

VAR-PZ: Constraining the photometric redshifts of quasars using variability

Sarath Satheesh-Sheeba^{1*} , Roberto J. Assef² , Timo Anguita^{1,3} , Paula Sánchez-Sáez⁴ , Raphael Shirley⁵ , Tonima T. Ananna⁶ , Franz E. Bauer⁷ , Alexey Bobrick^{8,9} , Carlos Guillermo Bornancini^{10,11} , Sarah E. I. Bosman^{12,13} , W.N. Brandt^{14,15,16} , Demetra De Cicco^{17,3,18} , Bozena Czerny¹⁹ , Marta Fatović^{20,21,18} , Kohei Ichikawa^{22,23} , Dragana Ilić^{24,25} , Andjelka B. Kovačević²⁴ , Guodong Li²⁶ , Mai Liao² , Alejandra Rojas-Lilayú²⁷ , Marcin Marculewicz⁶ , Daniel Marsango²⁸ , Chiara Mazzucchelli² , Tatevik Mkrtchyan² , Swayamtrupta Panda²⁹ , Alessandro Peca^{30,31} , Bindu Rani^{32,33} , Claudio Ricci^{34,2} , Gordon T. Richards³⁵ , Mara Salvato^{5,36} , Donald P. Schneider^{14,15} , Matthew J. Temple³⁷ , Francesco Tombesi^{38,39,40} , Weixiang Yu⁴¹ , Ilsang Yoon^{42,43} , and Fan Zou⁴⁴ 

(Affiliations can be found after the references)

Accepted XXX. Received YYY; in original form ZZZ

ABSTRACT

The Vera C. Rubin Observatory’s Legacy Survey of Space and Time (LSST) is expected to obtain observations of over ten million quasars. The survey’s exceptional cadence and sensitivity will enable a significant fraction of these objects to be monitored in the *ugrizy* bands, spanning observed wavelengths of approximately 0.3 – 1.0 μm . The unprecedented number of sources makes spectroscopic follow-up for the vast majority of them unfeasible in the near future, so most studies will have to rely on photometric redshift estimates which are traditionally much less reliable for Active Galactic Nuclei (AGNs) than for inactive galaxies. This work presents a novel methodology to constrain the photometric redshift of AGNs that leverages the effects of cosmological time dilation, and of the luminosity and wavelength dependence of AGN variability. Specifically, we assume that the variability can be modeled as a damped random walk (DRW) process, and we adopted a parametric model to characterize the DRW timescale (τ) and asymptotic amplitude of the variability (SF_{∞}) based on the redshift, the rest-frame wavelength, and the AGN luminosity. We constructed variability-based photometric redshift (photo- z) priors by modeling the observed variability using the expected DRW parameters at a given redshift. These variability-based photo- z (VAR-PZ) priors were then combined with traditional spectral energy distribution (SED) fitting to improve the redshift estimates from SED fitting.

Validation was performed using observational data from the Sloan Digital Sky Survey (SDSS), demonstrating significant reduction in catastrophic outliers by more than 10% in comparison with SED fitting techniques and improvements in redshift precision. The simulated light curves with both SDSS and LSST-like cadences and baselines confirm that VAR-PZ will be able to constrain the photometric redshifts of SDSS-like AGNs by bringing the outlier fractions down to below 15% from 32% (SED alone) at the end of the survey.

Key words. quasars: variability – quasars: Photometric redshifts – galaxies: active – methods: observational

1. Introduction

Active galactic nuclei (AGNs) are among the most luminous and distant objects in the universe. They are detected in wide redshift ranges, serving as light houses that can be observed over vast cosmic distances, providing crucial knowledge about the early universe (e.g. Richards et al. 2006; Fan et al. 2023). The intense luminosity of these objects arises from the accretion of matter onto the central supermassive black hole (SMBH), which is thought to exist in almost every massive galaxy. This accretion process drives the emission of intense radiation across the entire electromagnetic spectrum (e.g. see the review by Padovani et al. 2017).

Active galactic nuclei (AGNs) are intrinsically variable sources, displaying fluctuations in brightness. These variations in brightness are stochastic, occurring on timescales varying

from days to hundreds of years (Vanden Berk et al. 2004; Sartori et al. 2018). Their flux variability can be used as a tool to probe the central engine’s structure and the physical processes in its immediate environment, as reviewed in detail by Cackett et al. (2021). Various previous studies have shown that at least a portion of the observed variations originates directly from the accretion disk itself, rather than being solely a reprocessed response to the highly variable X-ray emission for both long and short timescales (e.g. Arévalo, P. et al. 2024). Several studies have made significant progress in linking the variability parameters to basic physical properties through the application of complex methods on large samples of longer and densely sampled light curves (e.g. Vanden Berk et al. 2004; MacLeod et al. 2010; Simm et al. 2016; Sánchez-Sáez et al. 2018; Li et al. 2018; Luo et al. 2020; Tachibana et al. 2020; Burke et al. 2021; Tang et al. 2023; Yu et al. 2025).

In order to place AGNs in a cosmological time frame and to study their intrinsic luminosities, black hole masses, accre-

* E-mail: 1998sarath.ss@gmail.com

tion rates and their evolution, accurate redshift measurements are necessary. Accurate redshifts are also important for studying AGN feedback and its regulatory role in star formation (Silk & Rees 1998; Fabian 2012), the coevolution of supermassive black holes and their host galaxies (Kormendy & Ho 2013), and their contribution to the ionizing background during the epoch of reionization (Madau & Haardt 2015). While current and upcoming multi-object spectroscopic surveys such as the Sloan Digital Sky Survey (SDSS)-V (York et al. 2000; Kollmeier et al. 2025), the 4-meter Multi-Object Spectroscopic Telescope (4MOST; de Jong et al. 2019), the Dark Energy Spectroscopic Instrument (DESI; DESI Collaboration et al. 2016), the Prime Focus Spectrograph (PFS; Tamura et al. 2016), and the Multi Object Optical and Near-infrared Spectrograph (MOONS; Cirasuolo et al. 2012) will significantly increase the number of observed sources with accurate redshifts, the vast majority of AGNs detected in multiwavelength all-sky and large surveys, such as WISE (Wright et al. 2010) and eROSITA (Merloni et al. 2012; Predehl et al. 2021), will remain without spectroscopic follow-up (Dahlen et al. 2013). This difference will only increase with the current deep and wide-area surveys such as Euclid (Euclid Collaboration et al. 2022) and the Vera C. Rubin Observatory Legacy Survey of Space and Time (LSST; Ivezić et al. 2019) covering unprecedented numbers (Newman & Gruen 2022; Savić et al. 2023; Li et al. 2025). Consequently, for most AGNs, photometric redshifts (photo- z) remain the only viable alternative to estimate their distances.

Spectral energy distribution (SED) fitting is a widely used method for estimating photo- z s for AGNs (e.g. Arnouts et al. 1999; Ilbert et al. 2006), as implemented in LePHARE, which has been optimized for AGNs (e.g. Shirley et al. 2025; Brammer et al. 2008; Salvato et al. 2011, 2019; Salvato et al. 2022). However, factors such as non simultaneous observations, intrinsic variability and mostly featureless SED continuum of the luminous type 1 AGNs affect the redshift estimation based on their broad spectral characteristics. However, the power-law continuum SED of luminous type 1 AGNs lacks the strong spectral features needed to accurately estimate the redshifts, with the notable exception of the Lyman break and Intergalactic medium (IGM) absorption features short of Ly α . Combined with the intrinsic variability of these objects, obtaining accurate photo- z s from SED modeling is very challenging. Arguably, large host galaxy fractions should make the task much easier, as they would add those missing features (Hsu et al. 2014; Temple et al. 2021), so the problem is most critical for luminous quasars. These issues result in many sources exhibiting very broad redshift probability distribution functions (PDFs) that are sometimes multi-peaked, leading to multiple possible redshift solutions. As reported by Saxena et al. (2024), SED fitting requires many photometric data points across a wide wavelength range to resolve these degeneracies (e.g. Brescia et al. 2019). SED fitting is able to perform well when it is combined with appropriate redshift priors, such as redshift-dependent absolute magnitude constraints and carefully selected spectral templates (Fotopoulou et al. 2012; Hsu et al. 2014; Ananna et al. 2017; Peca et al. 2021). Several studies have demonstrated that photometry from narrow and intermediate bands is more effective in this context (Salvato et al. 2009; Cardamone et al. 2010; Hsu et al. 2014; Laur et al. 2022; Lima et al. 2022). These bands are more sensitive to emission-line features, allowing for improved identification of line intensities, thereby distinguishing star-forming sources from AGNs. In addition to traditional SED fitting, machine learning (ML) has emerged as a powerful alternative for redshift estimation (e.g. Ananna et al. 2017; Saxena et al. 2024; Roster et al. 2024). ML

models can be sensitive to more subtle features than SED modeling, resulting in higher accuracy for sources that are well represented within the training dataset. However, they also tend to produce a significantly higher fraction of outliers when applied to sources outside the training sample's parameter space (Duncan et al. 2018; Salvato et al. 2022). As these methods often rely on deep and homogenized photometry with deep spectroscopy, they are typically most applicable in pencil-beam surveys.

In this work, we present VAR-PZ, a variability-based method to constrain the photometric redshifts of AGNs. It has been developed to help overcome the degeneracies of photo- z solutions, reducing the fraction of outliers and improve the accuracy of the redshifts. It models the AGN variability, including its dependence on luminosity, rest-frame wavelength, and the effects of cosmological time dilation. Assuming the variability follows a DRW process, this work adopts a parametric model from MacLeod et al. (2010, hereafter M10), where the DRW timescale and amplitude are functions of the redshift, the AGN luminosity, and the rest-frame wavelength. SED decomposition is used to estimate the intrinsic AGN luminosity and account for host galaxy contamination. These parameters are used to construct redshift-dependent priors. VAR-PZ generates PDFs, which are complementary to those derived from other established techniques and can be integrated with them, such as SED fitting and ML-based methods to improve redshift predictions. In particular, we explore the impact of using the VAR-PZ priors in two key quantities of photo- z estimates: the fraction of outliers, and the accuracy for the non-outliers.

The paper is structured as follows. Section 2 describes the data used to develop and test our algorithm. Section 3 gives an overview of our methodology and the way it is structured. A description of our code is discussed in Section 4. Section 5 illustrates the application of our algorithm in both simulated and real data for SDSS AGNs in the Stripe 82 of the Sloan survey. Section 6 discusses the potential impact of our methodology on the current LSST simulated data. Our conclusions are presented in the Section 7. We assume a flat Λ CDM cosmology with $H_0 = 69.8$ km s $^{-1}$ Mpc $^{-1}$, $\Omega_m = 0.28$.

2. Data

SDSS archival photometry data and light curves

The SDSS provides observations in the *ugriz* broad-band filters (Gunn et al. 1998), obtained using a dedicated 2.5m telescope (Fukugita et al. 1996) at the Apache Point Observatory. Our initial sample consists of 9254 spectroscopically confirmed quasars with recalibrated light curves studied by M10 from the Stripe 82 (S82) region (York et al. 2000; Annis et al. 2014). Of these, 8974 quasars are drawn from the SDSS Data Release 5 (DR5) Quasar catalog (Schneider et al. 2007), with the rest taken from the Data Release DR7 (Abazajian et al. 2009)¹. S82 is the only stripe within the SDSS footprint that offers extensive multi-epoch imaging, making it suitable for AGN variability studies. This region spans a 120°-long and 2.5°-wide stripe centered along the celestial equator. On average, this 290 deg² area has more than 60 epochs of observations per source for each filter. Since some observations were taken under non-photometric conditions, improved calibration methods have been applied to the sources in the S82 dataset by Ivezić et al. (2007) and Sesar et al. (2007). These calibrations correct for systematic errors and im-

¹ https://faculty.washington.edu/ivezic/macLeod/qso_dr7/Southern.html

prove the precision, which is particularly important for variability studies. These observations were done in annual “seasons”, lasting approximately three months, for approximately 11 years, although, the earlier seasons have very sparse sampling (~ 1 to 3 observations per season). The final three observing seasons, covering a baseline of about 800 days, provide a denser sampling of more than 10 observations per season due to the SDSS supernova survey (Frieman et al. 2008), which make them better suited for AGN variability studies. Therefore, our analysis is restricted to this well-sampled period taken later than Modified Julian Date (MJD) 53500 (after 2005).

We used the 5-band “BEST” PSF photometry from the DR7 quasar catalog (Schneider et al. 2010), calibrated following the procedure described by Richards et al. (2002). These photometric data points are corrected for galactic extinction using the u -band (A_u) galactic extinction value from Schlegel et al. (1998), while the galactic extinction values of $griz$ bands are $A_g, A_r, A_i, A_z = (0.736, 0.534, 0.405, 0.287)A_u$. To avoid the strong degeneracies in SED modeling at low redshifts, sources with $z < 0.3$ are excluded, resulting in a refined parent sample of 9210 quasars. Figure 1 shows the distribution of our parent




Fig. 1. Distribution of the parent-sample quasars in the luminosity (bolometric) and redshift space. The bolometric luminosities (L_{bol}) of the quasars in this sample were estimated by Shen et al. (2011).

sample in the redshift-luminosity plane. Note that only a small fraction of sources are found at $z < 0.3$ (0.3%), yet that is a part of the parameter space for SED modeling codes which can add significant degeneracies for the rest of the objects. For this reason, we limit our sample only to the 9210 quasars with $z \geq 0.3$.

3. Methodology

Previous studies (e.g. Collier & Peterson 2001; Bauer et al. 2009) confirm that the observed AGN variability can be statistically quantified using the Structure Function (SF). The SF is the Root Mean Squared (RMS) magnitude difference (Δm) as a function of the time difference (Δt) between observations, and is useful for analyzing sparsely or irregularly sampled light curves as it captures ensemble variability trends without relying on explicit modeling. For modeling AGN variability, the Damped Random Walk (DRW) process has been widely used to describe the SF. This process models the AGN variability as a stochastic process, defined by an exponential covariance matrix

that behaves as a random walk for short timescales and asymptotically reaches a finite SF amplitude dubbed SF_{∞} , at timescales which are comparable to or longer than the damping timescale (τ) (Kelly et al. 2009; Zu et al. 2011; Ivezić & MacLeod 2014; Suberlak et al. 2021). Although the physical origin is unclear, the DRW model can successfully describe many features of AGN light curves. For example, M10 modeled quasar light curves as a DRW process and studied whether the observed variability timescales of quasars could be linked to physical timescales in accretion disks and derived correlations between the best-fit parameters and the physical properties of the quasars. For a DRW process, the SF is described by

$$SF(\Delta t) = SF_{\infty}(1 - e^{-|\Delta t|/\tau})^{1/2}. \quad (1)$$

The observed variability of quasars is governed by their intrinsic properties such as the black hole mass (M_{BH}), the accretion rate, which is traced by the absolute i -band magnitude (M_i), and the rest-frame wavelength (λ_{RF}). Moreover, cosmological time dilation stretches the observed timescales and, consequently, the SF measured at a fixed observed-frame lag, making the variability appear redshift-dependent, although the asymptotic amplitude (SF_{∞}) remains unaffected. Additionally, the measured SF_{∞} is modulated by the AGN flux fraction in each filter j , given by $R_{\text{AGN}} = \frac{F_{\text{AGN}}^j}{F_{\text{AGN}}^j + F_{\text{Galaxy}}^j}$, since the host galaxy contribution is non-variable.




Fig. 2. Distribution of the variability parameters, SF_{∞} and τ , in the observed frame, measured from the $ugriz$ bands as a function of rest frame wavelength. The coefficients of the Eq. (3) are calculated by fitting these values, corresponding to both the amplitude and timescale of variability. The dotted line represents the linear fit with slopes -0.456 and 0.19 for SF_{∞} and τ , respectively.

To quantify how the DRW parameters vary with rest-frame wavelength (λ_{RF}) and to derive the corresponding coefficients in our variability model outlined in this section, we analyze their distributions across photometric bands. Figure 2 shows the distribution of the variability parameters SF_{∞} and τ as a function of the λ_{RF} , for each photometric band. While the scatter is large, the overall trend across filters is consistent with the expected decrease in variability amplitude at longer rest-frame wavelengths (Kubota & Done 2018) and the positive correlation in the characteristic timescale with longer rest-frame wavelengths (Frank et al. 2002). The model shown in Fig. 2 is included as a reference to illustrate the dependence of the parameters on λ_{RF} . The

amplitude is tailored to pass through the contours. The width of the contours is dependent on other physical parameters such as the M_i , etc, that are not included in the model, and therefore appears broader than the scatter associated with the wavelength dependence alone. As a result, the contour widths should not be interpreted as uncertainties of the fit. Note that while the correlation in τ is significant, it is hard to observe in the Figure owing to the wide dynamic range spanned by the measurements. The results from M10 showed that, even though τ and SF_∞ exhibit weak trends with redshift and luminosity, the most significant correlation are with M_{BH} and Eddington ratio (L_{bol}/L_{Edd}) (see their Figure 15), which is consistent with the thermal or viscous disk timescales. They accounted for the uncertainties in M_{BH} (~ 0.4 dex; Shen et al. 2024) and demonstrated that ignoring this uncertainty can lead to underestimation of the M_{BH} dependence. Additionally, since the L_{bol}/L_{Edd} depends on both L_{bol} and M_{BH} , it is probable that the observed trends in SF_∞ reflect an underlying dependence on the L_{bol}/L_{Edd} rather than luminosity or mass alone. However, because M_{BH} estimates are often unavailable or uncertain in large surveys, they also provided reduced models excluding it. Furthermore, DRW is a stationary Gaussian process, and Stone et al. (2022) showed that when the baseline extends to ~ 20 years, the best-fit τ continues to increase, suggesting that AGN variability may not be strictly stationary over multi-year or decade timescales. Thus, the DRW model should be regarded as a convenient simplification. The DRW parameters τ and SF_∞ are modeled by M10 as

$$\log f = A + B \log \left(\frac{\lambda_{RF}}{4000 \text{ \AA}} \right) + C(M_i + 23) + D \log \left(\frac{M_{BH}}{10^9 M_\odot} \right) + E \log(1 + z), \quad (2)$$

where f is either SF_∞ or τ , M_i , and A, B, C, D , and E are fit to the ensemble of objects. In our specific case, to model the AGN variability, we used the best-fit relation which does not have an explicit dependence over the M_{BH} , resulting in a simplified form solely dependent upon M_i and λ_{RF} . Specifically, we follow M10 and set $D = 0$ and $E = 0$, such that

$$\log_{10} f = A + B \log_{10} \left(\frac{\lambda_{RF}}{4000 \text{ \AA}} \right) + C(M_i + 23). \quad (3)$$

Note that although we used Eq. (3) for the modeling to estimate redshift priors through VAR-PZ (see below), the relation (Eq. 2) which incorporates the dependency on the M_{BH} was used only to simulate the realistic SDSS and LSST light curves (see Section 5.1). The A, B and C coefficients in this work are re-estimated to overcome the potential biases due to different algorithms for light curve modeling and in estimating the AGN continuum luminosities. Specifically, the M_i and R_{AGN} values are obtained through SED modeling of the SDSS photometry using the Low Resolution Templates (LRT) and algorithm of Assef et al. (2010, hereafter A10, see Section 4). We find the best-fit values of the A, B and C coefficients in Eq. (3) following a similar approach to M10. We first estimate B by fitting a power-law of the form $f \propto \left(\frac{\lambda_{RF}}{4000 \text{ \AA}} \right)^B$ for both SF_∞ and τ parameters to every quasar with at least two observations in multiple bands. We adopt as the value of B for all subsequent calculations as the median of all objects for which we were able to carry out this fit, corresponding to -0.456 ± 0.03 and 0.19 ± 0.01 for SF_∞ and τ , respectively (see Sect. 5.1 in M10). We then fix B and fit the values of coefficients A and C to the full ensemble of objects. This method of fixing coefficient B before other coefficients is done

to eliminate the degeneracies between the λ_{RF} and other physical parameters, as B is the only parameter in a term that depends on wavelength. This process effectively decouples B from the rest of the parameters and minimizes the covariances.

Table 1 shows the values of the coefficients A, B, C and D of the SF_∞ and τ relations in Eq. (2), respectively, in comparison to the values estimated by M10. The small variations in the coefficients with respect to those from M10 are likely due to methodological changes in both light curve modeling and M_i estimation.

In order to model the AGN variability as a DRW process and to obtain the photo- z prior, we used Gaussian Process (GP) along with the Maximum Likelihood Estimation (MLE) framework. GP is a flexible non-parametric method that can model the covariance structure of time series data and MLE is a statistical approach for estimating the likelihood of the fitted parameters, respectively. We use Eq. (3) to calculate the pair of DRW parameters (SF_∞ and τ) as a function of redshift for each photometric band using the M_i and R_{AGN} values obtained from the SED modeling at those redshifts. These redshift-dependent DRW parameters are used as inputs to the GP-based DRW model, where for each trial redshift, the corresponding DRW parameters are applied to fit the light curve in each photometric band. The MLE quantifies how likely are those DRW parameters to result in the observed light curve, providing the likelihood of that trial redshift. The resulting likelihood distributions from all bands are then multiplied to construct a prior PDF for constraining photo- z s. This PDF (hereafter VAR-PZ PDF) is then combined with the PDF derived from the SED modeling (hereafter SED-PDF) to produce a more accurate posterior probability distribution. This approach enables us to evaluate the contribution of the VAR-PZ PDF in improving the photo- z predictions.

Figure 3 schematically illustrates a characteristic example of this approach for an S82 quasar (“SDSS J000013.80-005446.8”, $z_{spec} = 1.8361$), demonstrating how VAR-PZ constrains the photo- z estimation. The top panel shows the SED fitting results to the stacked SDSS photometry for different trial redshifts (0.5, 1.0, 1.5, 1.8, and 2.0), with each redshift represented by a different color along with their corresponding χ^2 fit values. Note that for different redshifts, the best-fit AGN luminosities (parametrized by M_i) and the fraction of the flux coming from the AGN component (R_{AGN}) can vary significantly, which in turn affects the expected variability properties (see Section 4). In this example, the SED fitting strongly favors a low-redshift solution. The middle panel presents the steps occurring within VAR-PZ for constraining the photo- z s. The left panels show the M_i and R_{AGN} values estimated through the SED modeling as a function of redshift, while the right panels show the expected SF_∞ and τ values expected as a function of redshift to estimate the DRW model likelihood. When this technique is applied across the full redshift grid, it generates the PDFs shown in the bottom panel. Vertical dashed lines mark the trial redshifts from the top panel. The first subplot shows VAR-PZ PDF across all five bands, which strongly favors higher redshifts. The second subplot displays the PDF derived solely from SED fitting, which peaks at low redshift. The third subplot presents the posterior obtained through the combination of the SED and variability PDFs. This combined approach effectively discards the low-redshift solutions preferred by SED fitting alone, with the variability constraints driving the posterior toward the true spectroscopic redshift. This example demonstrates how VAR-PZ breaks degeneracies inherent in SED-only photo- z estimation and improves redshift accuracy.

Table 1. Best-fit coefficients for Eq. (3) from M10 and this work.

f	A	B (λ_{RF})	C (M_i)	D (M_{BH})	Reference
SF_∞	-0.51 ± 0.02	-0.479 ± 0.005	0.131 ± 0.008	0.18 ± 0.03	M10
SF_∞	-0.618 ± 0.007	-0.479 ± 0.005	0.090 ± 0.003	0	
τ	2.4 ± 0.2	0.17 ± 0.02	0.03 ± 0.04	0.21 ± 0.07	
τ	2.2 ± 0.1	0.17 ± 0.02	-0.01 ± 0.02	0	
SF_∞	-0.695 ± 0.001	-0.456 ± 0.03	0.086 ± 0.001	0.012 ± 0.001	This Work
SF_∞	-0.713 ± 0.003	-0.456 ± 0.03	0.071 ± 0.002	0	
τ	2.2 ± 0.004	0.19 ± 0.01	-0.027 ± 0.002	0.027 ± 0.003	
τ	2.2 ± 0.019	0.19 ± 0.01	-0.026 ± 0.011	0	

Notes. Coefficients are provided for both variability parameters, SF_∞ and τ , with and without M_{BH} dependency. The recalibration of these coefficients is done to avoid inconsistencies arising from using different M_i estimates and light curve modeling algorithms.

4. Implementation

In this section, we explain the VAR-PZ routine, detailing how the code models AGN variability using the DRW process and incorporates it into the photo- z estimation. We employed the Python-based Gaussian Process (GP) toolkit *Celerite* (Foreman-Mackey et al. 2017) to model AGN light curves as a DRW process. *Celerite*'s computationally efficient GP implementation enables faster modeling compared to traditional methods such as *Javelin* (Zu et al. 2013). It provides the flexibility to use MLE or Markov Chain Monte Carlo (MCMC), depending on our specific needs. Although, MCMC provides a thorough exploration of parameter space, it is often computationally more expensive compared to methods which can leverage MLE. We developed a custom DRW kernel to use *Celerite* for our goals. Unlike the default *Celerite* kernel (*RealTerm*) used for DRW fitting, which models the SF_∞ and τ independently, our kernel jointly fits for both of them using the relations outlined in Eq. (3) with the coefficients in Table 1.²

The following stage of the implementation details how the kernel integrates the redshift-dependent DRW parameters into the VAR-PZ framework to generate the variability prior. The kernel involves two main stages. An initialization computes the value of M_i for the AGN component as well as R_{AGN} to account for the host fractions in the photometric band of the light curve being modeled, as a function of redshift in a grid between $z=0$ and $z=5$ with a bin size of 0.01. These values are obtained by modeling the observed SED adopting each redshift of the grid. We nominally interpolate between redshifts for flexibility, but in practice we minimize the use of interpolation by matching the redshift grids of the VAR-PZ and SED PDFs. In the next stage, the kernel computes the expected SF_∞ and τ values based on the estimated M_i values for each band in the light curve, using Eq. (3). The value of τ is then modified to account for cosmological time dilation while the SF_∞ is scaled by the AGN flux fraction (R_{AGN}) to account for the contribution of the host galaxy. This renormalization approach helps to account for the damped variability observed in intermediate-type AGN where host galaxy contamination significantly affects the measured variability amplitude. While this work does not explicitly model the physical parameters of the host galaxy and their connection the AGN properties, future refinements could incor-

² While *Celerite* internally models variability using the variance parameter σ^2 , we apply a transformation to recover SF_∞ via the relation $\text{SF}_\infty = \sqrt{2}\sigma$, as commonly adopted in variability analyses.

porate additional host galaxy properties, such as the $M_{\text{BH}} - M_*$ relation, which influences the $L_{\text{bol}}/L_{\text{Edd}}$, which could potentially provide stronger constraints on the VAR-PZ PDF.

We then compute the likelihood distribution as described earlier, independently for each band and combine them to produce the VAR-PZ priors. To account for additional uncorrelated white-noise components in the light curves, we included a *JitterTerm* kernel in our model. This term reduces the effect of photometric noise on the AGN variability parameters fit. We describe the workflow in detail in Appendix A.1. We have made our implementation of the method publicly available here³ along with the required documentation.

5. Application to SDSS quasars

As a proof of concept, we apply our methodology to SDSS quasars in S82, first with simulated light curves and then with the observed light curves to estimate the improvement in photo- z s obtained by using the VAR-PZ priors. To assess performance of our photo- z s, the outlier fraction and precision of the estimates are calculated. Following Salvato et al. (2019), we define outliers as objects for which

$$\frac{|z_{\text{phot}} - z_{\text{spec}}|}{1 + z_{\text{spec}}} > 0.15$$

where z_{phot} is the photometric redshift estimate, and z_{spec} is the spectroscopic redshift. We evaluate the precision of the estimates through the Normalized Median Absolute Deviation (NMAD; e.g. Salvato et al. 2019) of the non-outlier objects. The NMAD is defined as

$$\sigma_{z_p} = 1.48 \times \text{median} \left(\frac{|z_{\text{phot}} - z_{\text{spec}}|}{1 + z_{\text{spec}}} \right).$$

5.1. Simulated light curves

We start by looking at the simplest scenario by simulating the light curves of our parent sample of S82 quasars assuming a DRW model and Eq. (2). Note that, we use this Eq. (2) that incorporates the M_{BH} dependency to simulate the light curves rather than Eq. (3), which is the one used to create the VAR-PZ priors, as it should provide a better approximation to the real

³ <https://github.com/SarathSS98/VAR-PZ/tree/main>

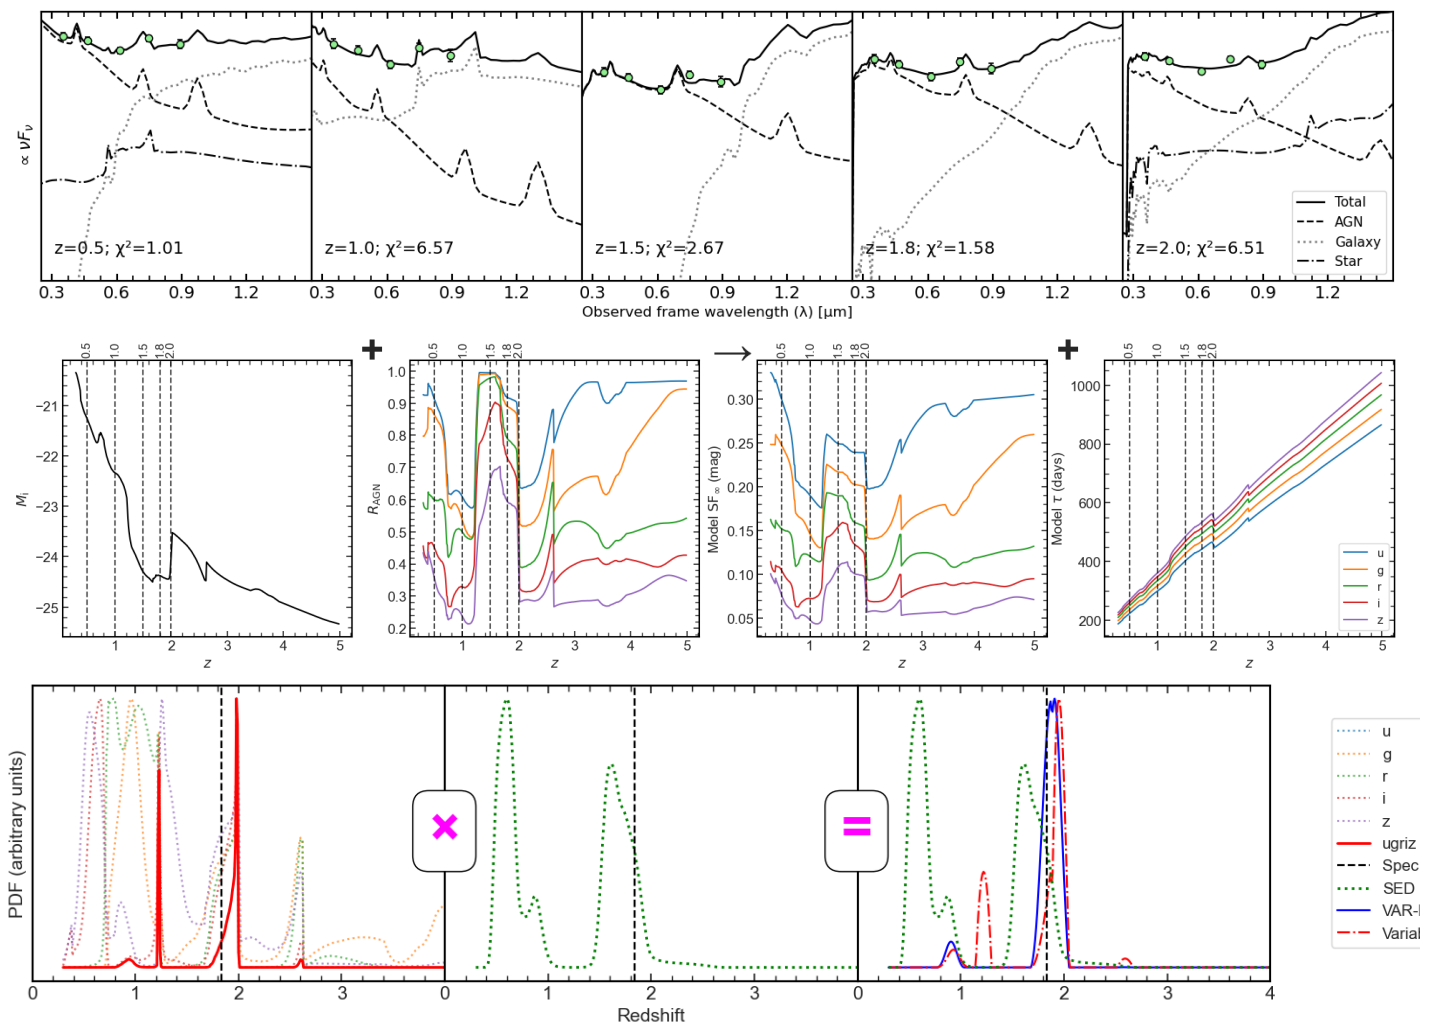


Fig. 3. *Top panel:* SED fits at trial redshifts $z = 0.5, 1.0, 1.5, 1.8,$ and 2.0 with corresponding χ^2 values. *Middle panel:* VAR-PZ redshift prior estimation. Left: M_i and R_{AGN} components estimated from the SED modeling as a function of redshifts. Right: Expected SF_∞ and τ using the Eq. (3) for the corresponding redshifts. *Bottom panel:* Photo- z priors from VAR-PZ along with SED fitting PDF, and the combined posterior. Vertical dashed lines mark the trial redshifts from the top panel.

SDSS light curves (see M10 for details). The light curves are simulated through *Celerite* using the implementation of Burke et al. (2021) for a grid of observing seasons and cadences, incorporating realistic photometric uncertainties and realistic seasonal observing gaps. The simulated light curves span a wide grid of cadence intervals ranging from 1 to 10 days, and baseline lengths ranging from 1 to 10 seasons. Each S82 observing season was assumed to last 90 days, with seasonal gaps of 275 days. Realistic SDSS photometric uncertainties are added to our simulated light curves according to the brightness of each target. This approach enables an isolation of the effect of varying cadences and baselines.

According to Suberlak et al. (2021), accurate recovery of the damping timescale, τ requires a light curve baseline at least three times longer than this timescale. Therefore, depending on the length of the baseline, VAR-PZ priors are constructed using the method outlined earlier in the range of redshifts where the baselines are at least three times longer than the expected value of τ . The remaining part of the PDF is simply flat, with a value given by the maximum of the PDF, where the baseline $\geq 3\tau$. This strategy enables an efficient application of the variability constraints only in the part of the parameter space where there is constraining power for our method.

Figure 4 presents a heatmap, which quantifies the impact of priors generated in this range of the baselines and observational cadences on both the outlier fraction and redshift precision of the photometric redshifts obtained using LRT (see Section 3). For reference, without using priors, LRT obtains $\sigma_{zp} = 0.0554$ and an outlier fraction of 0.33. As expected, the largest improvements are achieved with a combination of long baselines and high observational cadences. While the improvements in the precision are modest, the reduction in the outlier fraction is substantial. In fact, for the highest cadence and largest number of seasons probed, VAR-PZ is able to constrain photo- z s, with an accuracy (σ_{zp}) = 0.0435 and an outlier fraction of 0.06.

In contrast, with limited temporal coverage, specifically, fewer than two seasons and a cadence ≥ 7 days, the addition of VAR-PZ does not improve upon SED-only methods for redshift estimation, and thus requires careful application to avoid degrading performance. The diagonal structure of the heatmap reveals a partial degeneracy between cadence and baseline; that is, higher cadence can partially compensate for shorter baselines and vice versa, underlining potential observational strategies.

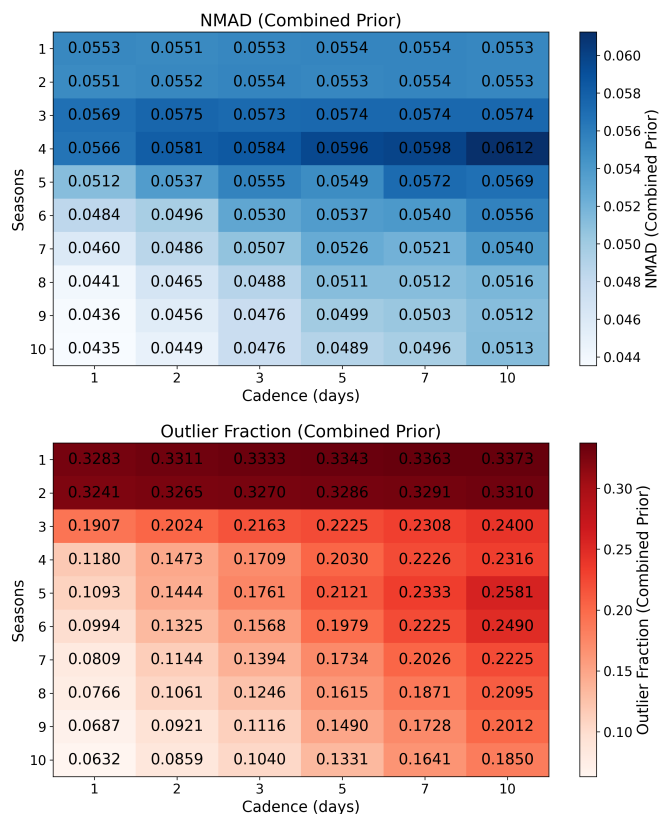


Fig. 4. Heatmap illustrating the influence of observational cadence and baseline length (in SDSS seasons) on photometric redshift estimation metrics: the NMAD (top) as a measure of precision, and the outlier fraction (bottom). The color bar, presented on a logarithmic scale, maps these metrics such that regions depicting lower values correspond to the most promising observing conditions, yielding higher precision and a reduced fraction of outliers in photometric redshift predictions.

5.2. Observed light curves

We applied our methodology to the real SDSS Stripe 82 light curves (see Section 2). To ensure the use of reliable and high-quality data, the light curves were filtered to exclude flux measurements with a signal-to-noise ratio (S/N) below 3 relative to the zero-flux level, thereby ensuring that the median flux of each light curve is not biased by measurements at the faint end.⁴ Using this median as a reference, we then exclude observations which deviate by more than 2.5σ , thereby eliminating strong outliers. Finally, a quality cut is applied by retaining only epochs with $\text{SNR} > 10$. This filtering ensures that only high-quality measurements are included in the light curves, though we note that they may also introduce biases in the measured light curve properties. We only use light curves for a given band to contribute to the VAR-PZ prior if they have more than 35 measurements (which corresponds to the peak of the data point distribution in the light curves). Out of 9210 total sources, 6493 met these criteria in all five bands, and, 8413 do so in at least one band. We remove the remaining 797 sources from our analysis. As discussed in the previous section, the relatively short temporal baselines of the SDSS light curves restrict their constraining power at higher redshifts. Figure 5 illustrates the distribution of median τ values, estimated using Eq. (2), as a function of redshift for our AGN sample. The black dashed horizontal line represents one-third of the SDSS S82 baseline length. This highlights a

⁴ <https://www.sdss4.org/dr17/algorithms/magnitudes/>

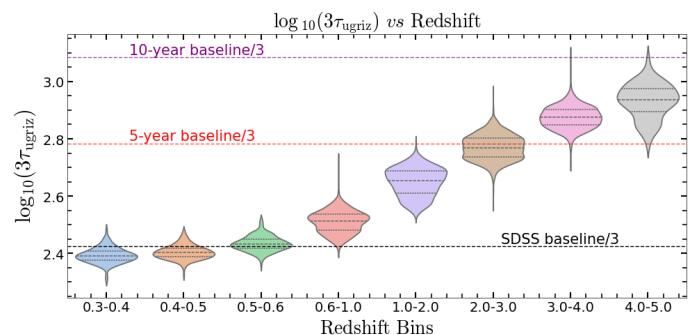


Fig. 5. Violin plot showing the median AGN variability timescale (τ) estimated from M10 over the SDSS bands as a function of redshift. The dashed black line represents one-third of the SDSS baseline duration, showing the redshift-dependent threshold beyond which the condition $3\tau < \text{baseline}$ is no longer satisfied. This requirement defines the redshift range over which VAR-PZ can constrain redshifts with the current data. The red and purple line represent one-third of the 5 and 10 year LSST baseline, respectively. LSST’s enhanced sensitivity will enable the detection of fainter objects, resulting in lower overall timescale values (τ) and consequently providing greater constraining power.

redshift-dependent limitation: beyond $z \sim 0.5 - 0.6$, the S82 data do not provide a sufficiently long baseline to reliably constrain the DRW parameters, and consequently, the accuracy of the redshift (VAR-PDF) priors diminishes at higher redshifts. Hence, we followed the same implementation strategy mentioned before to produce the posterior distribution function, namely, we replace all likelihood values for redshifts where the expected τ from Eq. (3) exceeds one-third of the maximum likelihood observed for those points. The red and purple lines in Figure 5 represent one-third of the 5 and 10 year LSST baseline, respectively, compared to these objects, although note that the higher photometric sensitivity of LSST will extend AGN detection to much fainter flux limits compared to SDSS, systematically shifting the population toward lower τ values. The increased survey depth will yield improved constraints for more AGNs even when considering similar shorter observational baselines in the initial years of the survey.

Figure 6 compares the photometric and spectroscopic redshifts before and after applying the variability priors. The inclusion of VAR-PZ priors significantly improves the SED-derived probability distributions by reducing catastrophic failures and resolving degeneracies inherent to SED fitting alone. The accuracy estimates of this method on the SDSS observations are shown in Table 2. For our sample, the inclusion of variability reduces the outlier fraction from 32% to 22%. This effect is particularly pronounced in the regime where SED fitting predicts low redshifts ($z_{\text{SED}} < 0.6$), where variability plays a critical role in mitigating catastrophic failures: the outlier fraction in this regime drops from 81% to 33%, as shown in Figure 7. Note that, as described in Section 2, we restricted our analysis to sources with $z > 0.3$ to avoid significant degeneracies arising from the LRT SED modeling. Although this removes the low-redshift regime, it represents only $\sim 0.3\%$ of the parent sample and therefore has a negligible impact on the overall statistics. Retaining these low- z sources would lead to the overestimation of the VAR-PZ’s performance, as the LRT SED modeling algorithm tends to produce degenerate redshift solutions in this low- z regime, resulting in an outlier fraction of $\sim 47\%$ for the SED-only case. Excluding them thus ensures that the reported improvement of VAR-PZ accurately reflects its true performance rather than being driven by SED-induced degeneracies.




Fig. 6. Binned scatter diagrams comparing photometric redshifts derived from SED fitting using LRT, independently (top) and combined with our variability model (bottom), against spectroscopic redshifts. The dashed line represents the one-to-one correspondence between the axes.

As an independent test, we study the impact of the VAR-PZ approach by applying its priors to the SED PDFs computed by the more sophisticated LePHARE SED fitting framework (Arnouts et al. 1999; Ilbert et al. 2006) which has been optimized for AGNs by (Shirley et al. 2025) to estimate the improvements. Note, however, that the computation of M_i and R_{AGN} for VAR-PZ is still computed using the LRT SED modeling as the host-AGN decomposition is not currently computed by LePHARE. When applying the VAR-PZ priors, the outlier fraction decreases from 23% to 19% overall for the sample, while for sources where LePHARE predicts a $z < 0.6$ the outlier fraction improves from 66% to 37%, further demonstrating the robustness of variability as a complementary prior. We note that a, more pronounced improvement by the addition of VAR-PZ, potentially even comparable to those seen in LRT, could be possible, if the M_i and R_{AGN} values were derived self-consistently with LePHARE. While introducing luminosity function (LF) priors on the host galaxies (see A10 for details on the implementation) can overcome some




Fig. 7. Distribution of photometric redshifts for sources with $z_{\text{phot}} < 0.6$, comparing LRT SED-only (gray) and VAR-PZ-enhanced (red) solutions, illustrating how variability priors redistribute sources and remove catastrophic photo- z failures in this regime. The blue histogram represents the spectroscopic redshifts of those sources.

of the degeneracies resolved by VAR-PZ, they often compromise the accuracy of redshift estimates for sources that are otherwise well-constrained, highlighting the need for a more sophisticated LF prior which closely matches the distribution of quasar host galaxies.

The short baseline of the SDSS observations limits the achievable improvements, so in the next section we study the potential impact when considering the observations of the upcoming Rubin LSST, which will provide longer and denser light curves spanning a decade (Ivezić et al. 2019). Note that while, in principle, data from Zwicky Transient Facility (ZTF) could be used to extend the temporal baseline of the sources in S82, their relatively large photometric uncertainties (with a median value of about ~ 0.1 mag for the objects in our sample) compared to SDSS (~ 0.02 mag) result in only small improvements. In a similar context to our study, Suberlak et al. (2021) found that increased noise effectively offsets the potential benefits of the longer time coverage.

6. Prospects for Rubin LSST

Rubin LSST is expected to revolutionize our understanding of AGNs. With between 50 and 200 visits per source in each of the “ugrizy” filters, the ten-year survey will observe at least 10 million AGNs spanning $18,000 \text{ deg}^2$ of the sky up to redshift $z \sim 7.5$, with average luminosities of $L_{\text{bol}} \sim 10^{44} \text{ erg s}^{-1}$ (LSST Science Collaboration 2009; Li et al. 2025).⁵ We also expect to detect an additional $\sim 40,000$ fainter AGNs with more than 1000 samplings per band in the Deep Drilling Fields (DDFs), a smaller sky area (60 deg^2) with very high cadence (Brandt et al. 2018). LSST’s long-term monitoring will provide a complete census of AGNs, and its exceptional cadence with small photometric uncertainties, will greatly enhance variability studies. This coverage should be ideal for VAR-PZ to improve the accuracy of the redshift predictions.

⁵ <http://lsst.org>

Table 2. Comparison of photo- z performance for SED and SED + variability models.

Method	Total Sources				
	Number of Sources	Number of Outliers	Outlier Fraction	NMAD Without Outliers	NMAD With Outliers
LRT SED	8413	2745	0.3263	0.0558	0.1051
LRT SED + Variability	8413	1918	0.2280	0.0579	0.0850
LePHARE SED	8413	1954	0.2323	0.0483	0.0705
LePHARE SED + Variability	8413	1642	0.1952	0.0485	0.0663
$z_{\text{SED}} < 0.6$					
LRT SED	1300	1060	0.8154	0.0410	0.6105
LRT SED + Variability	1300	429	0.3300	0.0579	0.1077
LePHARE SED	734	491	0.6689	0.0379	0.6059
LePHARE SED + Variability	734	278	0.3787	0.0499	0.1242

Notes. NMAD represents the Normalized Median Absolute Deviation. Values are shown for the entire AGN sample and for the subsample with SED photo- z s less than 0.6.

We used the expected bandpasses provided by the observatory⁶ to simulate realistic photometry. The synthetic stacked LSST photometry were simulated by using the best-fit SED model to the observed SDSS *ugriz* photometric data obtained with LRT from A10 at their corresponding spectroscopic redshifts convolved with the LSST bandpasses. Ivezić et al. (2019) provide comprehensive LSST performance measurements. The anticipated photometric error in magnitudes for a single visit can be expressed as

$$\sigma_{\text{source}} = \sqrt{\sigma_{\text{rand}}^2 + \sigma_{\text{sys}}^2} \quad (4)$$

where σ_{sys} is the systematic photometric error (caused, for instance, by imprecise PSF modeling, but excluding uncertainties in the absolute photometric zero-point) and σ_{rand} is the random photometric error (see Ivezić et al. (2019) for details). The survey's calibration system is developed to limit the systematic error to $\sigma_{\text{sys}} < 0.005$ mag. For point sources, Ivezić et al. (2019) adopts the following expression based on Sesar et al. (2007) for the σ_{rand} in magnitudes,

$$\sigma_{\text{rand}}^2 = (0.04 - \gamma)x + \gamma x^2 \quad (5)$$

where $\log_{10}(x) = 0.4(m - m_{5\sigma})$ with m being the observed magnitude and $m_{5\sigma}$ being the 5σ depth of the coadded magnitude limit of the LSST filters⁷. The parameter γ is filter-dependent, with $\gamma_u = 0.038$ for the *u*-band and $\gamma = 0.039$ for the *g*, *r*, *i*, *z*, and *y* bands (Ivezić et al. 2019). Note that, LSST observations in different bands are not simultaneous, and the associated variability-induced color offsets are not included in our simulations. Realistic LSST light curves are simulated for the parent sample using the current LSST baseline v5.0.0 of the Operation-Simulator (Op-Sim)⁸ using the Metrics Analysis Framework (MAF; Marshall et al. 2017)^{9 10}.

⁶ <https://github.com/lstt/throughputs/tree/main/baseline>

⁷ <https://rubinobservatory.org/for-scientists/rubin-101/key-numbers>

⁸ https://github.com/lstt/rubin_sim

⁹ <https://www.lsst.org/scientists/simulations/maf>

¹⁰ https://github.com/lstt/rubin_sim_notebooks/tree/main

Light curves were generated as a function of the LSST yearly data release. Figure 8 illustrates the simulated 10-year DRW light curve for an AGN (SDSS J000008.13+001634.6), in the Wide-Fast Deep (WFD)¹¹ survey. The improvement in the photo- z performance of LRT obtained by adding the VAR-PZ priors are presented in Figure 9. As expected, the photo- z precision improves and the outlier fraction decreases significantly over time, with the most limited performance seen during the first year. The figure also includes a comparison with photo- z estimates derived using LePHARE SED templates, providing an independent validation. Note that, to avoid inconsistencies, the SED PDFs are obtained by fitting the Sloan *ugriz* filters, instead of the simulated LSST stacked photometry (which relies on the LRT fits). Importantly, both LePHARE and LRT based routines, incorporating VAR-PZ variability priors, show consistent trends of improvement with increasing temporal baseline, further confirming the robustness of variability-informed photo- z estimation. To assess the performance of VAR-PZ for the fainter AGNs that LSST will discover, we constructed a faint sample by dimming our original sources by a factor of three in flux. This setup approximates the parameter space where LSST will detect sources that are about three times fainter than those in SDSS, but with comparable *S/N* owing to Rubin observatory's sensitivity and depth. A large fraction of faint AGNs detected by LSST will be host-dominated, for which SED-based photo- z s are expected to be relatively accurate even without variability information. In this experiment, however, we restrict ourselves to faint quasars. For this faint sample, we simulated LSST light curves at the fainter magnitudes while keeping R_{AGN} fixed to the values inferred for the original sample. We then combined the resulting VAR-PZ priors with the SED PDFs. For the faint quasars, our method reduces the outlier fraction from 32% (SED alone) to 17%, with an NMAD of 0.04 when computed excluding outliers.

Additionally, we also test the level of improvement that would be obtained if these objects were instead within the Deep Drilling Fields (DDFs). Specifically, we assume the observing cadence of the from the XMM-Large Scale Structure Field (XMM-LSS¹²; $\alpha_{J2000} = 35.57^\circ$, $\delta_{J2000} = -4.82^\circ$). The DDF cadence yields improved photo- z performance, with re-

¹¹ <https://survey-strategy.lsst.io/baseline/wfd.html>

¹² <https://survey-strategy.lsst.io/baseline/ddf.html>

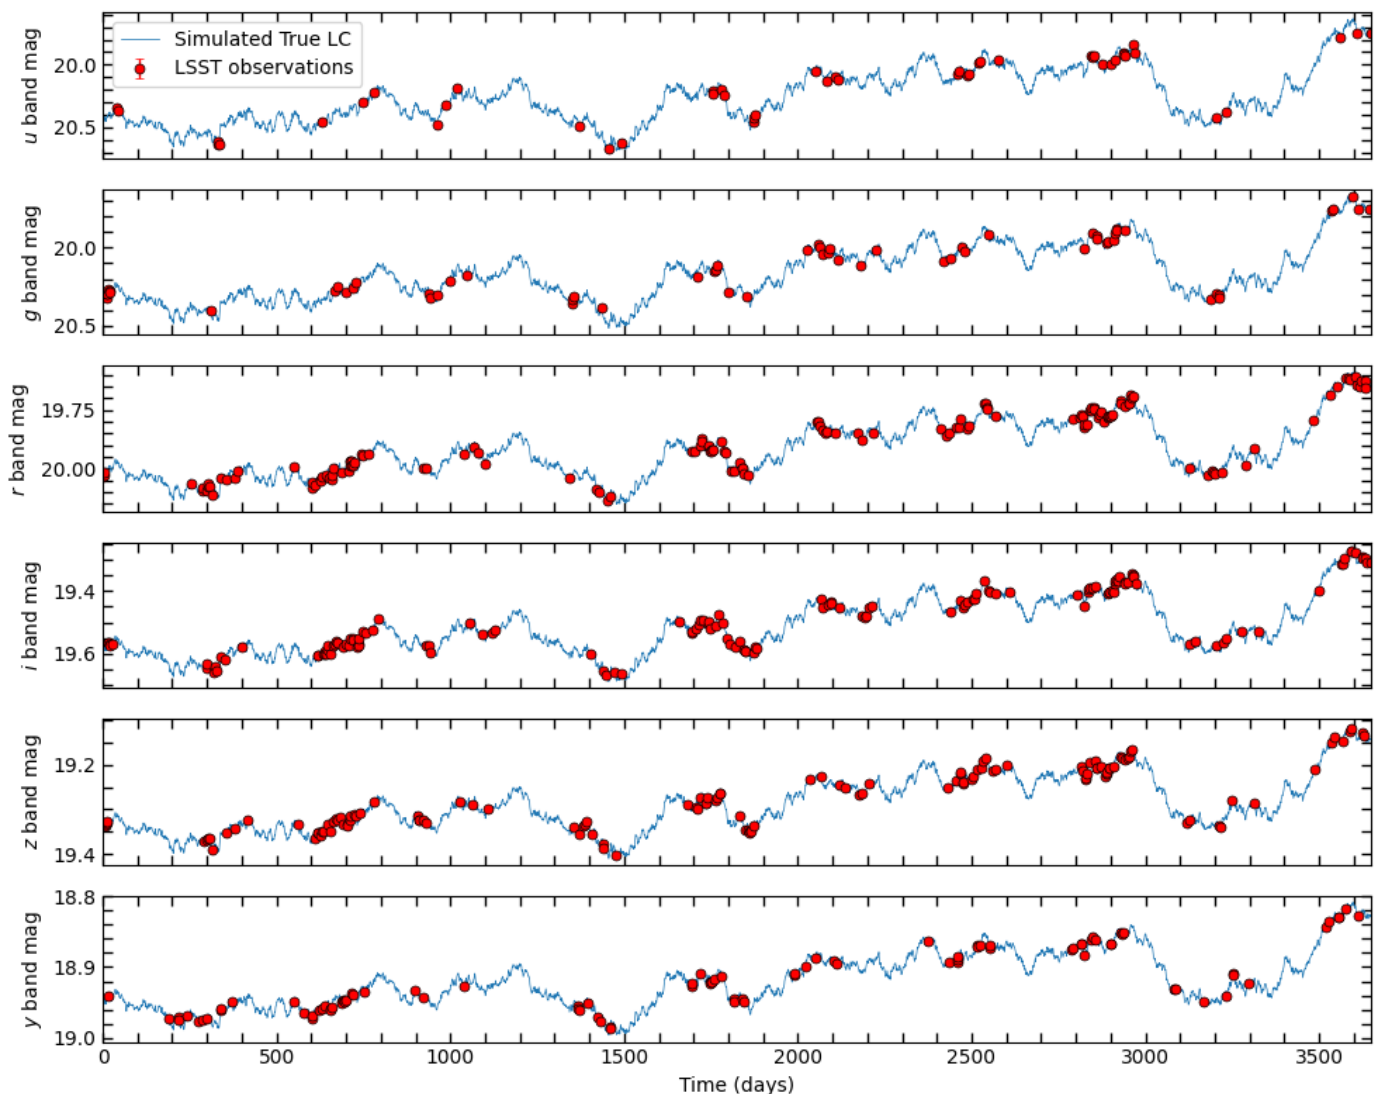


Fig. 8. Simulated LSST light curve for a source using DRW parameters fit (SDSS J000008.13+001634.6, $z=1.836$), as observed under the LSST Wide-Fast-Deep (WFD) survey strategy over the 10-year baseline. The underlying blue curve represents the light curve sampled with a uniform 1-day cadence. Overlaid red points correspond to LSST observations in all six bands (*ugrizy*), incorporating realistic survey cadence and photometric uncertainties based on the LSST observing strategy.

duced outlier fractions and improved precision compared to the WFD strategy. Nevertheless, the WFD cadence still provides reasonably accurate redshift estimates for AGN populations, even though its typical outlier rate ($\sim 15\%$) does not reach the higher fidelity achieved with the DDF fields ($\sim 9\%$).

7. Summary

In this work, we demonstrate that AGN variability can be effectively used to improve photometric redshift estimates. Specifically, we implement a method we refer to as VAR-PZ to generate redshift PDFs that can be applied as priors to, for example, photo- z estimates based on SED modeling, serving not as a stand-alone solution for accurate AGN photo- z s but as a complementary tool to existing techniques.

This work serves as a proof of concept, motivating the development of other physically motivated variability models, which would likely enable tighter redshift constraints. As detailed in Section 3, we model the observed stochastic variability of AGNs as a DRW process following the dependence of the DRW pa-

rameters on the physical properties of the AGNs proposed by M10. The parameters from M10 relation (Eq. (3)) are refitted to account for differences in light curve modeling and the AGN continuum luminosities, specifically M_i which is derived from SED modeling, along with host dilution fractions (R_{AGN}). While the DRW is a reasonable model for quasar variability, it should be emphasized that it is only an approximate description. Regardless, we show that it can yield broad but informative photo- z PDFs.

To generate variability priors from VAR-PZ, we first estimate the values of the expected DRW parameters in each band as a function of redshift based on Eq. (3) in a redshift grid, performing a full SED decomposition to obtain M_i and R_{AGN} , following the methodology outlined in Section 4. We then compute the likelihood of those parameters being able to describe the observed light curves, and express the likelihood as a function of redshift as a PDF. We refer to this PDF as the VAR-PZ PDF.

To assess the effect of survey cadences and baselines, we took all AGN with well sampled light curves in the SDSS S82 (see Section 2) and we simulated SDSS-like light curves across

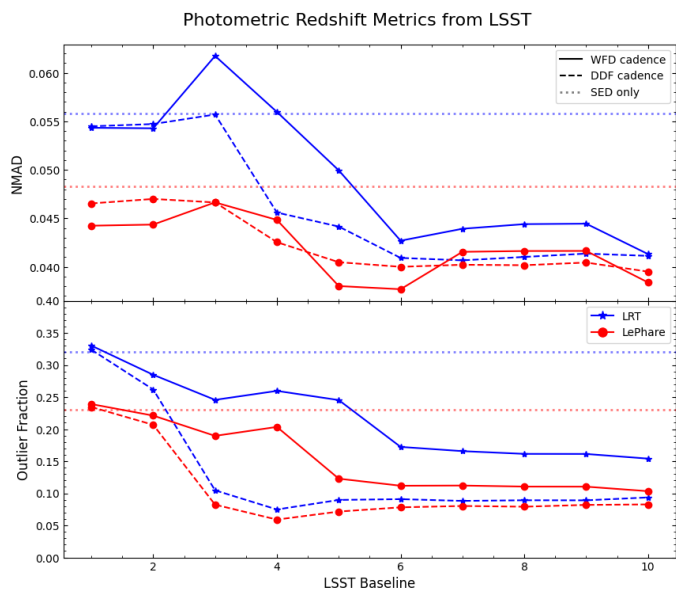


Fig. 9. Evolution of photo- z precision and outlier fraction for our parent sample as a function of the LSST temporal baseline. The panels also include a comparison with the results obtained using the LSST DDF cadence. This figure illustrates the improvements achieved by incorporating the VAR-PZ variability priors into both the LRT and LePHARE SED fitting routines, highlighting the impact on two independent methods. The LePHARE SED fitting is performed using the Sloan *ugriz* filters, whereas both LRT and VAR-PZ utilize synthetic LSST *ugrizy* filters.

a range of cadences and baselines (see Section 5). Figure 4, shows that the performance of photo- z s improve with longer and denser sampling. The results of SDSS observations (Section 5.2), presented in Figure 6 and Table 2, demonstrate that the inclusion of VAR-PZ priors significantly improves redshift estimation. In comparison to the LRT SED modeling method by itself, the outlier fraction decreases from 27% to 19% for the full sample, and from 81% to 33% for sources with $z_{\text{SED}} < 0.6$. When implementing VAR-PZ within the LePHARE SED-fitting framework, the performance is also improved albeit by a smaller factor, with outlier fractions reduced from 19% to 16% for the full sample and from 66% to 37% for low-redshift solutions. At higher redshifts, the limitations of SDSS observations, particularly the relatively short temporal baselines, reduce the reliability of variability-based redshift estimates (Figure 5). In these regimes, the DRW timescales fall below one-third of the available light curve baseline, leaving the variability prior essentially unconstrained (See Suberlak et al. (2021), Section 2.3). This highlights that the gain from variability priors strongly depends on light curve baseline. Complementary simulations using Rubin LSST observing strategies (WFD and DDF) further confirm that LSST’s increased temporal coverage and cadence substantially enhance photometric redshift accuracy. These findings, as shown in Figure 9, demonstrate a reduction in the outlier fraction from $\sim 32\%$ for SED-only estimates to $\sim 15\%$ for WFD and $\sim 9\%$ for DDF cadences by the end of the survey. Additionally, for faint LSST objects, the WFD cadence yields an outlier fraction of $\sim 17\%$ in separate tests. This underscores the importance of survey strategy and establishes the significance of incorporating variability-based priors into hybrid photometric redshift frameworks for upcoming large-scale AGN surveys.

Acknowledgements. We thank the anonymous referee for the thoughtful comments that helped us improve the paper. We gratefully acknowledge the support of the ANID BASAL project FB210003 (S.S.S., R.J.A., T.A., F.E.B., C.M.,

T.M. and C.R.), FONDECYT Regular 1231718 (S.S.S. and R.J.A.), 1240105 (T.A.), 1241005 (F.E.B.), 1230345 (C.R) and ANID Millennium Science Initiative AIM23-0001 (T.A., F.E.B.). S.S.S. acknowledges the travel support from the Royal Astronomical Society to attend the meeting ‘Supermassive Black Hole studies in the Legacy Survey of Space and Time’ in Durham, UK. T.T.A acknowledges support from NASA ADAP Grant 80NSSC24K0692. A.B. acknowledges support from the Australian Research Council (ARC) Centre of Excellence for Gravitational Wave Discovery (OzGrav), through project number CE230100016. DD acknowledges PON R& I 2021, CUP E65F21002880003, and Fondi di Ricerca di Ateneo (FRA), linea C, progetto TORNADO. A.B.K. and D.I. acknowledge funding provided by the University of Belgrade - Faculty of Mathematics (the contract 451-03-136/2025-03/200104) through the grants by the Ministry of Science, Technological Development and Innovation of the Republic of Serbia. C.M. acknowledges support from Fondecyt Iniciación grant 11240336. C.G.B. acknowledges support from the Consejo Nacional de Investigaciones Científicas y Técnicas (CONICET) and the Secretaría de Ciencia y Tecnología de la Universidad Nacional de Córdoba (SeCyT). C.M. acknowledges support from FONDECYT Iniciación grant 11240336. D.M. acknowledges financial support from CAPES – Finance Code 001. D.D. and M.F. acknowledge the financial contribution from PRIN-MIUR 2022 and from the Timedomes grant within the “INAF 2023 Finanziamento della Ricerca Fondamentale”. W.N.B. acknowledges the support of USA NSF grant AST-2407089. S.E.I.B. is supported by the Deutsche Forschungsgemeinschaft (DFG) under Emmy Noether grant number BO 5771/1-1. S.P. is supported by the international Gemini Observatory, a program of NSF NOIRLab, which is managed by the Association of Universities for Research in Astronomy (AURA) under a cooperative agreement with the U.S. National Science Foundation, on behalf of the Gemini partnership of Argentina, Brazil, Canada, Chile, the Republic of Korea, and the United States of America. C.R. acknowledges support from SNSF Consolidator grant F01-13252 and the China-Chile joint research fund. M.J.T. acknowledges funding from UKRI grant ST/X001075/1.

Funding for the SDSS and SDSS-II has been provided by the Alfred P. Sloan Foundation, the Participating Institutions, the National Science Foundation, the U.S. Department of Energy, the National Aeronautics and Space Administration, the Japanese Monbukagakusho, the Max Planck Society, and the Higher Education Funding Council for England. The SDSS Web site is <http://www.sdss.org/>. The SDSS is managed by the Astrophysical Research Consortium for the Participating Institutions. The Participating Institutions are the American Museum of Natural History, Astrophysical Institute Potsdam, University of Basel, University of Cambridge, Case Western Reserve University, University of Chicago, Drexel University, Fermilab, the Institute for Advanced Study, the Japan Participation Group, Johns Hopkins University, the Joint Institute for Nuclear Astrophysics, the Kavli Institute for Particle Astrophysics and Cosmology, the Korean Scientist Group, the Chinese Academy of Sciences (LAMOST), Los Alamos National Laboratory, the Max-Planck-Institute for Astronomy (MPIA), the Max-Planck-Institute for Astrophysics (MPA), New Mexico State University, Ohio State University, University of Pittsburgh, University of Portsmouth, Princeton University, the United States Naval Observatory, and the University of Washington.

Software packages: NumPy (Harris et al. 2020), Astropy (Astropy Collaboration et al. 2018), Matplotlib (Hunter 2007), Celerite (Foreman-Mackey et al. 2017), Pandas (Pandas Development Team 2020).

References

- Abazajian, K. N., Adelman-McCarthy, J. K., Agüeros, M. A., et al. 2009, *ApJS*, 182, 543
- Ananna, T. T., Salvato, M., LaMassa, S., et al. 2017, *ApJ*, 850, 66
- Annis, J., Soares-Santos, M., Strauss, M. A., et al. 2014, *ApJ*, 794, 120
- Arnouts, S., Cristiani, S., Moscardini, L., et al. 1999, *MNRAS*, 310, 540
- Arévalo, P., Churazov, E., Lira, P., et al. 2024, *A&A*, 684, A133
- Assef, R. J., Kochanek, C. S., Brodwin, M., et al. 2010, *ApJ*, 713, 970
- Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, 156, 123
- Bauer, A., Baltay, C., Coppi, P., et al. 2009, *ApJ*, 696, 1241–1256
- Brammer, G. B., van Dokkum, P. G., & Coppi, P. 2008, *ApJ*, 686, 1503
- Brandt, W. N., Ni, Q., Yang, G., et al. 2018, arXiv e-prints, arXiv:1811.06542
- Brescia, M., Salvato, M., Cavuoti, S., et al. 2019, *MNRAS*, 489, 663
- Burke, C. J., Liu, X., Shen, Y., Yang, Q., & Gammie, C. 2021, *Science*, 373, 789
- Cackett, E. M., Bentz, M. C., & Kara, E. 2021, *iScience*, 24, 102557
- Cardamone, C. N., van Dokkum, P. G., Urry, C. M., et al. 2010, *ApJS*, 189, 270
- Cirasuolo, M., Afonso, J., Bender, R., et al. 2012, in *Proc. SPIE*, Vol. 8446, 84460S
- Collier, S. & Peterson, B. M. 2001, *ApJ*, 555, 775

- Dahlen, T., Mobasher, B., Faber, S. M., et al. 2013, *ApJ*, 775, 93
- de Jong, R. S., Agertz, O., Berbel, A. A., et al. 2019, *The Messenger*, 175, 3
- DESI Collaboration, Aghamousa, A., Aguilar, J., et al. 2016, arXiv e-prints, arXiv:1611.00036
- Duncan, K. J., Brown, M. J. I., Williams, W. L., et al. 2018, *MNRAS*, 473, 2655
- Euclid Collaboration, Scaramella, R., Amiaux, J., et al. 2022, *A&A*, 662, A112
- Fabian, A. 2012, *ARA&A*, 50, 455
- Fan, X., Bañados, E., & Simcoe, R. A. 2023, *ARA&A*, 61, 373
- Foreman-Mackey, D., Hogg, D. W., & Morton, T. D. 2017, *AJ*, 154, 220
- Fontopoulou, S., Salvato, M., Hasinger, G., et al. 2012, *ApJS*, 198, 1
- Frank, J., King, A., & Raine, D. 2002, *Accretion Power in Astrophysics*, 3rd edn. (Cambridge University Press)
- Frieman, J. A., Bassett, B., Becker, A., et al. 2008, *AJ*, 135, 338
- Fukugita, M., Ichikawa, T., Gunn, J. E., et al. 1996, *AJ*, 111, 1748
- Gunn, J. E., Carr, M., Rockosi, C., Sekiguchi, M., & et al. 1998, *AJ*, 116, 3040
- Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Nature*, 585, 357
- Hsu, L.-T., Salvato, M., Nandra, K., et al. 2014, *ApJ*, 796, 60
- Hunter, J. D. 2007, *Computing in Science and Engineering*, 9, 90
- Ilbert, O., Arnouts, S., McCracken, H. J., et al. 2006, *A&A*, 457, 841
- Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2019, *ApJ*, 873, 111
- Ivezić, Ž. & MacLeod, C. 2014, in *IAU Symposium*, Vol. 304, 395
- Ivezić, Z., Sesar, B., Juric, M., & et al. 2007, *AJ*, 134, 973
- Kelly, L. O., Bechtold, J., & Siemiginowska, A. 2009, *ApJ*, 698, 895
- Kollmeier, J. A., Rix, H.-W., Aerts, C., et al. 2025, *AJ*, 171, 52
- Kormendy, J. & Ho, L. C. 2013, *ARA&A*, 51, 511–653
- Kubota, A. & Done, C. 2018, *MNRAS*, 480, 1247
- Laur, J., Tempel, E., Tamm, A., et al. 2022, *A&A*, 668, A8
- Li, G., Assef, R. J., Brandt, W. N., et al. 2025, arXiv e-prints, arXiv:2512.08654
- Li, Z., McGreer, I. D., Wu, X.-B., Fan, X., & Yang, Q. 2018, *ApJ*, 861, 6
- Lima, E. V. R., Sodré, L., Bom, C. R., et al. 2022, *Astronomy and Computing*, 38, 100510
- LSST Science Collaboration. 2009, *LSST Science Book*, Version 2.0
- Luo, Y., Shen, Y., & Yang, Q. 2020, *MNRAS*, 494, 3686
- MacLeod, C. L., Ivezić, Ž., Kochanek, C. S., et al. 2010, *ApJ*, 721, 1014
- Madau, P. & Haardt, F. 2015, *ApJ*, 813, L8
- Marshall, P., Clarkson, W., Shemmer, O., et al. 2017, *The LSST Strong Lens Time Delay Challenge*, Zenodo
- Merloni, A., Predehl, P., Becker, W., et al. 2012, *eROSITA Science Book: Mapping the Structure of the Energetic Universe*
- Newman, J. A. & Gruen, D. 2022, *ARA&A*, 60, 363
- Padovani, P., Alexander, D. M., Assef, R. J., et al. 2017, *A&A Rev.*, 25, 2
- Pandas Development Team. 2020, *Astrophysics Source Code Library [asc1:2002.018]*
- Peca, A., Vignali, C., Gilli, R., et al. 2021, *ApJ*, 906, 90
- Predehl, P., Andritschke, R., Arefiev, V., et al. 2021, *A&A*, 647, A1
- Richards, G. T., Fan, X., Newberg, H. J., et al. 2002, *AJ*, 123, 2945
- Richards, G. T., Lacy, M., Storrie-Lombardi, L. J., et al. 2006, *ApJS*, 166, 470
- Roster, W., Salvato, M., Krippendorf, S., et al. 2024, *A&A*, 692, A260
- Salvato, M., Hasinger, G., Ilbert, O., et al. 2009, *ApJ*, 690, 1250
- Salvato, M., Ilbert, O., Hasinger, G., et al. 2011, *ApJ*, 742, 61
- Salvato, M., Ilbert, O., & Hoyle, B. 2019, *Nature Astronomy*, 3, 212
- Salvato, M., Wolf, J., Dwelly, T., et al. 2022, *A&A*, 661, A3
- Sánchez-Sáez, P., Lira, P., Mejía-Restrepo, J., et al. 2018, *ApJ*, 864, 87
- Sartori, L. F., Schawinski, K., Trakhtenbrot, B., et al. 2018, *MNRAS*, 476, L34
- Savić, Đ. V., Jankov, I., Yu, W., et al. 2023, *ApJ*, 953, 138
- Saxena, A., Salvato, M., Roster, W., et al. 2024, *A&A*, 690, A365
- Schlegel, D. J., Finkbeiner, D. P., & Davis, M. 1998, *ApJ*, 500, 525
- Schneider, D. P., Hall, P. B., Richards, G. T., et al. 2007, *AJ*, 134, 102
- Schneider, D. P., Richards, G. T., Hall, P. B., et al. 2010, *AJ*, 139, 2360–2373
- Sesar, B., Ivezić, Ž., Lupton, R. H., et al. 2007, *AJ*, 134, 2236
- Shen, Y., Grier, C. J., Horne, K., et al. 2024, *ApJS*, 272, 26
- Shen, Y., Richards, G. T., Strauss, M. A., et al. 2011, *ApJS*, 194, 45
- Shirley, R., Salvato, M., Cohen-Tanugi, J., et al. 2025, submitted to *A&A*
- Silk, J. & Rees, M. J. 1998, *A&A*, 331, L1
- Simm, T., Salvato, M., Saglia, R., et al. 2016, *A&A*, 585, A129
- Stone, Z., Shen, Y., Burke, C. J., et al. 2022, *MNRAS*, 514, 164
- Suberlak, K. L., Ivezić, Ž., & MacLeod, C. 2021, *ApJ*, 907, 96
- Tachibana, Y., Graham, M. J., Kawai, N., et al. 2020, *ApJ*, 903, 54
- Tamura, N., Takato, N., Shimono, A., et al. 2016, *Ground-based and airborne instrumentation for astronomy VI*, 9908, 456
- Tang, J.-J., Wolf, C., & Tonry, J. 2023, *Nature Astronomy*, 7, 473–480
- Temple, M. J., Hewett, P. C., & Banerji, M. 2021, *MNRAS*, 508, 737
- Vanden Berk, D. E., Wilhite, B. C., Kron, R. G., et al. 2004, *ApJ*, 601, 692
- Wright, E. L., Eisenhardt, P. R. M., Mainzer, A. K., & et al. 2010, *AJ*, 140, 1868
- York, D. G., Adelman, J., Anderson, J. E., et al. 2000, *AJ*, 120, 1579
- Yu, W., Richards, G. T., Ruan, J. J., et al. 2025, *ApJ*, 992, 130
- Zu, Y., Kochanek, C. S., Kozłowski, S., & Udalski, A. 2013, *ApJ*, 765, 106
- Zu, Y., Kochanek, C. S., & Peterson, B. M. 2011, *ApJ*, 735, 80

Affiliations

- ¹Instituto de Astrofísica, Facultad de Ciencias Exactas, Universidad Andres Bello, Fernández Concha 700, 7591538 Las Condes, Santiago, Chile
- ²Instituto de Estudios Astrofísicos, Facultad de Ingeniería y Ciencias, Universidad Diego Portales, Av. Ejército Libertador 441, Santiago, Chile
- ³Millennium Institute of Astrophysics, Nuncio Monseñor Sótero Sanz 100, Providencia, Santiago, Chile
- ⁴European Southern Observatory, Karl-Schwarzschild-Strasse 2, 85748 Garching bei München, Germany
- ⁵Max-Planck-Institut für extraterrestrische Physik, Giessenbachstr. 1, 85748 Garching, Germany
- ⁶Department of Physics and Astronomy, Wayne State University, 666 W. Hancock St, Detroit, MI, 48201, USA
- ⁷Instituto de Alta Investigación, Universidad de Tarapacá, Casilla 7D, Arica, Chile
- ⁸School of Physics and Astronomy, Monash University, Clayton, Victoria 3800, Australia
- ⁹ARC Centre of Excellence for Gravitational Wave Discovery – OzGrav, Australia
- ¹⁰Instituto de Astronomía Teórica y Experimental, (IATE, CONICET-UNC), Córdoba, Argentina
- ¹¹Universidad Nacional de Córdoba, Observatorio Astronómico de Córdoba
- ¹²Institute for Theoretical Physics, Heidelberg University, Philosophenweg 12, D–69120, Heidelberg, Germany
- ¹³Max-Planck-Institut für Astronomie, Königstuhl 17, 69117 Heidelberg, Germany
- ¹⁴Department of Astronomy & Astrophysics, The Pennsylvania State University, University Park, PA 16802, USA
- ¹⁵The Institute for Gravitation and the Cosmos, The Pennsylvania State University, University Park, PA 16802, USA
- ¹⁶Department of Physics, 104 Davey Laboratory, The Pennsylvania State University, University Park, PA 16802, USA
- ¹⁷Department of Physics, University of Napoli "Federico II", via Cinthia 9, 80126 Napoli, Italy
- ¹⁸INAF - Osservatorio Astronomico di Capodimonte, via Moiaranello 16, 80131 Napoli, Italy
- ¹⁹Center for Theoretical Physics, Polish Academy of Sciences, Al. Lotników 32/46, 02-668 Warsaw, Poland
- ²⁰Dipartimento di Fisica “Ettore Pancini”, Università di Napoli Federico II, Via Cintia 80126, Naples, Italy
- ²¹Ruđer Bošković Institute, Bijenička Cesta 54, 10000 Zagreb, Croatia
- ²²Frontier Research Institute for Interdisciplinary Sciences, Tohoku University, Sendai 980-8578, Japan
- ²³Global Center for Science and Engineering, Faculty of Science and Engineering, Waseda University, 3-4-1, Okubo, Shinjuku, Tokyo 169-8555, Japan
- ²⁴Department of Astronomy, Faculty of Mathematics, University of Belgrade, Studentski trg 16, 11000 Belgrade, Serbia
- ²⁵Hamburger Sternwarte, Universität Hamburg, Gojenbergsweg 112, D-21029 Hamburg, Germany
- ²⁶Kavli Institute for Astronomy and Astrophysics, Peking University, Beijing 100871, People’s Republic of China
- ²⁷Departamento de Física, Universidad Técnica Federico Santa María, Vicuña Mackenna 3939, San Joaquín, Santiago, Chile
- ²⁸Universidade Federal de Santa Maria (UFSM), Centro de Ciências Naturais e Exatas (CCNE), Santa Maria, 97105-900, RS, Brazil
- ²⁹International Gemini Observatory/NSF NOIRLab, Casilla 603, La Serena, Chile

³⁰Eureka Scientific, 2452 Delmer Street, Suite 100, Oakland, CA 94602-3017, USA

³¹Department of Physics, Yale University, P.O. Box 208120, New Haven, CT 06520, USA

³²NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA

³³Center for Space Science and Technology, University of Maryland Baltimore County, MD 21250, USA

³⁴Department of Astronomy, University of Geneva, ch. d'Ecogia 16, 1290, Versoix, Switzerland

³⁵Department of Physics, Drexel University, 32 S. 32nd Street, Philadelphia, PA 19104, USA

³⁶Exzellenzcluster ORIGINS, Boltzmannstr. 2, D-85748 Garching, Germany

³⁷Centre for Extragalactic Astronomy, Department of Physics, Durham University, South Road, Durham DH1 3LE, United Kingdom

³⁸Physics Department, Tor Vergata University of Rome, Via della Ricerca Scientifica 1, 00133 Rome, Italy

³⁹INAF – Astronomical Observatory of Rome, Via Frascati 33, 00040 Monte Porzio Catone, Italy

⁴⁰INFN – Rome Tor Vergata, Via della Ricerca Scientifica 1, 00133 Rome, Italy

⁴¹Department of Physics & Astronomy, Bishop's University, 2600 rue College, Sherbrooke, QC, J1M 1Z7, Canada

⁴²National Radio Astronomy Observatory, 520 Edgemont Road, Charlottesville, VA 22904, USA

⁴³Department of Astronomy, University of Virginia, 530 McCormick Rd, Charlottesville, VA 22904, USA

⁴⁴Department of Astronomy, University of Michigan, 1085 S University, Ann Arbor, MI 48109, USA

Appendix A: VAR-PZ workflow

Figure A.1 presents the computational pipeline and algorithmic steps used in the VAR-PZ implementation. This flowchart serves as a visual reference for the complete VAR-PZ computational workflow.

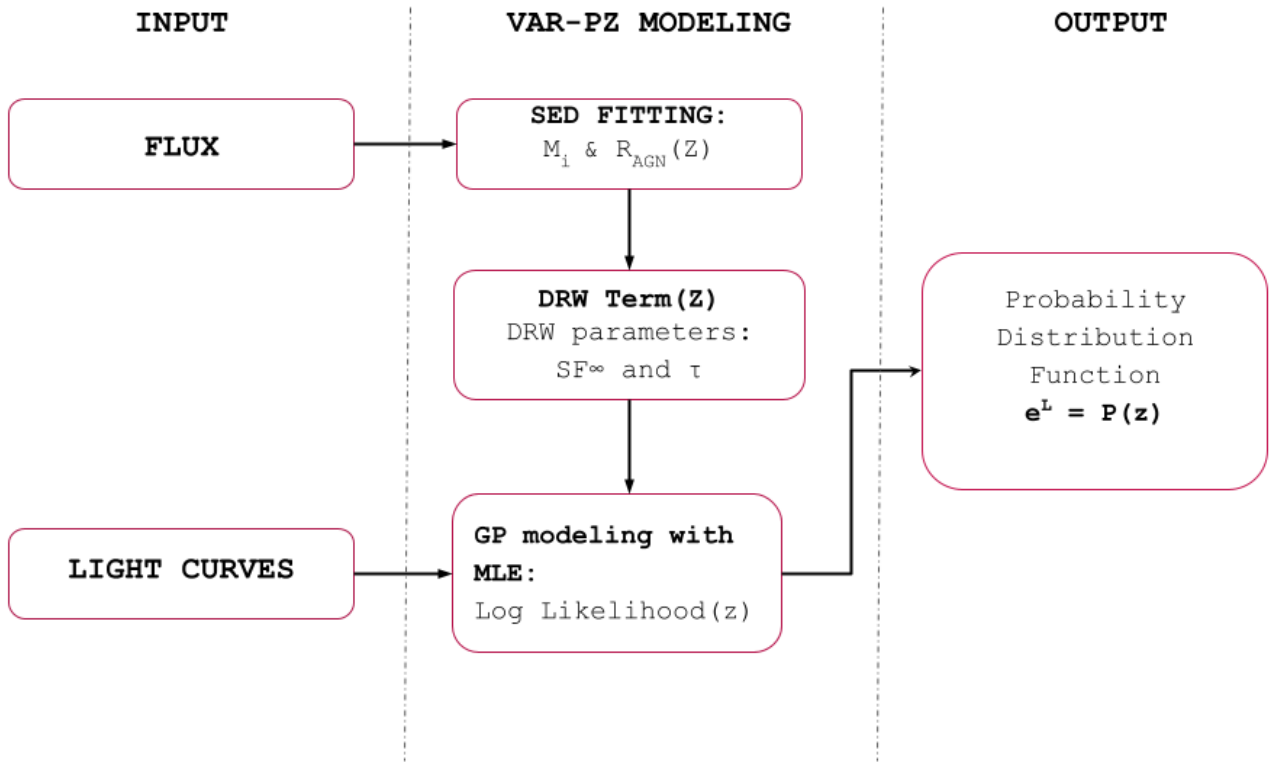


Fig. A.1. Flowchart illustrating the overall workflow of VAR-PZ: Input fluxes and light curves are used to model the SED using LRT templates, and the light curves are modeled as a DRW using the M10 kernel. The resulting kernel is incorporated into a Gaussian Process to produce redshift-dependent photometric redshift PDFs.

Appendix B: LSST DDF cadence

Figure B.1 a simulated light curve based on the XMM-LSS DDF observing cadence to illustrate the improved temporal sampling compared to the WFD. This simulation highlights the dense sampling of the DDFs, which constitute approximately 6.5% of the total survey time and are designed to provide intensive monitoring, with each field receiving over 20,000 visits over the survey duration.

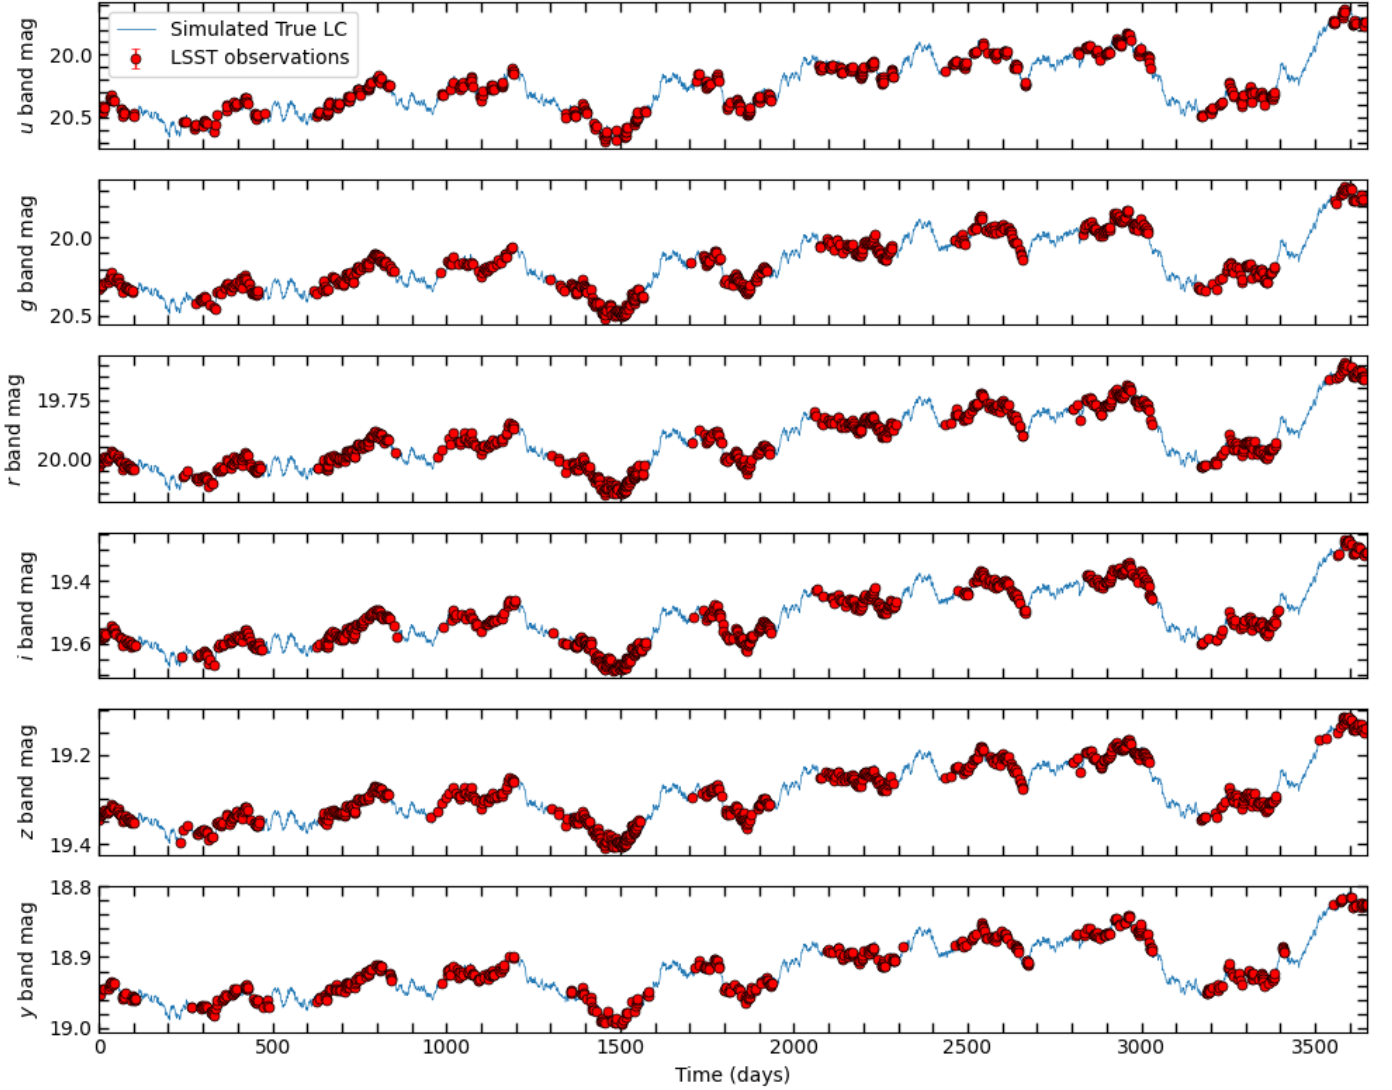


Fig. B.1. Simulated light curve for an AGN located in a Rubin LSST DDF. This example corresponds to the XMM-LSS field, centered at $\alpha_{J2000} = 35.57^\circ$, $\delta_{J2000} = -4.82^\circ$