 (2025).
- [34] B. X. Hu, D. J. D’Orazio, Z. Haiman, K. L. Smith, B. Snios, M. Charisi, and R. Di Stefano, *Mon. Not. R. Astron. Soc.* **495**, 4061 (2020).
- [35] E. Kun, S. Frey, and K. É Gabányi, *Mon. Not. R. Astron. Soc.* **496**, 3336 (2020).
- [36] N. M. Sorabella, S. Bhattacharya, S. G. T. Laycock, D. M. Christodoulou, and A. Massarotti, *Astrophys. J.* **927**, 234 (2022).
- [37] W. Kollatschny and D. Chelouche, *Astron. Astrophys.* **690**, L2 (2024).
- [38] L. Z. Kelley, D. J. D’Orazio, and R. Di Stefano, *Mon. Not. R. Astron. Soc.* **508**, 2524 (2021).
- [39] P. A. Abell, J. Allison, S. F. Anderson, J. R. Andrew *et al.* (LSST Science Collaboration), [arXiv:0912.0201](https://arxiv.org/abs/0912.0201).
- [40] D. Sijacki, M. Vogelsberger, S. Genel, V. Springel, P. Torrey, G. F. Snyder, D. Nelson, and L. Hernquist, *Mon. Not. R. Astron. Soc.* **452**, 575 (2015).
- [41] Á. Kis-Tóth, Z. Haiman, and Z. Frei, *Classical Quantum Gravity* **42**, 075007 (2025).
- [42] K. Porter, S. C. Noble, E. M. Gutierrez, J. Pelle, M. Campanelli, J. Schnittman, and B. J. Kelly, *Astrophys. J.* **979**, 155 (2025).
- [43] C. Xin, M. Isi, W. M. Farr, and Z. Haiman, [arXiv:2506.10846](https://arxiv.org/abs/2506.10846).
- [44] Z. Haiman, *Phys. Rev. D* **96**, 023004 (2017).
- [45] D. J. D’Orazio, P. C. Duffell, and C. Tiede, *Astrophys. J.* **977**, 244 (2024).
- [46] A. Lin, M. Charisi, and Z. Haiman, [arXiv:2505.14778](https://arxiv.org/abs/2505.14778).
- [47] T. Bogdanović, C. S. Reynolds, and M. C. Miller, *Astrophys. J.* **661**, L147 (2007).
- [48] Rubin Observatory, Baseline_v3.3 Survey Strategy Release, <https://community.lsst.org/t/baseline-v3-3-run-released/8042> (2024), Rubin Observatory Community Forum.
- [49] v. Ivezić, R. L. Jones, and R. Lupton, The LSST photon rates and SNR calculations, v1.2, Technical Report No. LSE-40, Department of Astronomy, University of Washington, 2010.
- [50] M. Siwek, R. Weinberger, and L. Hernquist, *Mon. Not. R. Astron. Soc.* **522**, 2707 (2023).
- [51] I. M. Sobol’, *Zh. Vychisl. Mat. i Mat. Phys.* **7**, 784 (1967).
- [52] Í. Zubeldia, A. Rotti, J. Chluba, and R. Battye, *Mon. Not. R. Astron. Soc.* **507**, 4852 (2021).
- [53] D. J. D’Orazio and R. Di Stefano, *Mon. Not. R. Astron. Soc.* **474**, 2975 (2018).
- [54] A. Lowe, R. W. Jones, and M. J. Harrison, *Journal of Intelligent Information Systems: Integrating Artificial Intelligence and Database Technologies* **13**, 27 (1999).
- [55] P. Virtanen *et al.* (SciPy 1.0 Contributors), *Nat. Methods* **17**, 261 (2020).

-
- [56] S. D. von Fellenberg *et al.*, *Astron. Astrophys.* **688**, L12 (2024).
- [57] T. E. Oliphant, *Comput. Sci. Eng.* **9**, 10 (2007).
- [58] K. J. Millman and M. Aivazis, *Comput. Sci. Eng.* **13**, 9 (2011).
- [59] S. van der Walt, S. C. Colbert, and G. Varoquaux, *Comput. Sci. Eng.* **13**, 22 (2011).
- [60] J. D. Hunter, *Comput. Sci. Eng.* **9**, 90 (2007).
- [61] Ž. Ivezić, A. Connolly, J. Vanderplas, and A. Gray, *Statistics, Data Mining and Machine Learning in Astronomy* (Princeton University Press, Princeton, NJ, 2014).
- [62] D. Foreman-Mackey, *J. Open Source Software* **1**, 24 (2016).
- [63] <http://iaifi.org/>.