LATENT SDES FOR MODELLING QUASAR VARIABILITY AND INFERRING BLACK HOLE PROPERTIES

Anonymous authors
Paper under double-blind review

ABSTRACT

Active galactic nuclei (AGN) are thought to be powered by the accretion of matter around supermassive black holes at the centers of galaxies. The variability of an AGN's brightness over time can provide valuable insights into the physical characteristics of its underlying black hole. The variability pattern is thought to closely follow a damped random walk, described by a stochastic differential equation (SDE). Upcoming wide-field telescopes are expected to observe 100 million AGN in multiple bandpass filters. Latent SDEs are well suited for modeling the AGN time series, as they explicitly model the underlying dynamics. We modify latent SDEs to jointly reconstruct the unobserved portions of multivariate AGN light curves as well as infer their physical properties, such as the black hole mass. We train our model on a realistic physics-based simulation of ten-year AGN light curves, and we demonstrate its ability to fit AGN light curves even in the presence of long seasonal gaps and irregular sampling across different bands. Our method is general and applicable to a wide range of time series, and has the potential to provide a deeper understanding of the physical properties of black holes in AGN.

1 INTRODUCTION

Active galactic nuclei (AGN) are among the brightest objects in the universe and play a crucial role in galaxy evolution. They are thought to be powered by the conversion of gravitational potential energy into thermal radiation by feeding on matter in the hot accretion disks of super massive black holes (SMBH) at the center of galaxies (Salpeter, 1964; Zel'dovich, 1964). Luminous AGN with unobscured accretion disks are known as quasars. These objects are so luminous they remain observable at extreme cosmological distances (Mortlock et al., 2011; Bañados et al., 2018) making them exceptional probes of the early universe.

The stochastic variability of quasars has been studied extensively since their discovery (Greenstein, 1963; Hazard et al., 1963; Matthews & Sandage, 1963; Oke, 1963; Schmidt, 1963). While the accretion disk of an AGN is too small to resolve at extragalactic distances, its light curve variability can provide insight into the intrinsic physical properties of the AGN. For instance, the amplitude increases with decreasing luminosity, rest-frame wavelength, and Eddington ratio (Wills et al., 1993; Giveon et al., 1999; Berk et al., 2004). Efforts have been made to constrain the correlation of variability parameters with black hole mass, but their robustness has been inconclusive, with studies claiming positive or negative relations depending on the degrees of observational bias present in the fit data (Wold et al., 2008; MacLeod et al., 2010; Simm et al., 2015). Inferring black hole physics from AGN light curves can teach us about the evolution of the universe and the nature of dark matter and dark energy (Khadka & Ratra, 2020).

The Rubin Observatory Legacy Survey of Space and Time (LSST) marks an unprecedented improvement in data quality and volume. The LSST main survey alone (10,000 $\deg^2$) is projected to yield 100 million AGN light curves with six optical bandpass filters (*ugrizY*) at $55 - 185$ samplings per band, out to redshift $z \sim 7.5$ and moderate luminosities of $L \sim 10^{44}$ erg/s over a ten year period (Abell et al., 2009). Machine learning algorithms are well suited to analyze the large amount of data expected from LSST; however, these light curves pose challenges for existing techniques such as multiple bands, long gaps of missing data, non-uniform sampling, and photometric and systematic noise.

Quasar variability is thought to closely follow a damped random walk (DRW), a type of Gaussian process also known as the Ornstein-Uhlenbeck process (Rasmussen & Williams, 2006; Zu et al., 2013). The DRW $X(t)$ is governed by the stochastic differential equation (SDE):

$$dX(t) = -\frac{1}{\tau}X(t)dt + \sigma\sqrt{dt}\epsilon(t) + bdt \tag{1}$$

where $\epsilon(t)$ is a white noise process with mean zero and variance one (Kelly et al., 2009). $X(t)$ has standard deviation $\mathrm{SF}_\infty = \sigma\sqrt{\tau/2}$ with characteristic time scale τ and mean $\overline{X} = b\tau$. Gaussian process regression (GPR) can be used with an absolute exponential kernel to measure the variability parameters SF_∞ and τ, which may be correlated to the black hole mass (Suberlak et al., 2021). A smaller τ yields more high frequency variations since the variability happens on a shorter time scale. Fig. 3 shows an example DRW for two different values of τ.

Here we model the quasar variability as a latent SDE (Li et al., 2020). Latent SDEs work similarly to variational auto-encoders (VAEs) (Kingma & Welling, 2013) with a neural SDE as the decoder. We use the SDE to model the light curve dynamics including at unobserved times. We simultaneously train our network to predict the black hole mass and variability parameters SF_∞ and τ from the latent space. Latent SDEs can be superior to GPR since they learn to model the variability across an entire training set instead of independently for each light curve, and we do not need to specify an analytic kernel. Previous works have modelled quasar variability using RNNs (Tachibana et al., 2020; Sá nchez-Sáez et al., 2021) and attentive neural processes (Park et al., 2021), but we are the first to model the variability as an SDE.

2 METHOD

2.1 TRAINING SET

Our training set is comprised of a realistic, physics based simulation of LSST 10-year light curves. We first generate a realization of the DRW for a given SF_∞, τ and mean flux. To simulate the reaction of the accretion disk to a driving DRW, a transfer function is calculated for each spectral band (Blandford & McKee, 1982). All of the physical properties of the black hole and quasar geometry are included in the transfer functions that are modelled using a general relativistic ray-traced accretion disk simulation (Best et al., 2022; Bursa, 2017; 2018). These transfer functions assume an X-ray source irradiating a vertically thin, optically thick accretion disk (Shakura & Sunyaev, 1973) (the irradiated disk model (Sergeev et al., 2005; Cackett et al., 2007)). The accretion disk is assumed to be a modified thin-disk where the temperature profile is extended to have a wind cooling parameter as described in Sun et al. (2019). The DRW is convolved with the transfer function kernel to obtain simulated light curves for a set of black hole parameters [1].

We then degrade the time series to LSST like observations and noise. The observation times and photometric noise are taken from `rubin_sim`[2] by sampling 10,000 mock observation schemes from across the sky with between 750 and 1000 total observations across the ten years. The systematic uncertainty is set to 0.005 mag, the maximum value expected for LSST (Ivezić & et al., 2019). The time series is discretized to three day intervals. We use a data set of 10,000 mock light curves per epoch with 9,000 used for training and 1,000 strictly for validation. The training set is randomly regenerated on the fly each epoch during training to avoid overfitting and improve performance.

2.2 MODEL ARCHITECTURE

The input into the network is the magnitude brightness and uncertainty for a total of 12 features at each time step (6 bands and 6 uncertainties). The network architecture can be broken down into three different parts: the RNN encoder, the neural SDE decoder, and the multi layer perceptron (MLP) parameter estimator. The light curves are normalized in the network such that each band has observed data with mean zero and variance one. We use a hidden size of 128, context size of 32, and latent size of 16. Each layer is followed by an activation function, layer norm (Ba et al., 2016), and

[1] The light curves are converted from flux to AB magnitude: $m_{\mathrm{AB}} = -2.5\log_{10}(f_\nu) - 48.60$ in cgs units

[2] https://github.com/lsst/rubin_sim

1% dropout (Srivastava et al., 2014). The network has 542,885 trainable parameters. The network uses LeakyReLU activation except the neural SDE which uses Softplus.

The encoder starts with a GRU-D (Che et al., 2016) layer, a modified version of the Gated Recurrent Unit (GRU) (Chung et al., 2014) that includes both masking and irregular time intervals. This is followed by a GRU layer and two dense layers that output the context vector. Two more dense layers are used to produce the mean μ_{q_z} and variance $\sigma^2_{q_z}$ of the variational distribution. We then draw the latent vector $z_0 \sim \mathcal{N}(\mu_{q_z}, \sigma^2_{q_z})$, the initial conditions of our neural SDE.

The decoder contains an itô neural SDE with diagonal noise that has drift, diffusion, and prior drift networks (Li et al., 2020). The drift and prior drift networks are both MLPs with a residual skip connection (He et al., 2015). There are three dense layers pre skip and two post skip. The diffusion network consists of two dense layers and is applied element wise to satisfy the diagonal noise. In order to model both the dynamics of the time series as well as produce uncertainties on predictions, we use an RNN to project the output of the SDE and the latent vector to the reconstructed light curve. The projector consists of one dense layer and two GRU layers with a one dense layer skip connection.

The parameter estimation is done using the same MLP architecture as the drift networks. The input into the parameter estimator is μ_{q_z}, $\sigma^2_{q_z}$, the output of the drift, diffusion, and prior drift networks at μ_{q_z}, and the mean and standard deviation of each band of the unnormalized light curve.

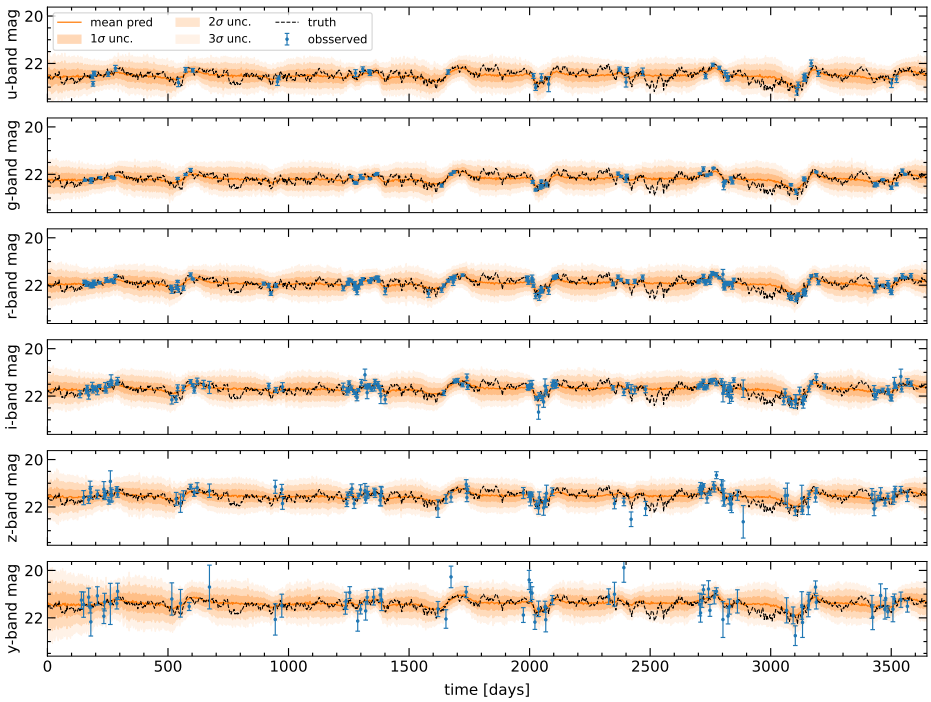

Figure 1: Reconstructed light curve using a latent SDE. The uncertainty is calculated by sampling 100 predictions from the posterior. The true light curve is shown in black and the mock observations (context points for the network) are shown in blue.

2.3 Uncertainty quantification and loss

The epistemic uncertainty is included using Monte Carlo dropout (Gal & Ghahramani, 2015). For the parameter estimation, the aleatoric uncertainty is found by minimizing the multivariate, negative Gaussian log-likelihood. For the light curve reconstruction, the evidence lower bound (ELBO) is minimized with negative Gaussian log-likelihood at each time step between the predicted and true light curve. The total loss is the sum of the parameter estimation and light curve reconstruction losses. Linear KL-annealing is used for the first 50 epochs as well as annealing for the parameter

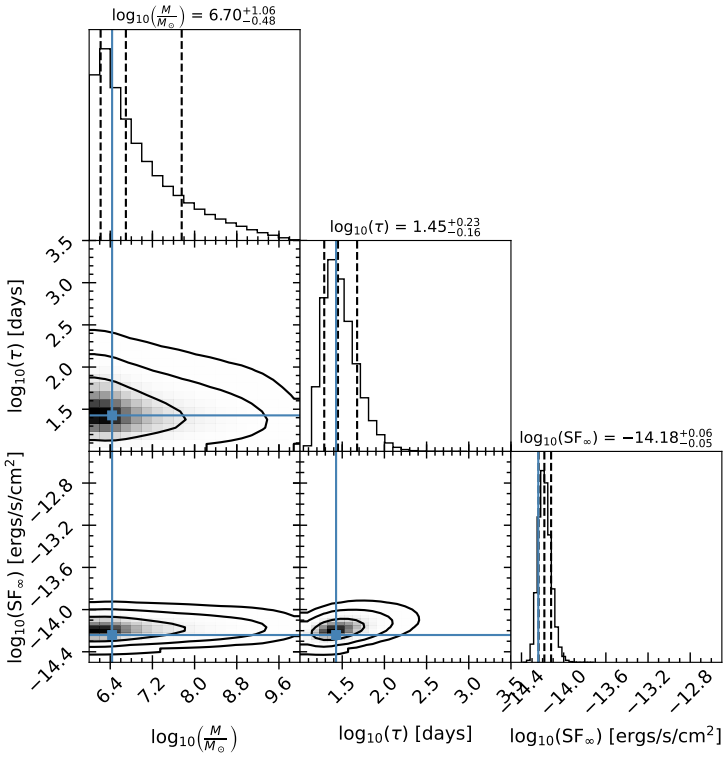

Figure 2: Example posterior for parameter predictions. The diagonal plots display marginal probability distributions with 68% probability and medium indicated by dotted lines. The central plots depict joint probabilities with 68%, 95%, and 99% enclosed by contours. The true parameter values are represented by the blue lines.

predictions. Our network is trained for 80 epochs using an Adam optimizer (Kingma & Ba, 2017) with an initial learning rate of 0.0025, an exponential decay of 0.95, and a batch size of 128. The parameter predictions are done in the logit space to constrain the predictions within their uniformly drawn range.

3 Results and discussion

Fig. 1 shows an example light curve reconstruction from the validation set. We can see the predicted uncertainty is much higher away from the context points, and the network is able to infer information between bands. Fig 2 shows an example posterior for the parameter predictions on the same light curve. A population of these measurements could be used to find the correlation between the quasar variability and black hole mass. Fig. 4 shows confusion matrices for our parameter predictions.

Here we presented latent SDEs as a way of simultaneously modelling quasar variability and predicting black hole properties. A pretrained Latent SDE is faster than training GPR on each light curve individually, which is important when applied to the 100 million light curves expected from LSST. Our network is trained to predict the variability parameters of the quasar source flux while GPR can only predict the effective parameters after the transfer function convolution, biasing the results. We are also able to directly predict black hole mass, and we hope to extend our parameter predictions to more black hole properties such as the spin and inclination angle. Latent SDEs also have the potential for anomaly detection to find quasars with variability that differs from the training set. The method presented here is completely general and can be applied to any multivariate time series with missing data and irregular sampling with applications to other astronomical time series and beyond.

REFERENCES

Paul A Abell, David L Burke, Mario Hamuy, Martin Nordby, Tim S Axelrod, David Monet, Bojan Vrsnak, Paul Thorman, DR Ballantyne, Joshua D Simon, et al. Lsst science book, version 2.0. Technical report, LSST, 2009.

Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. Layer normalization, 2016. URL https://arxiv.org/abs/1607.06450.

Eduardo Bañados, Bram P. Venemans, Chiara Mazzucchelli, Emanuele P. Farina, Fabian Walter, Feige Wang, Roberto Decarli, Daniel Stern, Xiaohui Fan, Frederick B. Davies, Joseph F. Hennawi, Robert A. Simcoe, Monica L. Turner, Hans-Walter Rix, Jinyi Yang, Daniel D. Kelson, Gwen C. Rudie, and Jan Martin Winters. An 800-million-solar-mass black hole in a significantly neutral Universe at a redshift of 7.5. , 553(7689):473–476, January 2018. doi: 10.1038/nature25180.

Daniel E Vanden Berk, Brian C Wilhite, Richard G Kron, Scott F Anderson, Robert J Brunner, Patrick B Hall, Željko Ivezić, Gordon T Richards, Donald P Schneider, Donald G York, et al. The ensemble photometric variability of˜ 25,000 quasars in the sloan digital sky survey. *The Astrophysical Journal*, 601(2):692, 2004.

Henry Best, Joshua Fagin, Georgios Vernardos, and Matthew O'Dowd. Resolving the vicinity of supermassive black holes with gravitational microlensing, 2022. URL https://arxiv.org/abs/2210.10500.

R. D. Blandford and C. F. McKee. Reverberation mapping of the emission line regions of Seyfert galaxies and quasars. , 255:419–439, April 1982. doi: 10.1086/159843.

Michal Bursa. Numerical implementation of equations for photon motion in Kerr spacetime. In *RAGtime 17-19: Workshops on Black Holes and Neutron Stars*, pp. 7–21, December 2017.

Michal Bursa. SIM5: Library for ray-tracing and radiation transport in general relativity. Astrophysics Source Code Library, record ascl:1811.011, November 2018.

Edward M. Cackett, Keith Horne, and Hartmut Winkler. Testing thermal reprocessing in active galactic nuclei accretion discs. , 380(2):669–682, September 2007. doi: 10.1111/j.13652966.2007.12098.x.

Zhengping Che, Sanjay Purushotham, Kyunghyun Cho, David Sontag, and Yan Liu. Recurrent neural networks for multivariate time series with missing values, 2016. URL https://arxiv.org/abs/1606.01865.

Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling, 2014. URL https://arxiv.org/abs/1412.3555.

Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning, 2015. URL https://arxiv.org/abs/1506.02142.

Uriel Giveon, Dan Maoz, Shai Kaspi, Hagai Netzer, and Paul S Smith. Long-term optical variability properties of the palomar—green quasars. *Monthly Notices of the Royal Astronomical Society*, 306(3):637–654, 1999.

Jesse L. Greenstein. Red-Shift of the Unusual Radio Source: 3C 48. , 197(4872):1041–1042, March 1963. doi: 10.1038/1971041a0.

C. Hazard, M. B. Mackey, and A. J. Shimmins. Investigation of the Radio Source 3C 273 By The Method of Lunar Occultations. , 197(4872):1037–1039, March 1963. doi: 10.1038/1971037a0.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition, 2015. URL https://arxiv.org/abs/1512.03385.

Željko Ivezić and Kahn et al. LSST: From Science Drivers to Reference Design and Anticipated Data Products. , 873(2):111, March 2019. doi: 10.3847/1538-4357/ab042c.

Brandon C. Kelly, J. Bechtold, and A. Siemiginowska. Are the Variations in Quasar Optical Flux Driven by Thermal Fluctuations? In *American Astronomical Society Meeting Abstracts #213*, volume 213 of *American Astronomical Society Meeting Abstracts*, pp. 422.06, January 2009.

Narayan Khadka and Bharat Ratra. Using quasar X-ray and UV flux measurements to constrain cosmological model parameters. *Monthly Notices of the Royal Astronomical Society*, 497(1): 263–278, 06 2020. ISSN 0035-8711. doi: 10.1093/mnras/staa1855. URL https://doi.org/10.1093/mnras/staa1855.

Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017.

Diederik P Kingma and Max Welling. Auto-encoding variational bayes, 2013. URL https://arxiv.org/abs/1312.6114.

Xuechen Li, Ting-Kam Leonard Wong, Ricky T. Q. Chen, and David Duvenaud. Scalable gradients for stochastic differential equations. *International Conference on Artificial Intelligence and Statistics*, 2020.

Ch L MacLeod, Ž Ivezić, CS Kochanek, S Kozłowski, B Kelly, E Bullock, A Kimball, B Sesar, D Westman, K Brooks, et al. Modeling the time variability of sdss stripe 82 quasars as a damped random walk. *The Astrophysical Journal*, 721(2):1014, 2010.

Thomas A. Matthews and Allan R. Sandage. Optical Identification of 3C 48, 3C 196, and 3C 286 with Stellar Objects. , 138:30, July 1963. doi: 10.1086/147615.

Daniel J. Mortlock, Stephen J. Warren, Bram P. Venemans, Mitesh Patel, Paul C. Hewett, Richard G. McMahon, Chris Simpson, Tom Theuns, Eduardo A. Gonzáles-Solares, Andy Adamson, Simon Dye, Nigel C. Hambly, Paul Hirst, Mike J. Irwin, Ernst Kuiper, Andy Lawrence, and Huub J. A. Röttgering. A luminous quasar at a redshift of z = 7.085. , 474(7353):616–619, June 2011. doi: 10.1038/nature10159.

J. B. Oke. Absolute Energy Distribution in the Optical Spectrum of 3C 273. , 197(4872):1040–1041, March 1963. doi: 10.1038/1971040b0.

Ji Won Park, Ashley Villar, Yin Li, Yan-Fei Jiang, Shirley Ho, Joshua Yao-Yu Lin, Philip J. Marshall, and Aaron Roodman. Inferring black hole properties from astronomical multivariate time series with bayesian attentive neural processes, 2021. URL https://arxiv.org/abs/2106.01450.

Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian processes for machine learning*. Adaptive computation and machine learning. MIT Press, 2006. ISBN 026218253X.

E. E. Salpeter. Accretion of Interstellar Matter by Massive Objects. , 140:796–800, August 1964. doi: 10.1086/147973.

P. Sá nchez-Sáez, H. Lira, L. Martí, N. Sánchez-Pi, J. Arredondo, F. E. Bauer, A. Bayo, G. Cabrera-Vives, C. Donoso-Oliva, P. A. Estévez, S. Eyheramendy, F. Förster, L. Hernández-García, A. M. Muñoz Arancibia, M. Pérez-Carrasco, M. Sepúlveda, and J. R. Vergara. Searching for changing-state AGNs in massive data sets. i. applying deep learning and anomaly-detection techniques to find AGNs with anomalous variability behaviors. *The Astronomical Journal*, 162 (5):206, oct 2021. doi: 10.3847/1538-3881/ac1426. URL https://doi.org/10.3847%2F1538-3881%2Fac1426.

M. Schmidt. 3C 273 : A Star-Like Object with Large Red-Shift. , 197(4872):1040, March 1963. doi: 10.1038/1971040a0.

S. G. Sergeev, V. T. Doroshenko, Yu. V. Golubinskiy, N. I. Merkulova, and E. A. Sergeeva. Lag-Luminosity Relationship for Interband Lags between Variations in B, V, R, and I Bands in Active Galactic Nuclei. , 622(1):129–135, March 2005. doi: 10.1086/427820.

N. I. Shakura and R. A. Sunyaev. Reprint of 1973A&A....24..337S. Black holes in binary systems. Observational appearance. *Astronomy and Astrophysics*, 500:33–51, June 1973.

T Simm, R Saglia, M Salvato, R Bender, WS Burgett, KC Chambers, PW Draper, H Flewelling, N Kaiser, R-P Kudritzki, et al. Pan-starrs1 variability of xmm-cosmos agn-i. impact on photometric redshifts. *Astronomy & Astrophysics*, 584:A106, 2015.

Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(56):1929–1958, 2014. URL http://jmlr.org/papers/v15/srivastava14a.html.

Krzysztof L. Suberlak, Ž eljko Ivezić, and Chelsea MacLeod. Improving damped random walk parameters for SDSS stripe 82 quasars with pan-STARRS1. *The Astrophysical Journal*, 907 (2):96, feb 2021. doi: 10.3847/1538-4357/abc698. URL https://doi.org/10.3847%2F1538-4357%2Fabc698.

Mouyuan Sun, Yongquan Xue, Jonathan R. Trump, and Wei-Min Gu. Winds can 'blow up' AGN accretion disc sizes. , 482(2):2788–2794, January 2019. doi: 10.1093/mnras/sty2885.

Yutaro Tachibana, Matthew J. Graham, Nobuyuki Kawai, S. G. Djorgovski, Andrew J. Drake, Ashish A. Mahabal, and Daniel Stern. Deep modeling of quasar variability. *The Astrophysical Journal*, 903(1):54, nov 2020. doi: 10.3847/1538-4357/abb9a9. URL https://doi.org/10.3847%2F1538-4357%2Fabb9a9.

Beverley J Wills, MS Brotherton, D Fang, Charles C Steidel, and Wallace LW Sargent. Statistics of qso broad emission-line profiles. i. the c iv lambda 1549 line and the lambda 1400 feature. *The Astrophysical Journal*, 415:563, 1993.

M Wold, MS Brotherton, and Z Shang. Black hole mass and variability in quasars. In *AIP Conference Proceedings*, volume 1053, pp. 55–58. American Institute of Physics, 2008.

Ya. B. Zel'dovich. The Fate of a Star and the Evolution of Gravitational Energy Upon Accretion. *Soviet Physics Doklady*, 9:195, September 1964.

Ying Zu, C. S. Kochanek, Szymon Kozłowski, and Andrzej Udalski. IS QUASAR OPTICAL VARIABILITY a DAMPED RANDOM WALK? *The Astrophysical Journal*, 765(2):106, feb 2013. doi: 10.1088/0004-637x/765/2/106. URL https://doi.org/10.1088%2F0004-637x%2F765%2F2%2F106.

A APPENDIX

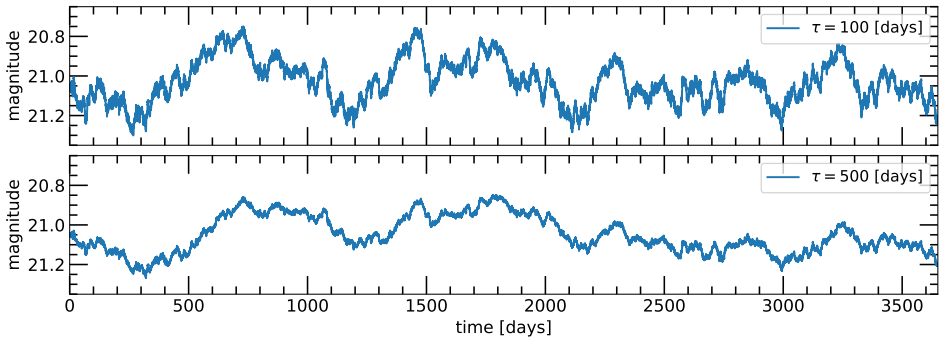

Figure 3: Damped random walk for two different values of τ. The damped random walk is generated as a flux and then converted to magnitude with mean 21 mag and SF_∞ of 0.1 mag. Both DRWs is generated using the same random state.

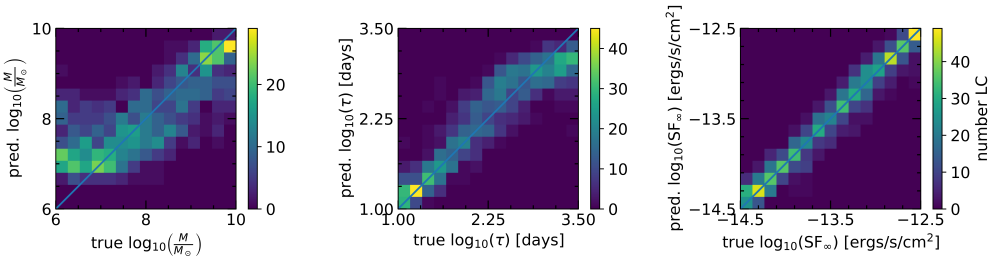

Figure 4: Confusion matrices for our mean parameter predictions vs the true parameters. The blue line across the diagonal represents perfect predictions.