

Estimating orbital parameters of direct imaging exoplanet using neural networks

Bo Liang,^{1,2,*} Chang Liu,^{3,*} Hanlin Song^{①,1,4,*†} Tianyu Zhao,^{1,‡} Yuxiang Xu,^{1,5} Zihao Xiao,^{1,3} Manjia Liang,¹ Minghui Du,¹ Wei-Liang Qian^{②,6}, Li-e Qiang^{③,1} Mingming Sun^{④,7} Peng Xu,^{1,8,9,§} and Ziren Luo^{1,8}

¹Center for Gravitational Wave Experiment, National Microgravity Laboratory, *Institute of Mechanics*, Chinese Academy of Sciences, Beijing 100190, China

²Taiji Laboratory for Gravitational Wave Universe (Beijing/Hangzhou), *University of Chinese Academy of Sciences (UCAS)*, Beijing 100049, China

³National Space Science Center, Chinese Academy of Sciences, Beijing 100190, China

⁴School of Physics, *Peking University*, Beijing 100871, China

⁵Shanghai Institute of Optics and Fine Mechanics, Chinese Academy of Sciences, Shanghai 201800, China

⁶Escola de Engenharia de Lorena, *Universidade de São Paulo*, Lorena 12602-810, SP, Brazil

⁷AGI Laboratory, *Beijing Institute of Mathematical Sciences and Applications*, Beijing, China

⁸Key Laboratory of Gravitational Wave Precision Measurement of Zhejiang Province, Hangzhou Institute for Advanced Study, UCAS, Hangzhou 310024, China

⁹Lanzhou Center of Theoretical Physics, *Lanzhou University*, Lanzhou 730000, China



(Received 3 December 2025; accepted 27 May 2026; published 25 June 2026)

We propose a flow-matching Markov chain Monte Carlo (FM-MCMC) algorithm for estimating the orbital parameters of exoplanetary systems, especially for those in which only one exoplanet is involved. Compared to traditional methods that rely on random sampling within the Bayesian framework, our approach first leverages flow-matching posterior estimation to efficiently constrain the prior range of physical parameters and then employs MCMC to accurately infer the posterior distribution. For example, in the orbital parameter inference of β Pictoris b, our model achieved a substantial speedup while maintaining comparable accuracy—running 77.8 times faster than parallel tempered MCMC and 365.4 times faster than nested sampling. Moreover, our FM-MCMC method also attained the highest average log-likelihood among all approaches, demonstrating its superior sampling efficiency and accuracy. This highlights the scalability and efficiency of our approach, making it well suited for processing the massive datasets expected from future exoplanet surveys. Beyond astrophysics, our methodology establishes a versatile paradigm for synergizing deep generative models with traditional sampling, which can be adopted to tackle complex inference problems in other fields, such as cosmology, biomedical imaging, and particle physics.

DOI: [10.1103/ykfh-cdzk](https://doi.org/10.1103/ykfh-cdzk)

I. INTRODUCTION

Since the first confirmed detection of an exoplanetary system in 1992 [1] and an exoplanet in 1995 [2], more than 4615 exoplanetary systems and 5989 exoplanets have been confirmed as of September 2025 [3]. These systems exhibit a wide range of characteristics, including variations in the number and composition of exoplanets, the dynamics of their orbital motions, and the types of their host stars. Notably, the majority of these systems contain only one confirmed planet, a pattern that is partly attributable to observational biases. For a recent

review, see, e.g., Ref. [4]. High-precision exoplanet detectors are essential for investigating the mechanisms underlying these characteristics, such as testing theories of planet formation and migration, and understanding the overall evolution of planetary systems [5,6]. Currently, several observatories remain in operation and are continuously expanding our exoplanet census. Operating projects like *Transiting Exoplanet Survey Satellite* [7] and *CHAracterising ExOPlanets Satellite* [8] conduct high-precision photometric surveys of nearby bright stars, enabling the detection and characterization of transiting exoplanets. In parallel, *Gaia* [9] ultraprecise astrometric data reveal subtle stellar motions caused by orbiting planets. Additionally, instruments like Spectro-Polarimetric High-contrast Exoplanet REsearch (SPHERE) [10] on the Very Large Telescope directly image exoplanets, offering valuable constraints on their orbital architectures and atmospheric properties. Upcoming missions such as PLANetary Transits and Oscillations of stars [11], Atmospheric Remote-sensing Infrared Exoplanet Large-survey [12,13], SPHERE + [14], and Nancy Grace Roman Space Telescope [15] are expected to significantly advance our understanding of exoplanet populations. By integrating large-scale photometric

*These authors contributed equally to this work.

†Contact author: hanlin@stu.pku.edu.cn

‡Contact author: zhaotianyu@imech.ac.cn

§Contact author: xupeng@imech.ac.cn




FIG. 1. Flow-matching MCMC framework. In the training section, there is first a prior distribution $p(\theta)$, from which the parameter θ is extracted, and the data are generated by a model x . These generated data are fed into the training process along with the parameters. The inference section is based on the trained model, inputs the real observation data, and uses the trained model to infer the initial proposals. The initial proposal is then provided to PTMCMC for likelihood calculations.

monitoring, atmospheric spectroscopy, and precise astrometric measurements, these missions are poised to deliver unprecedented insights into the formation, evolution, and diversity of planetary systems. In recent years, a variety of exoplanet discovery techniques have been developed and extensively employed [16], including transit photometry, radial velocity, microlensing, and direct imaging. These efforts have led to the detection of a large number of exoplanetary systems, significantly advancing our understanding of the statistical properties of planetary populations and providing critical constraints on models of planet formation and migration. In this work, we focus on exoplanetary systems detected using direct imaging method. Although the number of such detections is smaller than that obtained with the other three techniques [17,18], direct imaging provides unique advantages for determining the demographics of exoplanets around nearby young stars, characterizing their atmospheric properties, and exploring their circumstellar disks [14]. Through synergy with other facilities, such as *Gaia*, GRAVITY [19], and James Webb Space Telescope [20], more exoplanets will be directly imaged. This growing catalog of directly imaged exoplanets will yield essential observational data, which are vital for testing and validating theoretical models of planet formation and evolution, thereby elucidating the fun-

damental physical processes that determine planetary system architectures.

With the growing accumulation of observational data, a key challenge lies in the rapid and precise estimation of the physical parameters of exoplanetary systems, which are crucial for predicting exoplanet locations, calculating transit probabilities, and assessing the climates and habitability of Earth-like worlds in future missions. In the case of directly imaged exoplanets, several approaches have been proposed for estimating the physical parameters [5], such as linearized least-squares adjustment [21], the Orbits for the Impatient algorithm [22], and the parallel tempered Markov chain Monte Carlo (PTMCMC) algorithm [23,24]. Tools such as *orbitize!* implement nested sampling [25] and have been demonstrated in studies that combine astrometric and radial velocity data to constrain exoplanet orbits [5,6]. However, these traditional methods [21–24] suffer from limitations: Due to high parameter dimensionality and the potential multimodality of posterior distributions, convergence can be prohibitively slow without an optimal proposal distribution or favorable random initialization. In contrast, while machine-learning-based methods [26–28] offer a significant advantage in speed for rapid orbital parameter inference, the posterior distributions they generate typically exhibit larger errors than




FIG. 2. FM-MCMC vs PTMCMC vs nested sampling posterior comparison for β Pictoris b orbital elements. Posterior distributions of β Pictoris b's orbital parameters inferred by FM-MCMC (blue), PTMCMC (orange), and nested sampling (red) show statistical consistency, with nearly identical 1σ and 2σ credible regions.

traditional MCMC methods and often fail to capture multimodal features accurately. Consequently, achieving rapid inference of exoplanet orbital parameters while maintaining high precision remains a pressing challenge.

In this work, we present an approach, flow-matching-enhanced MCMC (FM-MCMC), that leverages continuous normalizing flows (CNFs) [29] trained via a flow-matching objective [30–32] to accelerate MCMC sampling. Our method achieves posterior estimates that closely approximate those obtained by PTMCMC and nested sampling, while outperforming neural posterior estimation [33,34] in preci-

sion. FM-MCMC combines the robustness and interpretability of traditional MCMC sampling with the computational speed of normalizing flow, yielding higher posterior accuracy than stand-alone simulation-based inference methods [35,36].

II. METHODOLOGY

In this section, we introduce the FM-MCMC algorithm to accelerate posterior sampling. The overall workflow of the training and inference procedure is shown in Fig. 1. Based on the CNFs and flow-matching algorithm, we first train a neural

network for initial inference. Then, along with the PTMCMC, we obtain the accurate posterior distributions. Unlike traditional MCMC methods—which often suffer from slow convergence in high-dimensional and multimodal parameter spaces due to inefficient proposal mechanisms and sensitivity to initial conditions—FM-MCMC leverages a learned flow-based proposal that drastically reduces the number of likelihood evaluations and shortens burn-in phases.

A. Continuous normalizing flows

Generative models are designed to learn invertible transformations that map high-dimensional data distributions to simpler reference distributions. Notably, normalizing flows achieve exact density estimation by applying the chain rule to compositions of invertible functions. However, their discrete layerwise architecture limits their capacity to model highly complex data distributions. To overcome this limitation, CNFs were introduced. CNFs construct a time-dependent probability density path $p_t(\theta)$ governed by a neural-network-parametrized vector field $v_t(x; \theta)$. The evolution of samples in CNFs is described by an ordinary differential equation (ODE):

$$\frac{d}{dt}\phi_{t,x}(\theta) = v_{t,x}(\phi_{t,x}(\theta)), \quad v_{t=0,x}(\theta) = \theta_0, \quad v_{t=1,x}(\theta) = \theta_1. \quad (1)$$

Traditional methods optimize the model via maximum likelihood estimation by minimizing the negative log-likelihood. However, backpropagation through an ODE solver incurs a computational cost of $\mathcal{O}(d^3)$, where d denotes dimensionality, hindering scalability to high-dimensional data. FM offers a solution by introducing a more efficient optimization objective. In contrast to conventional CNFs, FM decouples the probability density path from the underlying dynamics, substantially reducing computational overhead. It explicitly constructs a target probability path $p_t(\theta)$ and its associated velocity field $u_t(\theta)$, and trains a parametric velocity field $v_{t,x}(\theta)$ to approximate $u_t(\theta)$. Dax *et al.* applied FM to simulation-based inference, proposing a highly efficient training objective that alleviates the computational demands of traditional methods. The method begins by defining a Gaussian probability path conditioned on a target parameter θ_1 :

$$p_t(\theta | \theta_1) = \mathcal{N}(\theta | \theta_1 t, (1-t)^2 I_d), \quad (2)$$

which ensures the distribution starts as a standard normal $\mathcal{N}(0, I_d)$ at $t = 0$ and converges to a Gaussian centered at θ_1 when $t = 1$. The corresponding conditional velocity field that generates this path is derived as

$$u_t(\theta | \theta_1) = \theta_1 - \theta, \quad (3)$$

enabling linear transport of samples from the base to the target distribution over time. To approximate the posterior, the framework marginalizes over target parameters $\theta_1 \sim p(\theta)$ and conditions on observed data $x \sim p(x | \theta_1)$, yielding an aggregate probability path $p_t(\theta)$ and a composite velocity field $u_t(\theta)$ such that $p_1(\theta) = q(\theta | x)$. The training objective minimizes the mean-squared error between the learned and true velocity fields:

$$\begin{aligned} \mathcal{L}_{\text{FMPE}} = & \mathbb{E}_{t \sim \mathcal{U}(0,1)} \mathbb{E}_{\theta_1 \sim p(\theta)} \mathbb{E}_{x \sim p(x|\theta_1)} \\ & \times \mathbb{E}_{\theta_t \sim p_t(\theta_t|\theta_1)} \|v_{t,x}(\theta_t) - u_t(\theta_t | \theta_1)\|^2. \end{aligned} \quad (4)$$




FIG. 3. P-P plot validating the unbiasedness of the CNF model. The near-perfect alignment of all CDF curves with the theoretical diagonal indicates statistical consistency between the CNF posterior estimates and the true distribution.

This regression-based approach avoids expensive backpropagation through ODE solvers, greatly improving scalability in high-dimensional settings. Once trained, posterior samples are generated by solving the ODE numerically, effectively transforming base distribution samples into draws from the approximate posterior.

B. FM-MCMC sampling

The FM-MCMC algorithm first employs CNFs trained through flow matching to enable rapid obtaining of the initial proposals to efficiently generate initial proposals tailored to specific scientific problems. These proposals are then refined using PTMCMC to obtain accurate posterior distributions. For example, when applying the FM-MCMC algorithm to estimate the orbital parameters of imaged exoplanet, i.e., β Pictoris b, a two-step inference is performed. Step 1: Obtain the initial proposal for the orbital parameters of β Pictoris b using a pretrained neural network. Step 2: Obtain the accurate posterior distribution with the injected initial proposal using the PTMCMC algorithm. The CNFs' sample production dynamically adapts to PTMCMC's chain configurations and temperature ladder parameters. This integration delivers two key enhancements: (1) using CNFs' posterior samples as the initialization proposal distribution and (2) burn-in phases are reduced to 1% of the original requirement. Final posterior distributions are derived from thermally equilibrated PTMCMC chains.

C. Data and model

Since our ultimate goal is to estimate the orbital parameters of β Pictoris b with the trained neural network, it is

TABLE I. Comparison of efficiency and effectiveness of different posterior sampling approaches. Bold values indicate the best performance in each column.

Method	Time (s) $\uparrow$	Speedup	Log $\mathcal{L}$ (mean) $\uparrow$
FM-MCMC (ours)	348.9	1	-133.0
PTMCMC	27133.9	77.8	-285.4
Nested sampling	127478.8	365.4	-234.3

essential to construct a training dataset consisting of simulated observations that closely resemble systems analogous to β Pictoris b. In this work, we generate a dataset comprising 16 million simulated orbital configurations. The relative astrometry data of exoplanets are often expressed as the star-planet separation and position angle, or equivalently as the relative right ascension (ΔRA) and relative declination (ΔDec) [6,37]. These observed orbital data can be parametrized with eight quantities: the six standard Keplerian elements plus parallax π and total system mass M_{tot} . The specific priors and noise models used for training are detailed in the Supplemental Material [38].

III. RESULTS

With the flow-matching-based neural network, the orbital parameters of β Pictoris b can be inferred through the two-step procedure described above. Our FM-MCMC results demonstrate two major improvements over previous parameter estimation methods that rely solely on either machine learning or traditional sampling algorithms.

A. Efficiency improvement

To rigorously evaluate the effectiveness of FM-MCMC, we performed posterior estimation for β Pictoris b using both PTMCMC and nested sampling. The PTMCMC sampler was configured with 20 temperature tiers and 1000 parallel chains. After burn-in, we retained the final 20 000 samples for posterior analysis. The PTMCMC computation completed in 452 min, while the nested sampling algorithm required 2124 min to achieve convergence. As shown in Table I, FM-MCMC demonstrates high efficiency by integrating machine learning with traditional sampling. Under identical configurations (20 temperature tiers, 1000 parallel chains), FM-MCMC completes sampling in 348.9 s (≈ 5.8 min), outperforming PTMCMC (27 133.9 s/7.5 h) and nested sampling (127 478.8 s/35.4 h) by 77.8 $\times$ and 365.4 $\times$, respectively. Moreover, likelihood comparisons show that FM-MCMC achieves a mean negative log-likelihood of -133.0, which is 53.4% better than that of PTMCMC (-285.4) and 43.2% better than that of nested sampling (-234.3). Critically, its maximum likelihood (-129.6) equals that of PTMCMC (-129.6) and deviates by only 0.08% from that of nested sampling (-129.5), demonstrating that FM-MCMC achieves enhanced efficiency without compromising posterior accuracy. The offline training of the flow-matching posterior estimation network was completed in about 22 min on an NVIDIA RTX 4090 GPU using early stopping.

B. Accuracy

FM-MCMC demonstrates statistical robustness in parameter inference. As shown in Fig. 3, both the 2D contours and 1D marginal distributions of all FM-MCMC parameters are consistent with results obtained from PTMCMC and nested sampling, confirming the fidelity of the inferred distributions. Detailed posterior estimates and comparisons with purely neural-based methods (neural posterior estimation (NPE)) are provided in the Supplemental Material [38].

C. Robustness verification

To statistically validate the calibration of the CNFs' architecture across the full parameter space, Fig. 2 presents a probability-probability (P-P) plot demonstrating the model's unbiasedness, where all parameterwise cumulative distribution functions (CDFs) align closely with the theoretical diagonal. This consistency highlights the unbiased nature of the trained model, confirming its reliability and robustness across diverse parameters.

D. Scientific implications

The capability to reduce inference time from hours to minutes enables rapid updates of orbital posteriors for dynamically active systems. For targets like β Pictoris b, fast orbital inference is useful for rapid follow-up observations. Newly obtained astrometric data can be incorporated quickly to update the predicted separation, position angle, and observability of a planet during future observing windows. Such rapid updates will be valuable for prioritizing follow-up spectroscopy in future large surveys of directly imaged planets. Furthermore, the rapid posterior sampling facilitates on-the-fly N -body simulations to probe disk-planet interactions [39,40] and provides informed priors to accelerate the search for inner companions [41]. More broadly, this efficiency helps remove a computational bottleneck in characterizing the large exoplanet datasets expected from future surveys and astrometric missions such as Roman and *Gaia*.

IV. CONCLUSION

We have presented FM-MCMC, a flow-matching-enhanced MCMC sampler for astrometric orbit fitting, and demonstrated its application to the well-studied β Pictoris b system. FM-MCMC produces posterior distributions that are statistically consistent with those obtained via fully converged PTMCMC and nested sampling, while achieving a roughly 100 $\times$ reduction in wall-clock run-time. Importantly, this substantial speedup is achieved without sacrificing the accuracy of the posterior estimates. By integrating learned transport maps into a conventional MCMC sampler, FM-MCMC unites the statistical rigor of Bayesian inference with the high computational efficiency characteristic of modern generative models. While our study focused on β Pictoris b using priors consistent with previous works, FM-MCMC shows promise for application to broader prior ranges, multiplanet systems, and entirely new scientific problems.

In the present implementation, the network is trained for the specific observational setup and prior ranges adopted for the β Pictoris b analysis. Therefore, applying the current model to another directly imaged planet with different observing epochs, measurement uncertainties, or prior ranges requires a new simulated training set and retraining. This retraining step is computationally inexpensive; for the β Pictoris b analysis, it takes only about 22 min, allowing the present implementation to be efficiently adapted to other datasets. We also note that this source-specific training requirement is not an intrinsic limitation of FM-MCMC. A future generalized implementation could include observing epochs, measurement uncertainties, missing-data masks, and system-level quantities in the conditioning vector, allowing a single trained model to be applied to multiple sources and heterogeneous observation sets.

ACKNOWLEDGMENTS

We thank the anonymous referee for profound comments. This study was supported by the National

Key Research and Development Program of China (Grants No. 2021YFC2201901, No. 2021YFC2203004, No. 2020YFC2200100, and No. 2021YFC2201903) and International Partnership Program of the Chinese Academy of Sciences (Grant No. 025GJHZ2023106GC). We also gratefully acknowledge the financial support from Brazilian agencies Fundação de Amparo à Pesquisa do Estado de São Paulo (FAPESP), Fundação de Amparo à Pesquisa do Estado do Rio Grande do Sul (FAPERGS), Fundação de Amparo à Pesquisa do Estado do Rio de Janeiro (FAPERJ), Conselho Nacional de Desenvolvimento Científico e Tecnológico (CNPq), and Coordenação de Aperfeiçoamento de Pessoal de Nível Superior (CAPES). This work was also supported by High-performance Computing Platform of Peking University.

DATA AVAILABILITY

The code and trained neural-network model used in this work are publicly available in the FM-MCMC GitHub repository [42].

-
- [1] A. Wolszczan and D. A. Frail, A planetary system around the millisecond pulsar PSR1257+12, *Nature (London)* **355**, 145 (1992).
 - [2] M. Mayor and D. Queloz, A Jupiter-mass companion to a solar-type star, *Nature (London)* **378**, 355 (1995).
 - [3] NASA Exoplanet Archive, NASA Exoplanet Archive by the data 2025.07.09, 2025, <https://science.nasa.gov/exoplanets/>, accessed 5 September 2025.
 - [4] A. R. Howe, J. C. Becker, C. C. Stark, and F. C. Adams, Architecture classification for extrasolar planetary systems, *Astron. J.* **169**, 149 (2025).
 - [5] T. D. Brandt, T. J. Dupuy, Y. Li, G. M. Brandt, Y. Zeng, D. Michalik, D. C. Bardalez Gagliuffi, and V. Raposo-Pulido, orvara: An efficient code to fit orbits using radial velocity, absolute, and/or relative astrometry, *Astron. J.* **162**, 186 (2021).
 - [6] S. Blunt, J. J. Wang, I. Angelo, H. Ngo, D. Cody, R. J. De Rosa, J. R. Graham, L. Hirsch, V. Nagpal, E. L. Nielsen, L. Pearce, M. Rice, and R. Tejada, orbitize!: A comprehensive orbit-fitting software package for the high-contrast imaging community, *Astron. J.* **159**, 89 (2020).
 - [7] J. N. Winn, The transiting exoplanet survey satellite, [arXiv:2410.12905](https://arxiv.org/abs/2410.12905).
 - [8] A. Fortier, A. E. Simon, C. Broeg, G. Olofsson, A. Deline, T. G. Wilson, P. F. Maxted, A. Brandeker, A. C. Cameron, M. Beck, *et al.*, Cheops in-flight performance—A comprehensive look at the first 3.5 yr of operations, *Astron. Astrophys.* **687**, A302 (2024).
 - [9] E. Pancino, *Gaia*: The galaxy in six (and more) dimensions, *Adv. Space Res.* **65**, 1 (2020).
 - [10] J.-L. Beuzit, A. Vigan, D. Mouillet, K. Dohlen, R. Gratton, A. Boccaletti, J.-F. Sauvage, H. M. Schmid, M. Langlois, C. Petit, *et al.*, Sphere: The exoplanet imager for the very large telescope, *Astron. Astrophys.* **631**, A155 (2019).
 - [11] H. Rauer, C. Aerts, J. Cabrera, M. Deleuil, A. Erikson, L. Gizon, M. Goupil, A. Heras, J. Lorenzo-Alvarez, F. Marliani, C. Martin-Garcia, J. M. Mas-Hesse, L. O'Rourke, H. Osborn, I. Pagano, and G. Piotto, The Plato mission, *Experimental Astron.* **59**, 26 (2025).
 - [12] Q. Changeat and K. H. Yip, ESA-Ariel data challenge NeurIPS 2022: Introduction to exo-atmospheric studies and presentation of the Atmospheric Big Challenge (ABC) database, *RAS Techniques and Instruments* **2**, 45 (2023).
 - [13] T. Zingales, G. Tinetti, I. Pillitteri, J. Leconte, G. Micela, and S. Sarkar, The Ariel mission reference sample, *Exp. Astron.* **46**, 67 (2018).
 - [14] A. Boccaletti, G. Chauvin, D. Mouillet, O. Absil, F. Allard, S. Antonucci, J.-C. Augereau, P. Barge, A. Baruffolo, J.-L. Baudino, *et al.*, Sphere+: Imaging young Jupiters down to the snowline, [arXiv:2003.05714](https://arxiv.org/abs/2003.05714).
 - [15] V. P. Bailey, E. Bendek, B. Monacelli, C. Baker, G. Bedrosian, E. Cady, E. S. Douglas, T. Groff, S. R. Hildebrandt, N. J. Kasdin, *et al.*, Nancy Grace Roman Space Telescope coronagraph instrument overview and status, *Techniques and Instrumentation for Detection of Exoplanets XI* **12680**, 283 (2023).
 - [16] A.-L. Maire, L. Delrez, F. J. Pozuelos, J. Becker, N. Espinoza, J. Lillo-Box, A. Revol, O. Absil, E. Agol, J. M. Almenara, *et al.*, Workshop summary: Exoplanet orbits and dynamics, *Publ. Astron. Soc. Pac.* **135**, 106001 (2023).
 - [17] D. A. Fischer, A. W. Howard, G. P. Laughlin, B. Macintosh, S. Mahadevan, J. Sahlmann, and J. C. Yee, Exoplanet detection techniques, in *Protostars and Planets VI*, edited by H. Beuther, R. S. Klessen, C. P. Dullemond, and T. Henning (University of Arizona Press, Tucson, AZ, 2014), pp. 715–737.
 - [18] H. J. Deeg and J. A. Belmonte, *Handbook of Exoplanets* (2018).
 - [19] R. Abuter, M. Accardo, A. Amorim, N. Anugu, G. Avila, N. Azouaoui, M. Benisty, J.-P. Berger, N. Blind, H. Bonnet, *et al.*, First light for gravity: Phase referencing optical interferometry for the very large telescope interferometer, *Astron. Astrophys.* **602**, A94 (2017).

- [20] P. A. Lightsey, C. Atkinson, M. Clampin, and L. D. Feinberg, James Webb Space Telescope: Large deployable cryogenic telescope in space, *Opt. Eng.* **51**, 011003 (2012).
- [21] T. Forveille, J. L. Beuzit, X. Delfosse, D. Segransan, F. Beck, M. Mayor, C. Perrier, A. Tokovini, and S. Udry, Accurate masses of very low mass stars: I Gl 570bc ($0.6+0.4$ Msol), *Astron. astrophys.* **351**, 619 (1999).
- [22] S. Blunt, E. L. Nielsen, R. J. De Rosa, Q. M. Konopacky, D. Ryan, J. J. Wang, L. Pueyo, J. Rameau, C. Marois, F. Marchis, B. Macintosh, J. R. Graham, G. Duchêne, and A. C. Schneider, Orbits for the impatient: A Bayesian rejection-sampling method for quickly fitting the orbits of long-period exoplanets, *Astron. J.* **153**, 229 (2017).
- [23] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, emcee: The MCMC hammer, *Publ. Astron. Soc. Pac.* **125**, 306 (2013).
- [24] W. D. Vouden, W. M. Farr, and I. Mandel, Dynamic temperature selection for parallel tempering in Markov chain Monte Carlo simulations, *Mon. Not. R. Astron. Soc.* **455**, 1919 (2016).
- [25] J. S. Speagle, dynesty: A dynamic nested sampling package for estimating Bayesian posteriors and evidences, *Mon. Not. R. Astron. Soc.* **493**, 3132 (2020).
- [26] H. Sun, K. L. Bouman, P. Tiede, J. J. Wang, S. Blunt, and D. Mawet, Alpha-deep probabilistic inference (α -dpi): Efficient uncertainty quantification from exoplanet astrometry to black hole feature extraction, *Astrophys. J.* **932**, 99 (2022).
- [27] M. Ruth, Exoplanet orbital characterization using simulation-based inference, Ph.D. thesis, University of Liège, Liège, Belgium, 2024.
- [28] C. Chen, L. Kong, G. Li, and M. Tao, DeepTTV: Deep learning prediction of hidden exoplanet from transit timing variations, [arXiv:2409.04557](https://arxiv.org/abs/2409.04557).
- [29] R. T. Q. Chen, B. Amos, and M. Nickel, Learning neural event functions for ordinary differential equations, in *International Conference on Learning Representations* (ICLR, Virtual Event, Austria, 2021).
- [30] X. Liu, C. Gong, and Q. Liu, Flow straight and fast: Learning to generate and transfer data with rectified flow, in *The Eleventh International Conference on Learning Representations* (2022).
- [31] J. Wildberger, M. Dax, S. Buchholz, S. R. Green, J. H. Macke, and B. Schölkopf, Flow matching for scalable simulation-based inference, *Adv. Neural Inf. Process. Syst.* **36**, 16837 (2023).
- [32] Y. Lipman, R. T. Q. Chen, H. Ben-Hamu, M. Nickel, and M. Le, Flow matching for generative modeling, in *The Eleventh International Conference on Learning Representations* (ICLR, Kigali, Rwanda, 2023).
- [33] G. Papamakarios and I. Murray, Fast ε -free inference of simulation models with Bayesian conditional density estimation, in *Advances in Neural Information Processing Systems* (Curran Associates, Red Hook, NY, 2016), Vol. 29, pp. 1028–1036.
- [34] J.-M. Lueckmann, P. J. Goncalves, G. Bassetto, K. Öcal, M. Nonnenmacher, and J. H. Macke, Flexible statistical inference for mechanistic models of neural dynamics, [arXiv:1711.01861](https://arxiv.org/abs/1711.01861).
- [35] M. Gabrié, G. M. Rotskoff, and E. Vanden-Eijnden, Adaptive Monte Carlo augmented with normalizing flows, *Proc. Natl. Acad. Sci. USA* **119**, e2109420119 (2022).
- [36] K. W. k. Wong, M. Gabrié, and D. Foreman-Mackey, flowMC: Normalizing flow enhanced sampling package for probabilistic inference in JAX, *J. Open Source Softw.* **8**, 5021 (2023).
- [37] E. L. Nielsen, R. J. De Rosa, J. J. Wang, J. Sahlmann, P. Kalas, G. Duchêne, J. Rameau, M. S. Marley, D. Saumon, B. Macintosh, *et al.*, The Gemini Planet Imager Exoplanet Survey: Dynamical mass of the exoplanet β Pictoris b from combined direct imaging and astrometry, *Astron. J.* **159**, 71 (2020).
- [38] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/ykfh-cdzk> for additional details on the priors, noise models, training procedure, posterior estimates, and comparisons with neural posterior estimation methods.
- [39] R. G. Belleman, J. Bédorf, and S. F. P. Zwart, High performance direct gravitational N -body simulations on graphics processing units II: An implementation in CUDA, *New Astron.* **13**, 103 (2008).
- [40] A. Lacquement, H. Beust, V. Faramaz-Gorka, and G. Duchêne, Dynamics of the β Pictoris planetary system and possibility of an additional planet, *Astron. Astrophys.* **694**, A236 (2025).
- [41] S. Lacour, J. Wang, L. Rodet, M. Nowak, J. Shangguan, H. Beust, A.-M. Lagrange, R. Abuter, A. Amorim, R. Asensio-Torres, *et al.*, The mass of β Pictoris c from β Pictoris b orbital motion, *Astron. Astrophys.* **654**, L2 (2021).
- [42] B. Liang, C. Liu, H. Song, T. Zhao, Y. Xu, Z. Xiao, M. Liang, M. Du, W.-L. Qian, L.-e Qiang, *et al.*, FM-MCMC: Estimating orbital parameters of direct imaging exoplanets using neural networks, GitHub repository (2026), <https://github.com/aspirinliang/Estimating-Orbital-Parameters-of-Direct-Imaging-Exoplanet-Using-Neural-Network>, accessed 26 May 2026.