



Real-time Likelihood-free Inference of Roman Binary Microlensing Events with Amortized Neural Posterior Estimation

Keming Zhang (张可名)¹, Joshua S. Bloom^{1,2}, B. Scott Gaudi³, François Lanusse⁴, Casey Lam¹, and Jessica R. Lu¹

¹Department of Astronomy, University of California, Berkeley, CA 94720-3411, USA; kemingz@berkeley.edu

²Lawrence Berkeley National Laboratory, 1 Cyclotron Road, MS 50B-4206, Berkeley, CA 94720-3411, USA

³Department of Astronomy, The Ohio State University, Columbus, OH 43210, USA

⁴AIM, CEA, CNRS, Université Paris-Saclay, Université Paris Diderot, Sorbonne Paris Cité, F-91191, Gif-sur-Yvette, France

Received 2021 February 10; revised 2021 March 22; accepted 2021 March 30; published 2021 May 11

Abstract

Fast and automated inference of binary-lens, single-source (2L1S) microlensing events with sampling-based Bayesian algorithms (e.g., Markov Chain Monte Carlo, MCMC) is challenged on two fronts: the high computational cost of likelihood evaluations with microlensing simulation codes, and a pathological parameter space where the negative-log-likelihood surface can contain a multitude of local minima that are narrow and deep. Analysis of 2L1S events usually involves grid searches over some parameters to locate approximate solutions as a prerequisite to posterior sampling, an expensive process that often requires human-in-the-loop domain expertise. As the next-generation, space-based microlensing survey with the Roman Space Telescope is expected to yield thousands of binary microlensing events, a new fast and automated method is desirable. Here, we present a likelihood-free inference approach named amortized neural posterior estimation, where a neural density estimator (NDE) learns a surrogate posterior $\hat{p}(\theta|x)$ as an observation-parameterized conditional probability distribution, from pre-computed simulations over the full prior space. Trained on 291,012 simulated Roman-like 2L1S simulations, the NDE produces accurate and precise posteriors within seconds for any observation within the prior support without requiring a domain expert in the loop, thus allowing for real-time and automated inference. We show that the NDE also captures expected posterior degeneracies. The NDE posterior could then be refined into the exact posterior with a downstream MCMC sampler with minimal burn-in steps.

Unified Astronomy Thesaurus concepts: Binary lens microlensing (2136); Gravitational microlensing exoplanet detection (2147)

1. Introduction

When the apparent trajectory of a foreground lens star passes close to a more distant source star, the gravitational field of the lens will perturb the light rays from the source, which results in a time-variable magnification. Such are single-lens, single-source (1L1S) microlensing events. Binary microlensing events occur when the lens is a system of two masses: either a binary star system or a star–planet configuration. Observation of such events provides a unique opportunity for exoplanet discovery as the planet-to-star mass ratio may be inferred from the light curve without having to detect light from the star–planet lens itself (see Gaudi 2010 for a review). A next-generation microlensing survey with the Roman Space Telescope (Spergel et al. 2015; hereafter Roman) is estimated to discover thousands of binary microlensing events over the 5 yr mission span, many with planetary-mass companions (Penny et al. 2019), which is roughly an order of magnitude more than events previously discovered (see Gaudi 2012 for a review).

While single-lens microlensing events are described by a simple analytic expression (“Paczynski light curve”), binary microlensing events require numerical forward models that are computationally expensive. In addition, binary microlensing light curves exhibit extraordinary phenomenological diversity, owing to the different geometrical configurations for which magnification could take place. This translates to a parameter space for which the likelihood surface suffers from a multitude of local minima that are disconnected, narrow, and deep; this issue significantly hampers any attempt of direct sampling-based inference such as MCMC where the chains are initialized

from a broad prior. As a result, binary microlensing events thus far have generally been analyzed on a case-by-case basis.

For some planetary-mass-ratio events, heuristics could be used to “read off” an approximate solution from the planetary anomaly in the light curve (Gould & Loeb 1992; Gaudi & Gould 1997). Khakpash et al. (2019) applied the heuristics described in Gaudi & Gould (1997) on simulated Roman light curves and found that the projected binary separation can be recovered very well for low-mass-ratio events, and the binary mass-ratios within an order of magnitude for events with wide and close caustic topologies.

More generally, an expensive grid search is usually conducted over a subset of parameters to which the magnification pattern is hyper-sensitive: i.e., binary separation, mass ratio, and the source trajectory angle of approach (e.g., Herrera-Martín et al. 2020). At each grid-point, the remaining parameters are searched for with simple Nelder–Mead optimization (Nelder & Mead 1965) or MCMC. The fixed-grid solutions are then used to seed full MCMC samplings to refine solutions and sample the posteriors. This status quo approach, which is both computationally expensive and requires domain expertise in the loop, thus presents a great challenge to analyze the thousands of binary microlensing events expected to be discovered by Roman.

Recent progress in deep learning provides a promising path for a solution. In particular, both Convolutional (CNN; LeCun et al. 2015) and Recurrent Neural Networks (RNN; Hochreiter & Schmidhuber 1997; Cho et al. 2014) have emerged as powerful alternatives to feature engineering of astronomical time-series (e.g., Naul et al. 2018). Given sufficient training

data, CNN/RNNs could learn to compress the “high-dimensional” raw observations into “low-dimensional” feature vectors—automatically learning to produce features that are useful for downstream tasks such as classification or regression. Vermaak (2003) applied a more basic form of the neural network—the multilayer perceptron (MLP)—to predict 2LIS parameters on simulated noise-free light curves, and achieved a success rate of 68% when the MLP results were further refined with Nelder–Mead optimization (Nelder & Mead 1965). However, there remains a large gap between the proof-of-concept work of Vermaak (2003) and application to real data due to the omission of noise and restrictions in parameter space. Additionally, machine learning has also been previously applied to discover and classify microlensing events (Wyrzykowski et al. 2015; Godines et al. 2019; Mróz 2020).

In addition to advances in this “representation learning,” neural networks have also enjoyed significant progress in modeling probability distributions, otherwise known as neural density estimation (NDE), where the fundamental task is to learn distributions from samples of that distribution. Both autoregressive models (Germain et al. 2015; Oord et al. 2016) and flow-based models (Dinh et al. 2017; Papamakarios et al. 2017) are NDEs that are highly capable of modeling complicated and multimodal distributions, which can not only evaluate probability densities, but also sample from that distribution. NDEs thus allow for flexible uncertainty quantification and degenerate solutions, which were not possible in Vermaak (2003).

The advancement in feature learning and NDE has allowed for accelerated progress in the field of likelihood-free inference (LFI), also known as simulation-based inference, which has been motivated by inference problems without a tractable likelihood. LFI is an umbrella term that encompasses a wide range of inference algorithms that do not require explicit evaluation of the likelihood. Under our particular LFI approach called amortized neural posterior estimation, an NDE learns a surrogate posterior as an observation-parameterized conditional probability distribution, from precomputed simulations over the full prior space of interest. A “featurizer” neural network is employed to compress raw observation into a feature vector that parameterizes the NDE. Inference is amortized in that all of the computation cost of simulation is paid upfront—likelihood evaluation with the slow forward simulator is no longer required, thus allowing for fast inference. For other neural LFI instances, neural networks could learn the likelihood (Papamakarios et al. 2019) or the likelihood-ratio (Thomas et al. 2020) as surrogates to accelerate sampling-based inference algorithms like MCMC (see Cranmer et al. 2020 for an overview).

In this paper, we present an LFI approach for binary microlensing where an NDE learns a surrogate posterior $\hat{p}(\theta|\mathbf{x})$ as an observation-parameterized conditional distribution from (\mathbf{x}^i, θ^i) samples of simulated microlensing light curves (\mathbf{x}^i) with the associated microlensing parameters (θ^i) . After training, the NDE can automatically generate posterior samples for future observations effectively in real-time. Because of the speed and performance without supervision by domain experts, the approach we introduce here has great potential for batch inference tasks such as those posed by Roman. Our preliminary results were reported as an extended abstract in Zhang et al. (2020). The work herein supersedes and expands upon that work.

We first lay out our inference framework in Section 2. Training set construction under the context of Roman is discussed in Section 3. In Section 4, we demonstrate the ability of the NDE to capture degenerate solutions and also present a systematic evaluation of the NDE performance over a large number of test events. In Section 5, we suggest future directions including a potential addition of a downstream MCMC algorithm to refine the NDE posterior into the exact posterior, with minimal additional computation time.

2. Method

NDEs are neural networks that are capable of learning distributions from samples. We train an NDE to learn a surrogate posterior $\hat{p}(\theta|\mathbf{x})$ as an observation-parameterized conditional distribution from (\mathbf{x}^i, θ^i) samples of simulated microlensing light curves, where θ^i are the physical parameters and $\mathbf{x}^i \in \mathbb{R}^N$ is the light curve with N data points. The training objective is to minimize the Kullback–Leibler (KL) divergence (D_{KL}), or relative entropy, which is a measure of how one probability distribution (Q) is different from a reference probability distribution (P):

$$D_{\text{KL}}(P||Q) = \mathbb{E}_{\mathbf{x} \sim p(\mathbf{x})} \left[\log \left(\frac{p(\mathbf{x})}{q(\mathbf{x})} \right) \right]. \quad (1)$$

In this case, we would like to train a neural network that minimizes the KL divergence from the NDE surrogate posterior $\hat{p}(\theta|\mathbf{x})$ to the true posterior $p(\theta|\mathbf{x})$:

$$\begin{aligned} \phi &= \text{argmin}(D_{\text{KL}}(p(\theta|\mathbf{x})||\hat{p}_\phi(\theta|\mathbf{x}))) \\ &= \text{argmin}(\mathbb{E}_{\theta \sim p(\theta), \mathbf{x} \sim p(\mathbf{x}|\theta)} [\log(p(\theta|\mathbf{x})) - \log(\hat{p}_\phi(\theta|\mathbf{x}))]) \\ &= \text{argmax}(\mathbb{E}_{\theta \sim p(\theta), \mathbf{x} \sim p(\mathbf{x}|\theta)} [\hat{p}_\phi(\theta|\mathbf{x})]), \end{aligned} \quad (2)$$

where ϕ represent parameters of the neural network, and \mathbb{E} denotes the mathematical expectation over the specified distribution.

In light of Equation (2), the NDE is therefore trained through maximum likelihood estimation (MLE) on a training set with physical parameters drawn from the prior $p(\theta)$ and light curves drawn from the likelihood function, which is the Poisson measurement noise model on top of the noise-free microlensing light curve $g(\theta)$ (in the number of photons) which, for simplicity, is assumed to be in the Gaussian limit:

$$p(\mathbf{x}|\theta) = \mathcal{N}(\mu = g(\theta), \sigma = \sqrt{g(\theta)}). \quad (3)$$

The noise-free light curve, in turn, is determined by the baseline source flux (F_{source}), the magnification time-series produced by the microlensing physical forward model $A(\theta)$, and the constant blend flux, which is the flux from the lens star and any other star that is unresolved from the source star:

$$g(\theta) = A(\theta) \cdot F_{\text{source}} + F_{\text{blend}}. \quad (4)$$

We use a 20-block Masked Autoregressive Flow (MAF; Papamakarios et al. 2017) to model $\hat{p}(\theta|\mathbf{x})$, and a ResNet-GRU network to extract features (\mathbf{h}) from the light curve (\mathbf{x}). In the following discussion, we do not distinguish between $\hat{p}(\theta|\mathbf{x})$ and $\hat{p}(\theta|\mathbf{h})$ where the former is meant to refer to the “featurizer + NDE” model and the latter is meant to refer to the NDE model alone that is explicitly conditioned on \mathbf{h} . Figure 1 presents a diagram of our neural posterior estimation framework. The ResNet-GRU network is comprised of an 16-layer




Figure 1. Schematic illustration of the inference framework based on conditional NDE. The bottom left shows a microlensing light curve in arbitrary units, which is abstracted into the length-7200 vector (\mathbf{x}) above. The featurizer composed of a combination of ResNet and GRU, shown in the trapezoid, compresses the light curve into a low-dimensional feature vector \mathbf{h} . The masked autoregressive flow (MAF), composed of K blocks of masked autoencoder for density estimation (MADE), is shown in the dashed box. Each MADE block takes in the feature vector \mathbf{h} and predicts scaling (α) and shifting (μ) factors, which parameterizes an invertible affine transformation between adjacent random variables (e.g., \mathbf{z}_0 and \mathbf{z}_1) shown in the dotted box. The leftmost random variable is the mixture-of-Gaussian base distribution whereas the rightmost random variable (\mathbf{z}_K) is the posterior (θ).

1D ResNet (Residual Convolutional Network; He et al. 2016) and a two-layer GRU (Gated Recurrent Network; Cho et al. 2014). We describe the neural networks in detail below.

2.1. Masked Autoregressive Flow

The MAF belongs to a class of NDE called normalizing flows, which models the conditional distribution $\hat{p}(\theta|\mathbf{x})$ as an invertible transformation f from a base distribution $\pi_z(\mathbf{z})$ to the target distribution $\hat{p}(\theta|\mathbf{x})$. The base density $\pi_z(\mathbf{z})$ is required to be fast to evaluate and is typically chosen to be either a standard Gaussian or a mixture of Gaussians for the MAF. The basic idea is that if the MAF, conditioned on the observation \mathbf{x} , could learn to map the posterior to a standard Gaussian, then the inverse transformation could enable sampling of the posterior by simply sampling from that standard Gaussian.

As binary microlensing events often exhibit degenerate, multimodal solutions, we use a mixture of eight standard multivariate Gaussians, each with eight dimensions, as the base distribution. The posterior probability density $\hat{p}(\theta|\mathbf{x})$ is evaluated by applying the inverse transformation f^{-1} from θ to \mathbf{z} :

$$\hat{p}(\theta|\mathbf{x}) = \pi_z(f^{-1}(\theta)) \left| \det \left(\frac{\partial f^{-1}}{\partial \theta} \right) \right|, \quad (5)$$

where $\pi_z(f^{-1}(\theta))$ represents the probability density for the base distribution (π_z) evaluated at $f^{-1}(\theta)$, while the second term—the determinant of the Jacobian—corresponds to the “compression” of probability space.

The MAF is built upon blocks of affine transformations where the scaling and shifting factors for each dimension are computed with a Masked Autoencoder for Distribution Estimation (MADE; Germain et al. 2015). For a simple one-block case, the inverse transformation from θ to \mathbf{z} is

expressed as:

$$\mathbf{z}_i = (\theta_i - \mu_i) \cdot \exp(-\alpha_i), \quad (6)$$

In the above equation,

$$\mu_i = f_{\mu_i}(\theta_{1:i-1}; \mathbf{x}) \quad (7)$$

$$\alpha_i = f_{\alpha_i}(\theta_{1:i-1}; \mathbf{x}) \quad (8)$$

are the scaling and shifting factors modeled by MADE subject to the autoregressive condition that the transformation of any dimension can only depend on those prior to it according to a predetermined ordering. This allows the Jacobian of f^{-1} to be triangular, whose absolute determinant can be easily calculated as:

$$\left| \det \left(\frac{\partial f^{-1}}{\partial \theta} \right) \right| = \exp \left(-\sum_i \alpha_i \right), \quad (9)$$

where $\alpha_i = f_{\alpha_i}(\theta_{1:i-1}; \mathbf{x})$.

To sample from the posterior, the forward transformation $\theta = f(\mathbf{z})$, where $\mathbf{z} \sim \pi_z$, is applied:

$$\theta_i = \mathbf{z}_i \cdot \exp \alpha_i + \mu_i. \quad (10)$$

In the above equation, μ_i and α_i are computed in the same manner as the inverse transformation.

The MAF is built by stacking many such affine transformation blocks, M_1, M_2, \dots, M_K , where M_K models the invertible transformation f_K between the posterior (\mathbf{z}_K) and intermediate random variable \mathbf{z}_{K-1} , M_{K-1} models that between intermediate random variables \mathbf{z}_{K-1} and \mathbf{z}_{K-2} and so on, and finally the base random variable \mathbf{z}_0 is modeled with the mixture-of-Gaussian distribution. M_1 also computes the mixture weights. The composite transformation can be written

as $f = f_1 \circ f_2 \circ \dots \circ f_K$ and the posterior probability density is now:

$$\hat{p}(\theta|\mathbf{x}) = \pi_z(f^{-1}(\theta)) \prod_{i=1}^K \left| \det \left(\frac{\partial f_i^{-1}}{\partial \mathbf{z}_{i-1}} \right) \right|, \quad (11)$$

where it is understood that \mathbf{z}_K defined to be θ . The log-probability of the posterior is then, by Equation (9),

$$\begin{aligned} \log \hat{p}(\theta|\mathbf{x}) &= \log[\pi_z(f^{-1}(\theta))] + \sum_{i=1}^K \log \left| \det \left(\frac{\partial f_i^{-1}}{\partial \mathbf{z}_{i-1}} \right) \right| \\ &= \log[\pi_z(f^{-1}(\theta))] - \sum_{i=1}^K \sum_{j=1}^N \alpha_j^i, \end{aligned} \quad (12)$$

where α_j^i is the j th component of the scale factor in M_i , as in Equation (6). This serves as the optimization objective (see Section 3.3).

Autoregressive models are sensitive to the order of the variables. The original MAF paper uses the default order for the autoregressive layer closest to θ and reverses the order for each successive layer. In this work, we adopt fixed random orderings for each MAF block, which we find to allow for better expressibility. The random seed of the ordering serves as a hyperparameter to be optimized on.

2.2. Featurizer Network

A custom 1D ResNet with a downstream two-layer GRU is used as the light curve featurizer, which takes preprocessed light curves (\mathbf{x}) as input and outputs a low-dimensional feature vector (\mathbf{h}). The ResNet used in this study shares the identical architecture as Zhang & Bloom (2020; except for hyperparameters) and consists of eight identical residual blocks, each of which is composed of two convolutions followed by layer normalization (Ba et al. 2016). A residual connection is added between each adjacent residual block, which acts as a “gradient highway” to assist network optimization. A MaxPool layer is applied in between every two ResNet layers, where the sequence length is reduced by half and the feature dimension is doubled until a specified maximum. This results in an output feature map of length $L = 56$ and dimension $D = 256$ that is then fed into the GRU network which sequentially processes information across the temporal dimension and outputs a single vector of $D = 256$, which then serves as the conditional input to the MAF.

3. Data

Training data is generated within the context of the Roman Space Telescope Cycle-7 design (see Penny et al. 2019). We first simulate 10^6 2L1S magnification sequences with the microlensing code *MulensModel* (Poleski & Yee 2019); each sequence contains 144 days at a cadence of 0.01 days corresponding to the planned Roman cadence of 15 minutes (Penny et al. 2019). These sequences are chosen to have twice the length of the 72-day Roman observation window to facilitate sampling from a $t_0 \sim \text{Uniform}(0, 72)$ prior (see Section 3.1). We then fit each simulated magnification time-series with a Paczyński 1L1S model (assuming $S/N_{\text{base}} = 200$ and $f_s = 1$; see Section 3.2) and discard those that are consistent with 1L1S ($\chi^2/\text{dof} < 1$). This results in a final data set of 291,012 light curves, among which 95% (276,461) are used as a training set and the remaining 5% (14,551) as a test set.




Figure 2. Fraction of the 10^6 simulations passing the $\chi^2_{\text{ILIS}}/\text{dof} < 1$ cutoff as a function of each parameter, shown in the gray histograms. The original analytic priors used to generate the 10^6 simulations are shown in red dashed lines up to a normalization factor. For parameters with a flat original prior, the gray histogram is also the effective training set prior up to a normalization factor. The t_0 and f_s distributions follow the original priors as they are sampled on the fly during training.

3.1. Prior

Assuming rectilinear relative motion of the observer, lens, and source, binary microlensing (2L1S) events are characterized by eight parameters: binary lens separation (s), mass ratio (q), angle of the source trajectory with respect to the projected binary lens axis (α), impact parameter (u_0), time of closest approach (t_0), Einstein ring crossing timescale (t_E), finite source size (ρ), and source flux fraction (f_s). α is the angle between the vector pointing from the primary to the secondary and the source trajectory vector, measured counterclockwise in degrees. u_0 and t_0 are defined with respect to the binary lens center of mass (COM). Where applicable, the parameters are normalized to the Einstein ring length-scale or the Einstein ring crossing timescale of the total mass of the lens system. t_0 and t_E are in units of days. We simulate 2L1S events based on the following analytic priors:

$$\begin{aligned} s &\sim \text{LogUniform}(0.2, 5) \\ q &\sim \text{LogUniform}(10^{-6}, 1) \\ \alpha &\sim \text{Uniform}(0, 360) \\ u_0 &\sim \text{Uniform}(0, 2) \\ t_0 &\sim \text{Uniform}(0, 72) \\ t_E &\sim \text{TruncLogNorm}(1, 100, \mu = 10^{1.15}, \sigma = 10^{0.45}) \\ \rho &\sim \text{LogUniform}(10^{-4}, 10^{-2}) \\ f_s &\sim \text{LogUniform}(0.1, 1). \end{aligned} \quad (13)$$

We note that because of the $\chi^2_{\text{ILIS}}/\text{dof} < 1$ cutoff, the effective prior is the parameter distribution for the 276,461 training set simulations, different from the prior above. As shown in Figure 2, large $\log q$ and small u_0 , which otherwise have flat priors, are strongly preferred.

During training, a random 72-day segment is chosen on the fly from each 144-day magnification sequence, equivalent to prescribing a uniform prior on t_0 . The truncated normal distribution for t_E is an approximation of a statistical analysis based on OGLE-IV data (Mróz et al. 2017). The lower limit of $q = 10^{-6}$ corresponds to the mass ratio between Mercury and a low-mass ($M \sim 0.1 M_\odot$) M-dwarf star, highlighting the superb sensitivity of Roman. The source flux fraction is defined as the

ratio between the source flux and the total baseline flux

$$f_s = \frac{F_{\text{source}}}{F_{\text{source}} + F_{\text{blend}}}. \quad (14)$$

3.2. Light-curve Realization

The magnification sequences are converted into light curves during training on the fly by multiplying with the baseline premagnification source flux before adding the constant blend flux and applying measurement noise. For simplicity, we only consider photon-counting noise from the lens and fixed blend flux, assumed to be in the Gaussian limit of the Poisson noise (Equation (3)), where the standard deviation of each photometric measurement is the square root of the flux measurement in photon counts. Studies of the bulge star population show that the apparent magnitude largely lies within the range of 20–25 mag (Penny et al. 2019; Figure 5). The Roman/WFIRST Cycle 7 design has the zero-point magnitude (1 count s^{−1}) at 27.615 mag for the W149 filter. With an exposure time at 46.8 s, the aforementioned magnitude range corresponds to signal-to-noise ratios (S/N_{base}) between 230 and 23 for the baseline flux, which we randomly and uniformly sample during training. On-the-fly sampling of S/N_{base} and f_s also serves as data augmentation, which refers to the process of expanding the effective size of the training set.

3.3. Preprocessing and Training

Network optimization is performed with ADAM (Kingma & Ba 2015) at an initial learning rate of 0.001 and batch size 512, which decays to 0 according to a cosine annealing schedule (Loshchilov & Hutter 2017) for 250 epochs, at which point the training terminates. To ensure that there is no overfitting, we first reserved 20% of the training set as a validation set. After confirming the absence of overfitting, we then proceed with the full training set. We apply data augmentation on α by changing the direction of the source trajectory: the temporal order of each sequence is reverted and α becomes $-(\alpha + 180) \bmod 360$. Each training epoch takes ~ 6 minutes on four NVidia GTX 2080 Ti GPUs with a total training time of around 25 hours. As an evaluation metric, the final average negative log-likelihood (NLL) is -16.316 on the training set and -16.177 on the test set, where a lower value represents a better model fit to the data.

4. Results

The trained model is able to generate accurate and precise posterior samples at a rate of 10⁵ per second on one GPU, effectively in real-time. This is much faster compared to the ~ 1 per second simulation speed of the forward model `MulensModel` on one CPU core. In this section, we first highlight the ability of the NDE to capture multimodal solutions by providing NDE posteriors of representative events where we set the baseline S/N_{base} = 200. Then, the quality of NDE posteriors is systematically analyzed by examining the accuracy and calibration properties on a test set of 14,511 simulated light curves.

4.1. Central-caustic Passing Event

Figure 3(a) shows the NDE posterior for an example central-caustic passing event where a classic “close-wide” degeneracy is clearly exhibited by the s - $1/s$ behavior (Griest & Safizadeh 1998;

Table 1
Solutions for the Example Central-caustic Passing Event

	Truth	Close	Wide
$\log_{10}(s)$	−0.200	−0.1923 ^{+0.0034} _{−0.0036}	0.2301 ^{+0.0049} _{−0.0040}
$\log_{10}(q)$	−2.000	−2.0252 ^{+0.0144} _{−0.0147}	−1.9761 ^{+0.0177} _{−0.0155}
α	300.000	299.8524 ^{+0.2457} _{−0.2658}	299.5919 ^{+0.2469} _{−0.3139}
u_0	0.050	0.0505 ^{+0.0002} _{−0.0002}	0.0403 ^{+0.0007} _{−0.0010}
t_0	26.000	26.0027 ^{+0.0051} _{−0.0063}	26.1054 ^{+0.0131} _{−0.0075}
$\log_{10}(t_E)$	1.301	1.2993 ^{+0.0007} _{−0.0008}	1.3020 ^{+0.0010} _{−0.0009}
$\log_{10}(\rho)$	−2.301	−2.2075 ^{+0.0044} _{−0.1133}	−2.3421 ^{+0.0737} _{−0.0210}
f_s	0.120	0.1211 ^{+0.0003} _{−0.0003}	0.1214 ^{+0.0004} _{−0.0003}

Note. t_E and t_0 are in units of days, α is in units of degrees, and u_0 , s , and ρ are in units of θ_E . Uncertainties are 1 σ marginal uncertainties.

Table 2
Solutions for the Example Resonant-caustic Passing Event

	Truth	Close	Resonant
$\log_{10}(s)$	0.000	−0.0484 ^{+0.0017} _{−0.0006}	0.0018 ^{+0.0010} _{−0.0003}
$\log_{10}(q)$	−3.301	−3.3008 ^{+0.0098} _{−0.0201}	−3.2858 ^{+0.0194} _{−0.0121}
α	110.000	109.7839 ^{+0.2044} _{−0.1526}	109.9666 ^{+0.1657} _{−0.2093}
u_0	0.100	0.1004 ^{+0.0003} _{−0.0003}	0.1003 ^{+0.0003} _{−0.0003}
t_0	26.000	25.9980 ^{+0.0048} _{−0.0064}	26.0014 ^{+0.0047} _{−0.0062}
$\log_{10}(t_E)$	1.301	1.3012 ^{+0.0007} _{−0.0008}	1.3012 ^{+0.0008} _{−0.0007}
$\log_{10}(\rho)$	−2.301	−2.5335 ^{+0.0492} _{−0.2505}	−2.5325 ^{+0.0041} _{−0.4117}
f_s	0.200	0.1996 ^{+0.0006} _{−0.0005}	0.1995 ^{+0.0006} _{−0.0005}

Note. Same units as Table 1. Uncertainties are 1 σ marginal uncertainties.

Dominik 1999). Table 1 presents the ground truth 2L1S parameters of this event as well as the “close” and “wide” solutions, calculated as the modes of their respective distributions. The fact that f_s is slightly underestimated is related to a systematic effect as discussed in Section 4.4. Although the source is expected to pass the caustic center at the same distances for the two cases, Figure 3(a) shows a bimodal solution for u_0 as well because u_0 has been defined with respect to the COM, rather than the caustic center. While the caustic center is centered on the COM for close-separation events, for wide-separation events there is an offset from the COM of

$$\delta = \frac{s \cdot q}{1 + q} - \frac{q}{s \cdot (1 + q)}, \quad (15)$$

where the first term accounts for the offset of the primary from the COM, and the second term for the offset of the caustic center from the location of the primary star (Han 2008). Negative offset is directed toward the companion and vice versa. Plugging in the wide solution, we expect an offset of $\Delta u_0 = 0.0101$, which is close to the actual $\Delta u_0 = 0.0099$ of the NDE posterior solutions. Magnification curves of the two solutions, as well as the ground truth are plotted in Figure 3(b), which are hardly distinguishable from one another. Figure 3(c) shows the caustic structures of the two degenerate solutions.

4.2. Resonant-caustic Passing Event

We also highlight an example of a resonant-caustic passing event, whose parameters and solutions are shown in Table 2. As illustrated in Figure 4, the NDE finds an additional solution at

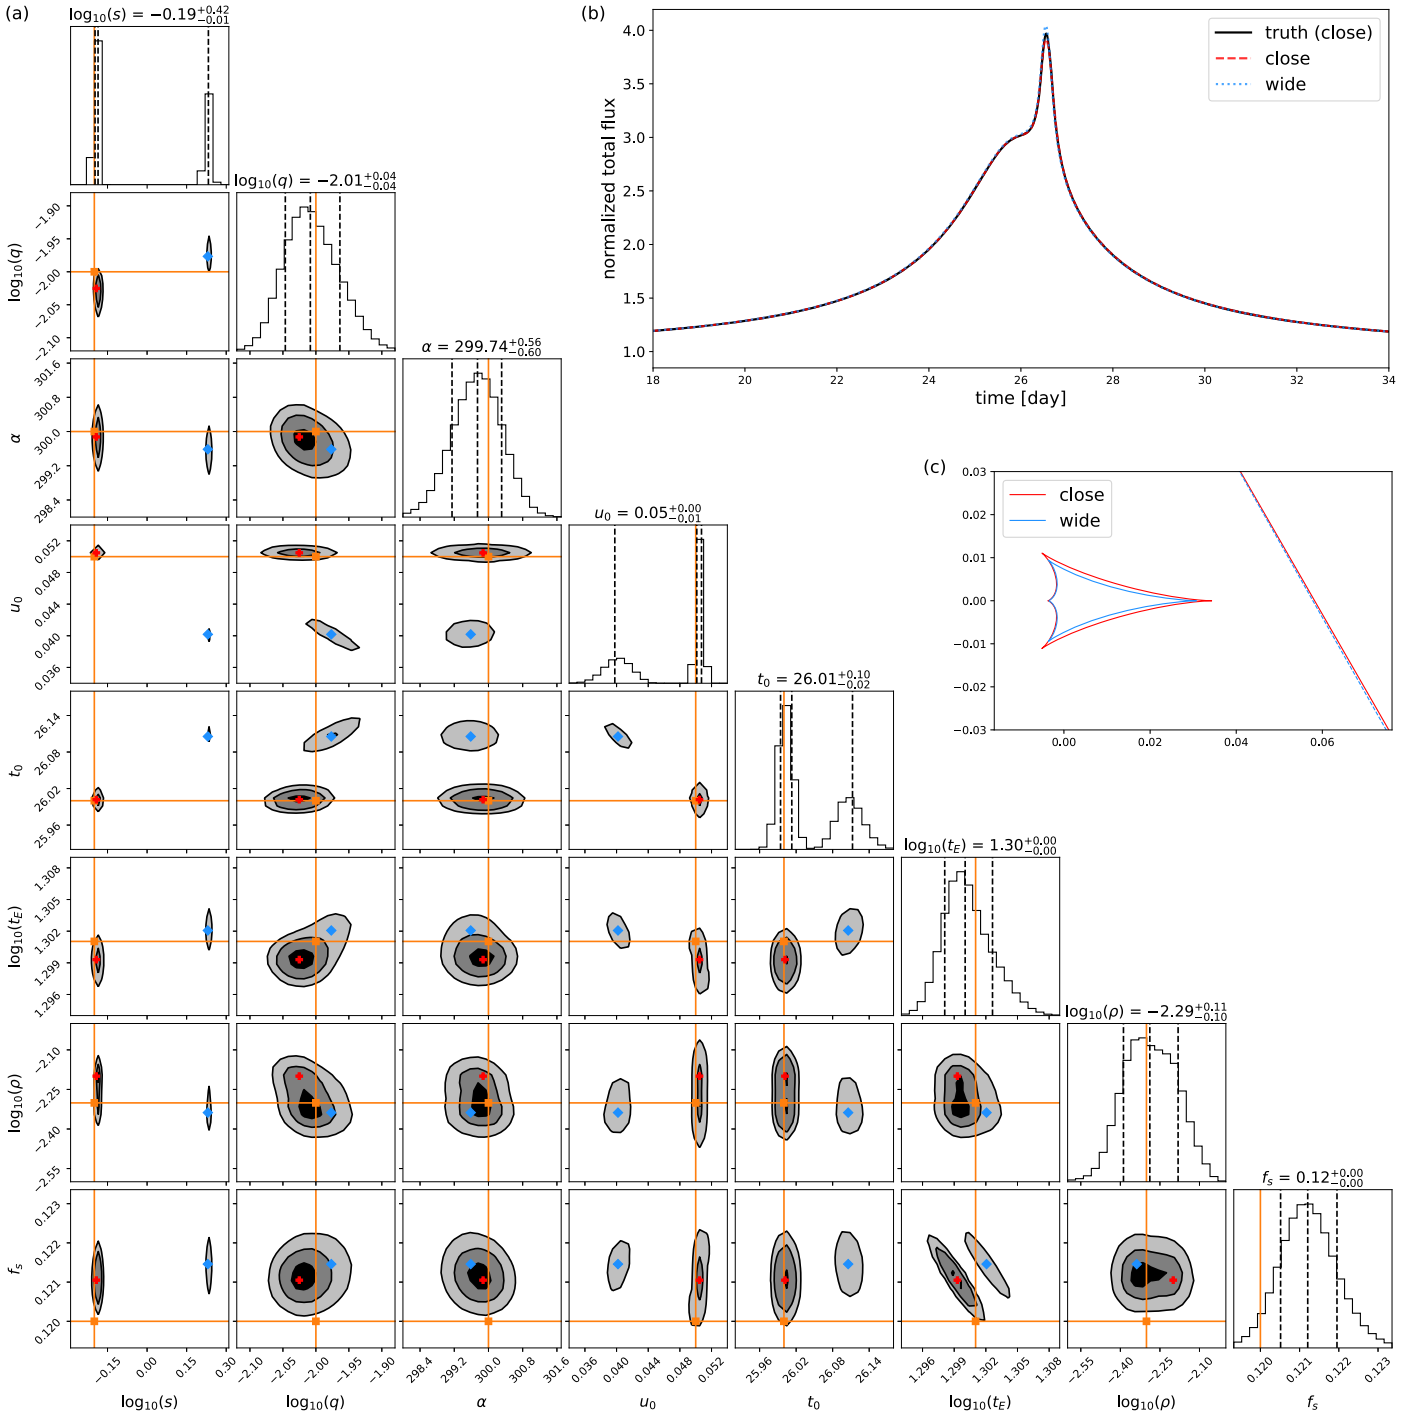


Figure 3. (a) NDE posterior for a central-caustic passing event. t_E and t_0 are in units of days, α is in units of degrees, and u_0 , s , and ρ are in units of θ_E . Filled contours show $1/2/3\sigma$ regions. The ground truth close solution is marked with orange cross-hairs. The close and wide solutions are marked with a red cross and a blue diamond, respectively. (b) Close-up view of the light-curve realizations normalized to the minimum fluxes for both solutions, in the same color-coding as the left panel. The 0.01 day cadence and measurement noise is negligibly small on the scale of the figure, and therefore not shown. (c) Caustic structures as well as trajectories for the two solutions in the same color-coding, centered on the center of caustic.

$s < 1$, whose triangular caustics are causing a similar weak demagnification as the resonant caustics (also see Figure 7 in Gaudi 2010). This type of degeneracy has been previously observed in the microlensing event OGLE-2018-BLG-0677Lb (Herrera-Martín et al. 2020). Additionally, strong covariances are seen among u_0 , t_E , and f_s , as are also seen in the previous example (Section 4.1). As first observed by Woźniak & Paczyński (1997), in the $f_s \ll 1$ and $u_0 \ll 1$ regime where the baseline flux is dominated

by the blend flux, there is strong degeneracy between the three parameters for 1L1S events. While the binary perturbations break some of that degeneracy, strong covariances remain.

4.3. Binary-planetary Degeneracy

We also provide a fascinating five-fold-degenerate example that is similar to the degeneracy reported in Choi et al. (2012) where a

light curve that is blunt and flat near the peak can be explained by either a binary case or a planetary case. Here, we simulate a close-topology, planetary-mass ratio ($q = 10^{-1.7}$) event where the source trajectory passes through the negative perturbation region toward the back end of the arrowhead-shaped central caustic as in the case of the “planetary A/B” caustic in Figure 5. Choi et al. (2012) noted that a similar perturbation can occur for the binary case when the source trajectory passes through the negative perturbation region between two adjacent cusps of the astroid-shaped central caustic, as in case of the “binary A/B/C” caustics in Figure 5; also see Figure 1 in their paper.

As shown in Figure 5, all five degenerate solutions cause magnification patterns that are hardly distinguishable. The solutions are provided in Table 3. The two planetary solutions exhibit a close-wide degeneracy. For the three binary solutions, the “binary B/C” solutions suggest two possible trajectories ($\sim 90/270$ deg) for the same lens system configuration whereas the “binary A” solution exhibits a smaller mass ratio and a wider binary separation than “binary B/C.” We note that this additional degeneracy in the mass ratio for the binary case was not reported in Choi et al. (2012). It is not clear if this is a discrete or continuous degeneracy, if it is an “accidental degeneracy” that arises because of the relatively weak perturbation, or if it is due to some underlying symmetry in the binary lens equation (e.g., Dominik 1999).

On the other hand, wide solutions for the binary case are largely absent from the NDE posterior, apart from an inkling of density near $\log_{10}(s) \sim 0.69$, which points to the expected close-wide degeneracy for the binary solution. We note that the reason those degenerate solutions are excluded is that, because of the offset between the COM and the central caustic (Equation (15)), wide-binary solutions would require $t_0 < 0$, which has a prior probability of zero.

4.4. Evaluating Performance

We present a systematic evaluation of all 14,551 test set events in the form of predicted versus true scatter plots (Figure 6). Each test event light curve is realized in the same fashion as training time. As the NDE returns potentially multimodal posteriors of arbitrary shape, we compute the mode(s) for the marginal 1D distributions of the posterior and consider the mode closest to the ground truth as the “predicted” value. The mode(s) is/are computed by first fitting each with a 1D histogram of 100 bins and then searching for local maxima defined as any bin count higher than that of the 20 adjacent bins. This limits the number of modes to 5. Considering the purpose of the NDE posterior is to allow ultra-fast convergence of a downstream sampling-based algorithm like MCMC to determine the exact posterior, as long as the correct solution has substantial density in the NDE posterior, it should not raise alarm if an alternative mode is mistakenly favored. Any degeneracies can be easily resolved downstream. Therefore, it is sensible to allow the correct mode to be used as the predicted value, even if another degenerate mode is incorrectly preferred.

As shown in the upper left corner of each subplot in Figure 6, all parameters are constrained at a rate of close to 100% except for the finite source effect for which only 14.2% is constrained, as the source trajectory is required to either cross or pass close to a caustic for ρ to be determined.⁵ We consider a

parameter to be constrained if the probability density of the 1D marginal distribution is more than twice the prior probability density at the global mode.

The second row in each upper left corner shows the frequency for which the correct mode is preferred by NDE, that is, the ground truth is the closest to the major mode compared to the minor mode(s), if any. If the ground truth is closer to a minor mode, the major mode is plotted in blue while the minor mode is shown in orange. We see clear degeneracy patterns in $\log_{10}(s)$ and α . For $\log_{10}(s)$, the “wide-close” degeneracy is exhibited by the cluster around the upper left to lower right diagonal. For α , there is also a cluster of events along the same diagonal, indicating a degeneracy between α and $-\alpha$. Such a degeneracy may happen for nearly symmetrical central caustics along the direction perpendicular to the lens axis.

The 1σ and 2σ ranges of prediction, shown by red shadows, are clustered closely around the diagonal for most parameters. We emphasize that the loose 1LIS-fitting cutoff ($\chi^2_{1LIS}/\text{dof} < 1$) means many of the test set light curves are only weakly perturbed by the binary nature of the lens, and should explain a number of cases in which the mass ratio is poorly constrained. Interestingly, we find that there is a tendency to overestimate the mass ratio in these cases. In addition, we notice that u_0 and f_s are underestimated for a large number of cases while t_E is correspondingly overestimated, though hardly visible in Figure 6. This bias could be explained by the combined effect of a known degeneracy for 1LIS events and a distribution mismatch.

First, there exists a well-known degeneracy between u_0 , f_s , and t_E for single-lens events, which in our case applies to events that are only weakly perturbed by the binary nature of the lens. As demonstrated by Woźniak & Paczyński (1997), this degeneracy is most severe for low magnification events ($u_0 \gg 1$), which is precisely where the biases occur as seen in Figure 6. Indeed, restricted to test events with $u_0 < 0.15$, the bias in f_s and t_E is largely removed. Figure 7 shows the NDE posterior for an example $u = 1.5$ 1LIS event that demonstrates the strong degeneracy among u_0 , f_s , and t_E .

In the presence of strong degeneracies as such, the effective likelihood implicitly provided by the featurizer is only marginally informative. In other words, the featurizer cannot distinguish among solutions within the continuous degeneracy, and only prescribes a region in parameter space where the observation is about equally likely. Therefore, the posterior is essentially dominated by the prior, which strongly favors small u_0 and f_s , as seen in Figure 7. Had the parameters for the weakly perturbed events in the test set been drawn from the same effective prior as the full training set, there would be little bias (under/overestimation) at all in Figure 6. However, quite the contrary, the distribution of the weakly perturbed is weighted toward the exact opposite direction of effective prior, e.g., toward large u_0 and small $\log_{10}(q)$ —those more likely to be excluded from the $\chi^2_{1LIS}/\text{dof} < 1$ cutoff. Because of this distribution mismatch, large u_0 and small $\log_{10}(q)$ occur much more often than expected by prior belief, thus resulting in the under/overestimation bias. And because of the strong covariances among u_0 , f_s , and t_E , the underestimation of u_0 translates into an underestimation of f_s and an overestimation for t_E (Figure 7), which explains the biases seen in Figure 6.

4.5. Calibration Properties

A perfectly calibrated posterior knows how often it is right or wrong. In other words, the quantile of the ground truth

⁵ Formally, effects on the light curve due to the finite size of the source are only significant if the gradient of the magnification across the source has a significant second derivative. In practice, this condition is only satisfied if the source passes within a few angular source radii of a caustic.

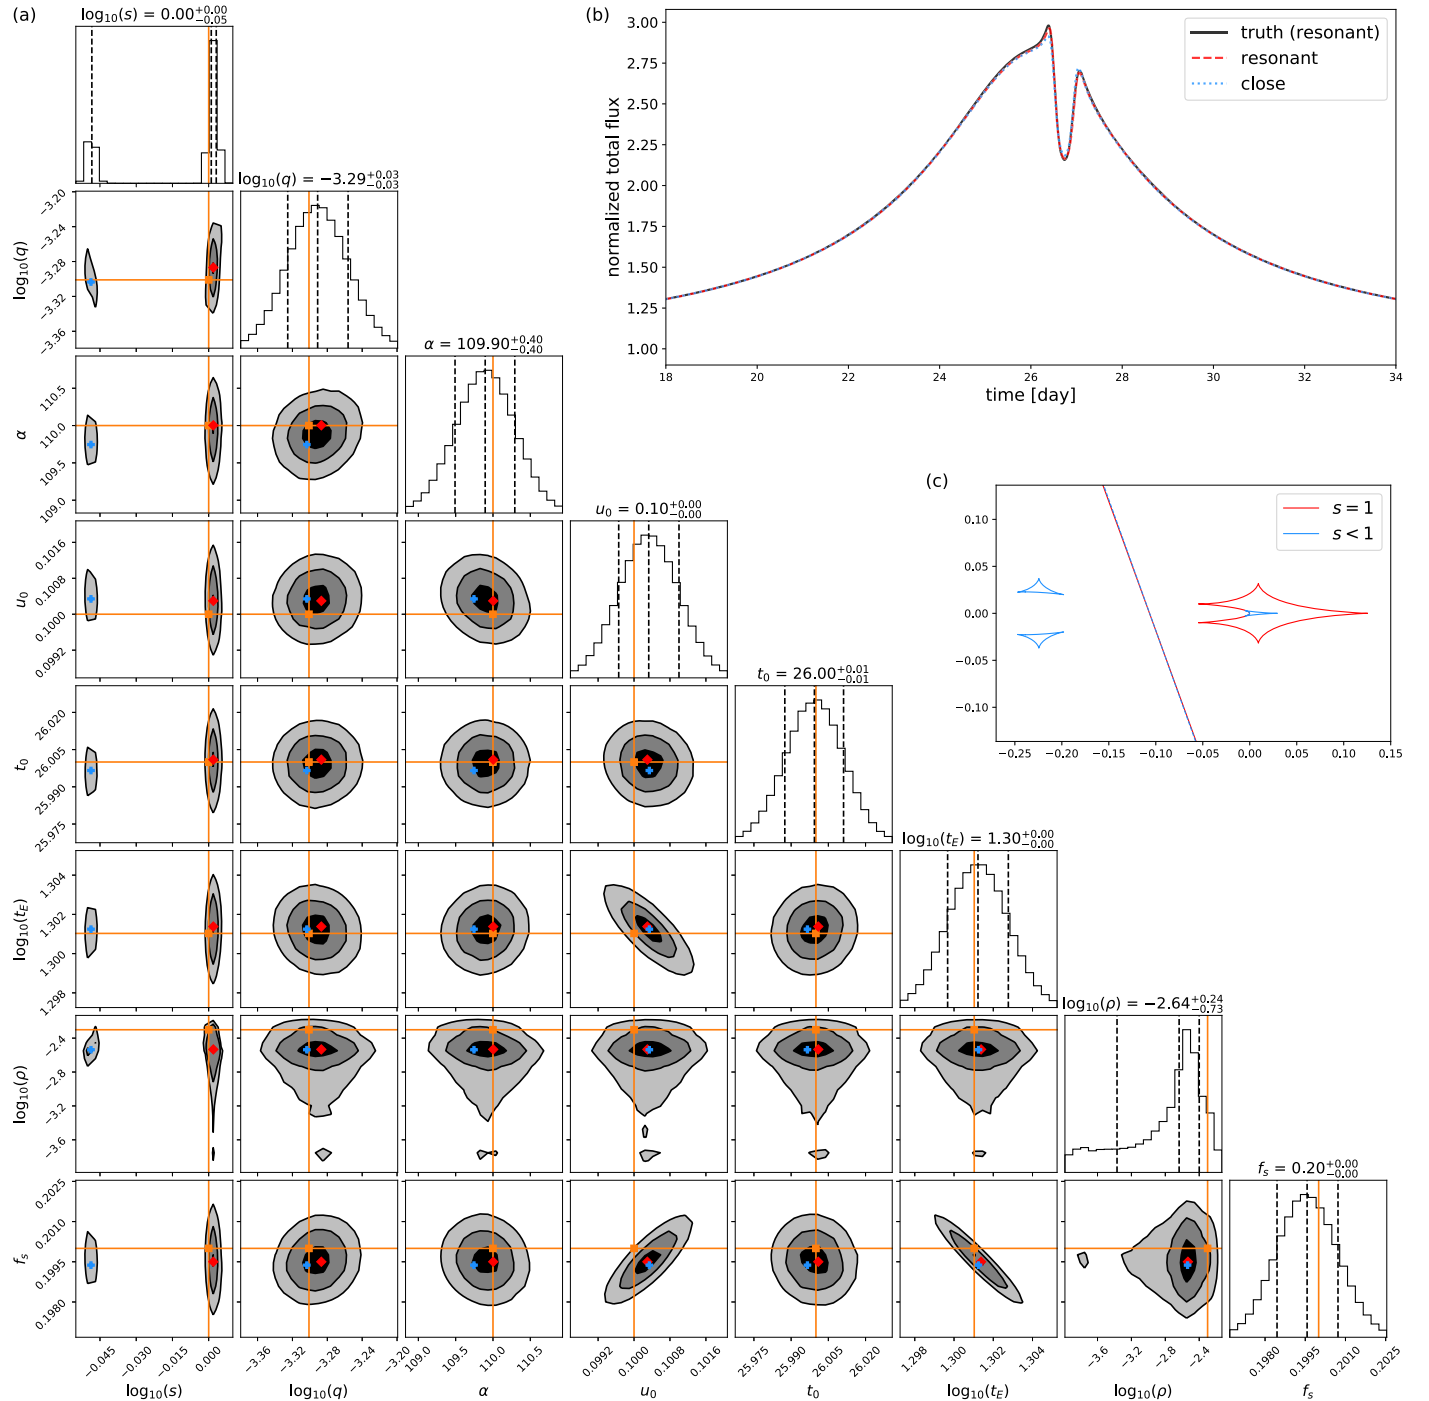


Figure 4. Resonant-caustic-passing event; same figure caption as Figure 3. Here, a degenerate solution is seen at $s < 1$, whose two triangular caustics cause a similar suppression pattern as the resonant caustic.

parameter under the NDE posterior should be expected to be distributed uniformly. Figure 8 shows the quantile distribution for the 1D marginal NDE posterior distributions for the same 14,551 test set inferences. The quantile distribution for $\log_{10}(q)$, $\log_{10}(\alpha)$, and t_0 is concave-up, indicating that the NDE uncertainty is overestimated and the true value lies closer to the center of the posterior more often than expected. This suggests that the NDE finds it hard to contract the posterior in those dimensions, possibly due to numerical optimization difficulties or insufficient neural network expressibility. On the other hand, distributions for the three parameters in the second

row— $\log_{10}(t_E)$, u_0 , and f_s —demonstrate the systematic under/overestimation as seen in Figure 6, where $\log_{10}(t_E)$ is systematically overestimated and u_0 and f_s are underestimated. The quantile distributions for $\log(s)$ and $\log_{10}(\rho)$ are consistent with uniform distributions and are thus well-calibrated.

5. Discussion and Conclusions

We have demonstrated that amortized neural posterior estimation, an LFI method that uses a conditional NDE to learn a surrogate posterior, $\hat{p}(\theta|\mathbf{x})$, greatly accelerates binary

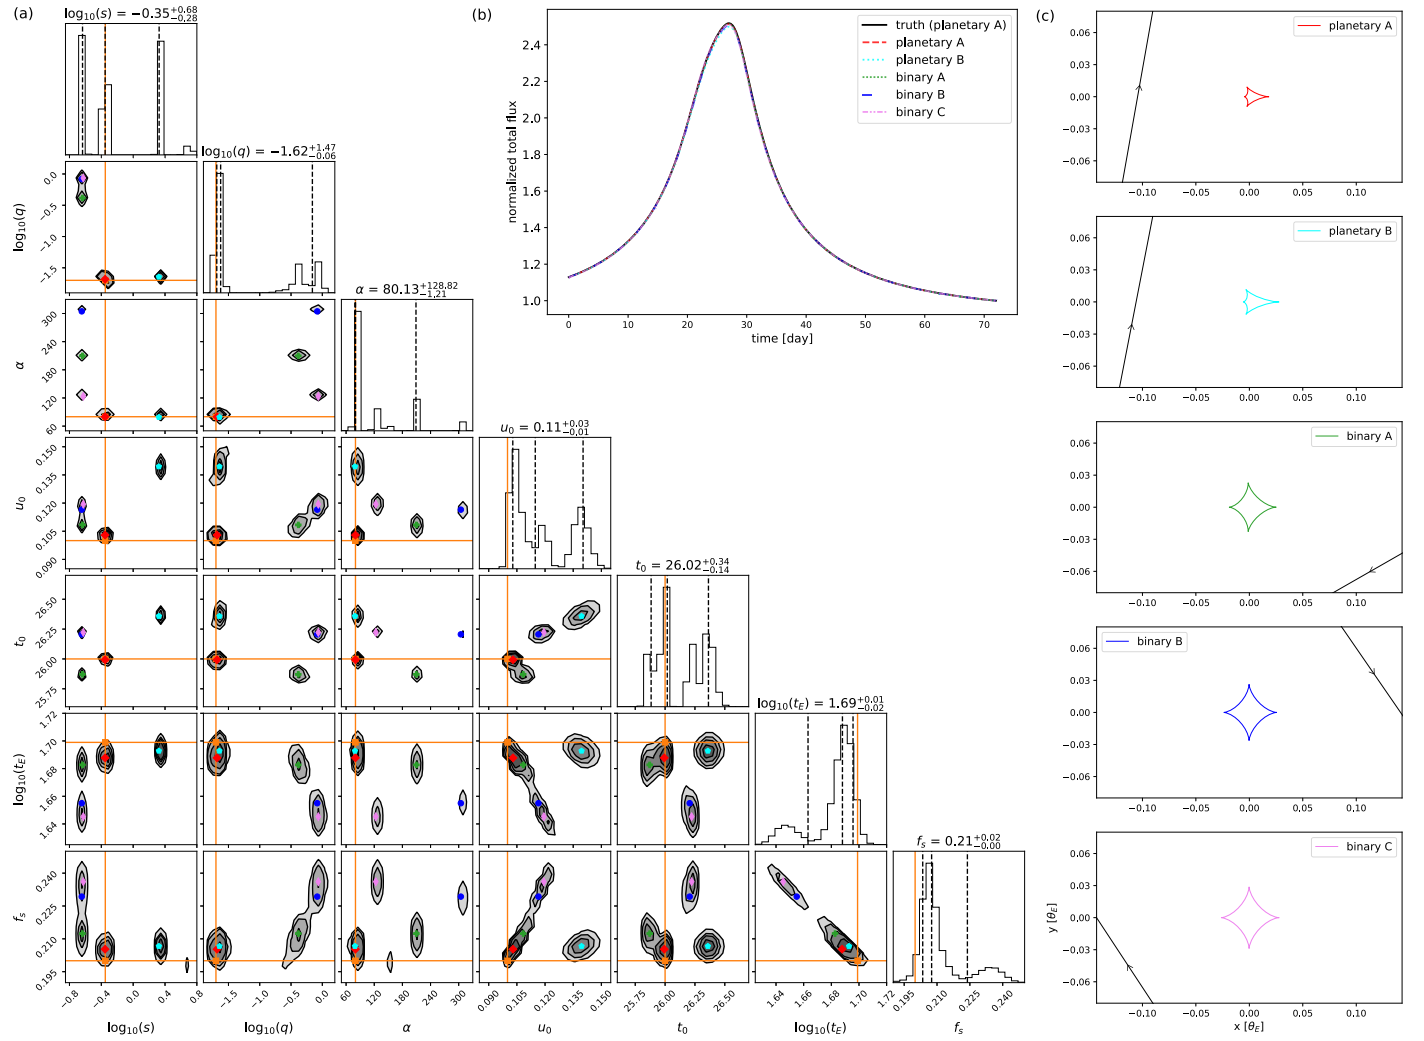


Figure 5. Example event exhibiting a blunt and flat light curve near the peak, which has a five-fold degenerate NDE posterior; same figure caption as Figure 3 for (a) and (b). (c) Caustic structures and source trajectories for the five solutions. The same color-coding is shared across the three panels.

Table 3
Degenerate Solutions for the Binary-planetary Degenerate Event Shown in Figure 5

	Truth	Planetary A	Planetary B	Binary A	Binary B	Binary C
$\log_{10}(s)$	−0.350	−0.3520 ^{+0.0037} _{−0.0049}	0.3242 ^{+0.0035} _{−0.0035}	−0.6373 ^{+0.0037} _{−0.0047}	−0.6450 ^{+0.0026} _{−0.0030}	−0.6267 ^{+0.0046} _{−0.0043}
$\log_{10}(q)$	−1.700	−1.6849 ^{+0.0275} _{−0.0140}	−1.6464 ^{+0.0190} _{−0.0110}	−0.3729 ^{+0.0250} _{−0.0273}	−0.0813 ^{+0.0297} _{−0.0250}	−0.0609 ^{+0.0162} _{−0.0244}
α	80.000	80.0207 ^{+0.2170} _{−0.3531}	79.0411 ^{+0.2682} _{−0.3430}	209.7867 ^{+0.8607} _{−0.8053}	304.7107 ^{+0.6110} _{−0.6557}	123.4429 ^{+0.2513} _{−1.3218}
u_0	0.100	0.1027 ^{+0.0009} _{−0.0006}	0.1390 ^{+0.0022} _{−0.0016}	0.1082 ^{+0.0010} _{−0.0011}	0.1160 ^{+0.0012} _{−0.0011}	0.1197 ^{+0.0009} _{−0.0013}
t_0	26.000	25.9926 ^{+0.0084} _{−0.0070}	26.3604 ^{+0.0245} _{−0.0169}	25.8701 ^{+0.0089} _{−0.0080}	26.2091 ^{+0.0065} _{−0.0128}	26.2230 ^{+0.0102} _{−0.0101}
$\log_{10}(t_E)$	1.699	1.6874 ^{+0.0035} _{−0.0016}	1.6927 ^{+0.0024} _{−0.0022}	1.6824 ^{+0.0038} _{−0.0025}	1.6550 ^{+0.0037} _{−0.0030}	1.6443 ^{+0.0045} _{−0.0022}
f_s	0.200	0.2048 ^{+0.0018} _{−0.0010}	0.2065 ^{+0.0015} _{−0.0011}	0.2114 ^{+0.0027} _{−0.0014}	0.2280 ^{+0.0030} _{−0.0015}	0.2357 ^{+0.0023} _{−0.0021}

Note. Same units as Table 1. Uncertainties are 1σ marginal uncertainties.

microlensing inference—an approximate posterior could be produced in seconds without the need for an expert in the loop. Our new approach is capable of capturing a variety of degeneracies. For future work, it is straightforward to extend to higher-level effects such as parallax and binary motion by introducing additional parameters. Application to more complex systems such as 3LIS may be fruitful, where the physical forward model is orders of magnitude slower. In addition, the photometric noise model adapted in our study is somewhat simplistic, and

future work can explore how to adapt models trained with ideal noise properties to fully realistic data with the help of image-based simulation pipelines such as ones used in Penny et al. (2019). We discuss two additional aspects of our work below.

5.1. A Hybrid NDE-MCMC Framework

The NDE posterior is easily validated and/or refined by a downstream MCMC sampler. While the NDE posterior is




Figure 6. Predicted vs. ground truth 2L1S parameters for 14,551 test set 2L1S events. t_E and t_0 are in units of days, α is in units of degrees, and u_0 , s , and ρ are in units of θ_E . Single-mode NDE posteriors are shown as black dots. For multimodel NDE posteriors, we color-code the solution as follows: those for which the global mode is closest to the ground truth are plotted in black; for cases where a minor mode is closest to the true value, this correct, minor mode is plotted in orange whereas the incorrect global mode is plotted in blue. Red shadows indicate 32–68th percentile (1σ) and 5–95th percentile (2σ) regions. Red dashed lines show the diagonal. In the upper left of each subplot, “constrain” refers to the percentage of events whose NDE posterior poses sufficient constraint—the peak posterior probability must be at least twice the prior probability. “Correct” refers to the percentage of constrained events whose true parameter lies closest to the global mode.




Figure 7. Corner plot for the marginal NDE posterior of a 1L1S event showing strong degeneracy among the three 1L1S parameters: u_0 in units of θ_E , t_E in units of days, and f_s . Filled contours show 1/2/3/4 σ regions. Small u_0 and f_s are strongly favored because of the effective priors (Section 3.1) and a marginally informative likelihood.

precise enough to allow for fast convergence of downstream MCMC typically within hundreds of steps, we do notice that the precision of the exact MCMC posterior could be more than




Figure 8. Calibration plot showing the test set distributions of the ground truth quantile for the 1D marginal NDE posteriors. Dashed lines indicated the uniform distribution as expected for a perfectly calibrated posterior.

order-unity higher in many cases. The precision of the NDE posterior is determined by two kinds of uncertainty: data uncertainty and the model uncertainty of the inference algorithm, the latter of which is negligible for MCMC. As neural networks in practice are not infinitely expressive, in the

limit of the highest-quality data, the NDE model uncertainty is expected to dominate over data uncertainty. This is the case for Roman data. Applied to much noisier and more sparsely sampled ground-based data, we expect that data uncertainties will dominate over model uncertainties, thus allowing the NDE posterior to converge toward the exact posterior.

5.2. Choice of Coordinate System

For all events in this work, we have adopted the COM coordinate system, which is the default in `MulensModel` but not the most efficient reference frame in the sense that more than 70% of the 1 million simulations turn out to be consistent with a 1L1S model. For example, most 2L1S configurations with large u_0 do not pass close to either the central caustics or the planetary caustics. For parts of the parameter space, alternative reference coordinates may be more descriptive or useful. For example, the caustic-center frame is preferred for binary and/or wide-separation events for which there is an offset of the caustic center from the COM. Doing so recovers the missing wide/binary solution in Section 4.3 without the need to expand the prior to include negative t_0 . Additionally, planetary-caustic passing events are also rare in the current COM coordinate set-up, because they require a narrow range of α . For future work, a hybrid and self-consistent coordinate system could be used.

K.Z. and J.S.B. are supported by a Gordon and Betty Moore Foundation Data-Driven Discovery grant. J.S.B. is partially sponsored by a faculty research award from Two Sigma. K.Z. thanks the LSSTC Data Science Fellowship Program, which is funded by LSSTC, NSF Cybertraining Grant 1829740, the Brinson Foundation, and the Moore Foundation; his participation in the program has benefited this work. K.Z. thanks Shude Mao and Tsinghua University for their hospitality during the COVID-19 pandemic. B.S.G. is supported by NASA grant NNG16PJ32C and the Thomas Jefferson Chair for Discovery and Space Exploration. This work is supported by the AWS Cloud Credits for Research program. We thank Yang Gao and Yu Sun for helpful discussions, and Weicheng Zang for pointing out issues with the initial figure presentation.

Software: `numpy` (van der Walt et al. 2011), `scipy` (Jones et al. 2001), `matplotlib` (Hunter 2007), `pytorch` (Paszke et al. 2017), `MulensModel` (Poleski & Yee 2018), `Jupyter` (Kluyver et al. 2016), `corner` (Foreman-Mackey 2016).

ORCID iDs

Keming Zhang (张可名)  <https://orcid.org/0000-0002-9870-5695>

Joshua S. Bloom  <https://orcid.org/0000-0002-7777-216X>

B. Scott Gaudi  <https://orcid.org/0000-0003-0395-9869>

François Lanusse  <https://orcid.org/0000-0001-7956-0542>

Casey Lam  <https://orcid.org/0000-0002-6406-1924>

Jessica R. Lu  <https://orcid.org/0000-0001-9611-0009>

References

- Ba, J. L., Kiros, J. R., & Hinton, G. E. 2016, arXiv:1607.06450
- Cho, K., van Merriënboer, B., Gulcehre, C., et al. 2014, in Proc. 2014 Conf. on Empirical Methods in Natural Language Processing (EMNLP) (Doha: Association for Computational Linguistics), 1724
- Choi, J.-Y., Shin, I.-G., Han, C., et al. 2012, *ApJ*, **756**, 48
- Cranmer, K., Brehmer, J., & Louppe, G. 2020, *PNAS*, **117**, 30055
- Dinh, L., Sohl-Dickstein, J., & Bengio, S. 2017, arXiv:1605.08803
- Dominik, M. 1999, *A&A*, **349**, 108
- Foreman-Mackey, D. 2016, *JOSS*, **1**, 24
- Gaudi, B. S. 2010, in Exoplanets, ed. S. Seager (Tucson, AZ: Univ. Arizona Press), 79
- Gaudi, B. S. 2012, *ARA&A*, **50**, 411
- Gaudi, B. S., & Gould, A. 1997, *ApJ*, **486**, 85
- Germain, M., Gregor, K., Murray, I., & Larochelle, H. 2015, in Proc. 32nd International Conf. on Machine Learning, ed. F. Bach & D. Blei (PMLR), 881, <http://proceedings.mlr.press/v37/germain15.html>
- Godines, D., Bachelet, E., Narayan, G., & Street, R. A. 2019, *A&C*, **28**, 100298
- Gould, A., & Loeb, A. 1992, *ApJ*, **396**, 104
- Griest, K., & Safizadeh, N. 1998, *ApJ*, **500**, 37
- Han, C. 2008, *ApJL*, **691**, L9
- He, K., Zhang, X., Ren, S., & Sun, J. 2016, in 2016 IEEE Conf. on Computer Vision and Pattern Recognition (CVPR), ed. L. O’Conner (Piscataway, NJ: IEEE), 770
- Herrera-Martín, A., Albrow, M. D., Udalski, A., et al. 2020, *AJ*, **159**, 256
- Hochreiter, S., & Schmidhuber, J. 1997, *Neural Computation*, **9**, 1735
- Hunter, J. D. 2007, *CSE*, **9**, 90
- Jones, E., Oliphant, T., Peterson, P., et al. 2001, SciPy: Open Source Scientific Tools for Python, <http://www.scipy.org/>
- Khakpash, S., Penny, M., & Pepper, J. 2019, *AJ*, **158**, 9
- Kingma, D. P., & Ba, J. 2015, arXiv:1412.6980
- Kluyver, T., Ragan-Kelley, B., Pérez, F., et al. 2016, in Positioning and Power in Academic Publishing: Players, Agents and Agendas, ed. F. Loizides & B. Schmidt (Amsterdam: IOS Press), 87
- LeCun, Y., Bengio, Y., & Hinton, G. 2015, *Natur*, **521**, 436
- Loshchilov, I., & Hutter, F. 2017, arXiv:1608.03983
- Mróz, P. 2020, *AcA*, **70**, 169
- Mróz, P., Udalski, A., Skowron, J., et al. 2017, *Natur*, **548**, 183
- Naul, B., Bloom, J. S., Pérez, F., & van der Walt, S. 2018, *NatAs*, **2**, 151
- Nelder, J. A., & Mead, R. 1965, *CompJ*, **7**, 308
- Oord, A. V. D., Dieleman, S., Zen, H., et al. 2016, arXiv:1609.03499
- Papamakarios, G., Pavlakou, T., & Murray, I. 2017, in Proc. Advances in Neural Information Processing Systems 30, ed. I. Guyon et al. (Red Hook, NY: Curran Associates, Inc.), 2338, <http://papers.nips.cc/paper/6828-masked-autoregressive-flow-for-density-estimation.pdf>
- Papamakarios, G., Sterratt, D. C., & Murray, I. 2019, arXiv:1805.07226
- Paszke, A., Gross, S., Chintala, S., et al. 2017, in 31st Conf. Neural Information Processing Systems (NIPS 2017), ed. V. von Luxburg, I. Guyon, & S. Bengio (Red Hook, NY: Curran Associates Inc.), <https://openreview.net/pdf?id=BJJsmfCZ>
- Penny, M. T., Scott Gaudi, B., Kerins, E., et al. 2019, *ApJS*, **241**, 3
- Poleski, R., & Yee, J. 2018, `MulensModel`: Microlensing light curves modeling. Version 1.17.4, Astrophysics Source Code Library, ascl:1803.006
- Poleski, R., & Yee, J. C. 2019, *A&C*, **26**, 35
- Spergel, D., Gehrels, N., Baltay, C., et al. 2015, arXiv:1503.03757
- Thomas, O., Dutta, R., Corander, J., Kaski, S., & Gutmann, M. U. 2020, arXiv:1611.10242
- van der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, *CSE*, **13**, 22
- Vermaak, P. 2003, *MNRAS*, **344**, 651
- Woźniak, P., & Paczyński, B. 1997, *ApJ*, **487**, 55
- Wyzykowski, L., Rynkiewicz, A. E., Skowron, J., et al. 2015, *ApJS*, **216**, 12
- Zhang, K., & Bloom, J. S. 2020, arXiv:2011.01243
- Zhang, K., Bloom, J. S., Gaudi, B. S., et al. 2020, *MNRAS*, in press (doi:10.1093/mnras/stab1248)