

FLOW MATCHING FOR POSTERIOR INFERENCE WITH SIMULATOR FEEDBACK

Anonymous authors

Paper under double-blind review

ABSTRACT

Flow-based models have shown great success in generative modeling, making them a promising candidate for solving inverse problems in physical sciences and allow for sampling and likelihood evaluation with much lower inference times than traditional methods. We propose to pretrain a neural network via flow matching and include control signals based on a simulator as an additional input for finetuning via a lightweight control network. Control signals can include gradients and a problem-specific cost function if the simulator is differentiable, or they can be fully learned from the simulator output. We motivate our design choices on several benchmark problems for simulation-based inference and evaluate flow matching with simulator feedback against classical MCMC methods for modeling strong gravitational lens systems, a challenging inverse problem in astronomy. We demonstrate that including simulator feedback improves the accuracy of reconstructed samples by 53%, making it competitive with traditional techniques while being up to 67x faster for inference. Upon acceptance, we will make our code publicly available.

1 INTRODUCTION

Acquiring posterior distributions given measurement data is of paramount scientific interest (Cranmer et al., 2020), with real-world applications ranging from particle physics (Baydin et al., 2019), over the inference of gravitational waves (Dax et al., 2021) to predictions of dynamical systems such as weather forecasting (Gneiting & Raftery, 2005). In Bayesian modeling, given an observation x_o and model parameters θ , we are interested in the posterior $p(\theta|x_o)$. Traditional likelihood-based methods can be expensive for high-dimensional data, when likelihood evaluations are costly or intractable and priors are difficult to represent mathematically. Simulation-based inference (Cranmer et al., 2020, SBI) addresses these challenges by including a learning-based component in the statistical inference process. In this paper, we focus on neural posterior estimation (NPE), which represents the posterior as a parametric function $q(\theta|x_o)$, which is a learnable conditional density estimator that can be trained purely by simulations $x \sim p(x|\theta)$ alone. By investing an upfront cost for training the density estimator, we can sample and compute likelihoods from $q(\theta|x_o)$ much faster than other methods, thereby amortizing the training cost over many observations. Traditionally, normalizing flows (Rezende & Mohamed, 2015; Dinh et al., 2017; Papamakarios et al., 2019) have been a popular class of density estimators used in many areas of science. To compute likelihoods and for sampling, normalizing flows transform a noise distribution to the target distribution via a bijective mapping. By conditioning the normalizing flow networks on the observation x_o obtained from the simulator, they can be trained as the conditional density estimator $q(\theta|x_o)$ for the posterior. The success of diffusion models (Ho et al., 2020; Dhariwal & Nichol, 2021; Song et al., 2021) has demonstrated that the mapping between sampling and posterior distribution can be specified by a corruption process that transforms any data distribution to a normal Gaussian. Diffusion models and normalizing flows can be linked via the probability flow ODE (Song et al., 2021), which has also influenced a class of flow-based models that can be trained via flow matching (Lipman et al., 2023) on more general mappings between sampling and target distribution than considered by diffusion models. The resulting continuous-time models outperform discrete, classical normalizing flows in many areas, and training larger models is much more scalable (Wildberger et al., 2023).

Despite the widespread success of flow-based models for generative modeling and density estimation, there is no direct feedback from the simulator between the model, the observation x_o and the sample

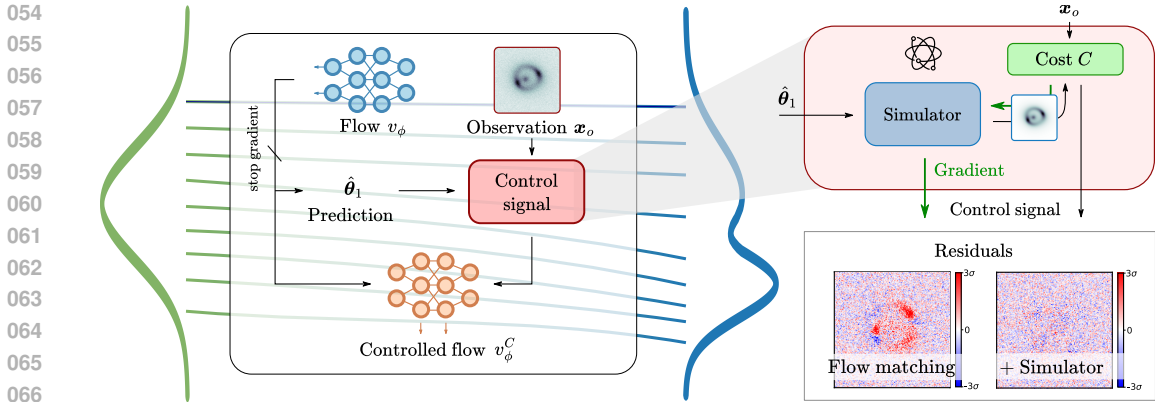


Figure 1: An overview of our proposed framework. We consider a pretrained flow network v_ϕ and use the predicted flow for the trajectory point θ_t at time t to estimate $\hat{\theta}_1$. On the right, we show a gradient-based control signal with a differentiable simulator and cost function C for improving $\hat{\theta}_1$. An additional network learns to combine the predicted flow with feedback via the control signal to give a new controlled flow. By combining learning-based updates with suitable controls, we avoid local optima and obtain high-accuracy samples with low inference times.

θ during training, which makes it very difficult to produce highly accurate samples based on learning alone.

We propose a simple strategy to reintroduce control signals using **simulators** into the flow network. We refine an existing pretrained flow-based model with a flexible control signal by aggregating the learned flow and control signals into a *controlled flow*, which requires only a minimal amount of additional parameters. To demonstrate how these refinements affect the accuracy of samples and the posterior, we consider modeling strong gravitational lens systems (Hezaveh et al., 2017; Cunha & Herdeiro, 2018; Legin et al., 2021), an inverse problem in astrophysics that is challenging and requires precise posteriors for accurate modeling of observations. In galaxy-scale strong lenses, light from a source galaxy is deflected by the gravitational potential of a galaxy between the source and observer, causing multiple images of the source to be seen. Since these images and their distortions are sensitive to the distribution of matter on small scales, this can act as a probe for different dark matter models. With upcoming and current sky surveys (Laureijs et al., 2011) expected to release large data catalogs in the near future, the number of known lenses will increase dramatically by several orders of magnitude. Traditional computational approaches require several minutes to many hours or days to model a single lens system. Therefore, there is an urgent need to reduce the compute and inference with learning-based methods. In this experiment, we demonstrate that using flow matching and our proposed control signals with feedback from a simulator, we obtain posterior distributions for lens modeling that are competitive with the posteriors obtained by MCMC-based methods but with much faster inference times.

Additionally, we evaluate different related variants of flow matching such as using problem-specific priors, self-conditioning (Chen et al., 2023) or different loss formulations in the context of SBI using several benchmark problems. We then analyze our proposed control signals for the Lotka-Volterra model, a system of coupled ordinary differential equations (ODEs) describing the population evolution of predators and prey over time. Our analysis underscores the essential role of simulator feedback for inference and that high accuracy is very challenging to achieve from scaling up datasets and model sizes alone.

To summarize, the main contributions of our work are:

- We propose a versatile strategy to improve pretrained flows with control signals based on feedback from a simulator. Control signals can be based on gradients and a cost function, if the simulator is differentiable, but they can also be learned directly from the simulator output.
- We assess different variants of flow matching in the context of SBI and demonstrate with the Lotka-Volterra model that performance gains due to simulator feedback are substantial and cannot be achieved by training on larger datasets alone.

- We demonstrate the efficacy of our proposed finetuning with control signals for inferring the parameter distributions of strong gravitational lens systems, a challenging inverse problem in astronomy that is sensitive to sample accuracy. We show that flow matching with simulator feedback is competitive with MCMC baselines and beats them significantly regarding inference time.

2 RELATED WORK

Solving inverse problems under a diffusion prior Diffusion models have been proposed to solve linear inverse problems (Kawar et al., 2021; 2022; Chung et al., 2022; Cardoso et al., 2023), as well as general inverse problems (Holzschuh et al., 2023; Song et al., 2023; Chung et al., 2023a;b), via stochastic optimization (Graikos et al., 2022; Mardani et al., 2024) or through amortization by reinforcement learning (Black et al., 2024; Fan et al., 2023). In most of these works, the diffusion model learns the prior distribution and sampling from the posterior is achieved through a modified inference procedure, which guides samples via a conditioning. The conditioning can be based on a class label, text input (Song et al., 2021; Ho & Salimans, 2022; Saharia et al., 2022; Wu et al., 2023) or directly on a differentiable measurement operator (Chung et al., 2023a;b). In contrast to these works, we finetune a pretrained flow and learn an optimal combination of the pretrained flow and feedback from a simulator via control signals in the broader flow matching context.

Flow matching Our work builds on top of prior work in flow matching (Lipman et al., 2023; Albergo et al., 2023a; Pooladian et al., 2023; Tong et al., 2023; Albergo et al., 2023b), particularly we adopt and evaluate conditional optimal transport paths (Lipman et al., 2023), [test problem-specific priors and rectification of flows to produce straighter paths](#) (Liu et al., 2023) for simulation-based inference. [Guiding flows has for example been explored by Zheng et al. \(2023\); Nisonoff et al. \(2024\).](#) We extend the existing literature by adding feedback from a simulator for scientific inverse problems.

Simulation-based inference Our work directly compares to neural posterior estimation approaches for simulation-based inference (Cranmer et al., 2020; Lueckmann et al., 2021, SBI). Contrary to static architectures (Dinh et al., 2017; Kingma & Dhariwal, 2018; Papamakarios et al., 2017; Durkan et al., 2019), our approach extends the continuous-time paradigm (Chen et al., 2018; Grathwohl et al., 2019). Wildberger et al. (2023) have applied flow matching to neural posterior estimation and Sharrock et al. (2022) have used conditional diffusion models and Langevin dynamics during sampling. In contrast to previous work, we include controls signals via problem-specific simulators and cost functions during training to significantly improve the sampling quality.

Strong lensing and parameter estimation Machine learning has been successfully applied to estimate parameters of lens and source models (Hezaveh et al., 2017; Levasseur et al., 2017), however, previous methods are usually restricted to point estimates, use simple variational distributions, Bayesian Neural Networks (Schuldt et al., 2021; Legin et al., 2021; Poh et al., 2022) that are not well suited to represent more complicated high-dimensional data distributions. [Legin et al. \(2023\) predict point estimates for the lensing parameters, which are utilized by mixture density networks to model their distribution in a likelihood-free inference framework.](#) In this paper, we combine flow matching with problem-specific simulators to obtain highly accurate samples via feedback from control signals.

3 FLOW MATCHING THEORY

Continuous-time flow models transform samples θ from a sampling distribution p_0 to samples of a target or posterior distribution p_1 . This mapping can be expressed via the ODE

$$d\theta_t = v_\phi(t, \theta_t)dt, \tag{1}$$

where $v_\phi(t, \theta_t)$ represents a neural network with parameters ϕ . Early works (Chen et al., 2018; Grathwohl et al., 2019) optimize $v_\phi(t, \theta)$ using maximum likelihood training, which is computationally demanding and difficult to scale to larger networks. Instead, in flow matching the network $v_\phi(t, \theta)$ is trained by regressing a vector field $u(t, \theta)$ that generates probability paths that map from p_0 to p_1 .

Generating probability paths We say that a [smooth](#)¹ vector field $u : [0, 1] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$, called *velocity*, generates the probability paths p_t , if it satisfies the continuity equation $\frac{\partial p}{\partial t} = -\nabla \cdot (p_t u_t)$

¹the vector field u is locally Lipschitz in θ and Bochner integrable in t

when viewed as a function $p : [0, 1] \times \mathbb{R}^d \rightarrow \mathbb{R}$. Informally, this means that we can sample from the distribution p_t by sampling $\theta_0 \sim p_0$ and then solving the ODE $d\theta = u(t, \theta)dt$ with initial condition θ_0 . In the following, we will denote $u(t, \theta)$ by $u_t(\theta)$. To regress the velocity field, we define the **flow matching** objective

$$\mathcal{L}_{\text{FM}}(\theta) := \mathbb{E}_{t \sim \mathcal{U}(0,1), \theta \sim p_t(\theta)} \|v_\theta(t, \theta) - u_t(\theta)\|^2. \quad (2)$$

In order to compute this loss, we need to sample from the probability distribution $p_t(\theta)$ and we need to know the velocity $u_t(\theta)$. However, in general $u_t(\theta)$ is not accessible.

Conditioning variable To solve this problem, we apply a trick by introducing a latent variable z distributed according to $q(z)$ and define the conditional likelihoods $p_t(\theta|z)$ that depend on the latent variable so that $p_t(\theta) = \int p_t(\theta|z)q(z)dz$. Interestingly, if the conditional likelihoods are generated by the velocities $u_t(\theta|z)$, then the velocity $u_t(\theta)$ can be written in terms of $u_t(\theta|z)$ and $p_t(\theta|z)$ with $u_t(\theta) := \mathbb{E}_{q(z)}[u_t(\theta|z)p_t(\theta|z)/p_t(\theta)]$. We can choose paths $p_t(\theta|z)$ that are easy to sample from and for which we know the generating velocities $u_t(\theta|z)$. Next, we define the **conditional flow matching** loss

$$\mathcal{L}_{\text{CFM}}(\phi) := \mathbb{E}_{t, q(z), p_t(\theta|z)} \|v_\phi(t, \theta) - u_t(\theta|z)\|^2. \quad (3)$$

In contrast to the flow matching loss eq. 2, this loss is tractable and can be used for optimization. Now, one can show (Tong et al., 2023) that if $p_t(\theta) > 0$ for all $\theta \in \mathbb{R}^d$, then

$$\nabla_\phi \mathcal{L}_{\text{FM}}(\phi) = \nabla_\phi \mathcal{L}_{\text{CFM}}(\phi). \quad (4)$$

This means that we can train $v_\theta(t, \theta)$ to regress $u_t(\theta)$ generating the mapping between p_0 and p_1 by optimizing the conditional flow matching loss eq. 3.

Couplings The above framework allows for many degrees of freedom when specifying the mapping from p_0 to p_1 via the conditioning variable z and the conditional likelihoods p_t . One particularly intuitive and simple choice is to consider the coupling $q(z) = p_1(\theta)$, i.e. the conditioning variable z is identified with the endpoint θ_1 (Lipman et al., 2023), together with conditional probability and generating velocity

$$p_t(\theta|\theta_1) = \mathcal{N}(\theta|t\theta_1, (1 - (1 - \sigma_{\min})t)I) \quad \text{and} \quad u_t(\theta|\theta_1) = \frac{\theta_1 - (1 - \sigma_{\min})\theta}{1 - (1 - \sigma_{\min})t}, \quad (5)$$

where $\sigma_{\min} > 0$. Conditioned on θ_1 , this coupling transports a point $\theta_0 \sim \mathcal{N}(0, I)$ from the sampling distribution to the posterior distribution on the linear trajectory $t\theta_1$ ending in θ_1 but decreasing the standard deviation from 1 to a smoothing constant σ_{\min} . In this case, the transport path coincides with the optimal transport between two Gaussian distributions.

4 CONTROLS FOR IMPROVED ACCURACY

While flow-based models $v_\phi(t, \theta)$ gradually transform samples from p_0 to p_1 in many steps during inference via solving the ODE eq. 1, there is no direct feedback loop between the underlying simulator, the current point on the trajectory θ_t , and the observation x_o . A central goal of our work is to reintroduce this feedback loop into inference and training by incorporating a control signal.

Conditioning of flows Flows $v_\phi(t, \theta)$ can be conditioned on an observation x_o through an additional input $v_\phi(t, \theta, x_o)$, therefore modeling the conditional densities $p_t(\theta|x_o)$ (Song et al., 2021). Models can be trained for both conditional and unconditional generation. This is achieved, for example, in classifier free-guidance (Ho & Salimans, 2022), by randomly dropping the conditioning and setting it to 0 during training.

A critical shortcoming here is that the conditioning x_o is static, whereas we propose to have a dynamic control mechanism that depends on the trajectory θ_t , the observation, and an underlying control signal. The latter should relate θ_t and observation using a physics-based model represented through a cost function C . As the accuracy of neural networks is inherently limited by the finite size of their weights, and smaller networks are attractive from a computational perspective, physics-based control has the potential to yield high accuracy with lean and efficient neural network models.

1-step prediction An additional issue is that the current trajectory θ_t might not be close to a good estimate of a posterior sample θ_1 , especially at the beginning of inference, where θ_0 is drawn from the sampling distribution. This issue is alleviated by applying the cost function C to the current estimate θ_t , we extrapolate θ_t forward in time to obtain an estimated $\hat{\theta}_1$

$$\hat{\theta}_1 = \theta_t + (1 - t)v_\phi(t, \theta_t, x_o). \quad (6)$$

This estimate is exact, if the trained model perfectly fits the conditional optimal transport paths.

Comparison with likelihood-guidance The 1-step prediction is conceptually related to diffusion sampling using likelihood-guidance (Chung et al., 2022; Wu et al., 2023). For inference in diffusion models, sampling is based on the conditional score $\nabla_{\theta_t} \log p(\theta_t|x_o)$, which can be decomposed into

$$\nabla_{\theta_t} \log p(\theta_t|x_o) = \nabla_{\theta_t} \log p(\theta_t) + \nabla_{\theta_t} \log p(x_o|\theta_t). \quad (7)$$

The first expression can be estimated using a pretrained diffusion model, whereas the latter is usually intractable, but can be approximated using $p(x_o|\theta_t) \approx p_{x_o|\theta_0}(x_o|\hat{\theta}(\theta_t))$, where the denoising estimate $\hat{\theta}(\theta_t) := \mathbb{E}_q[\theta_0|\theta_t]$ is usually obtained via Tweedie’s formula $(\mathbb{E}_q[\theta_0|\theta_t] - \theta_t)/t\sigma^2$. In practice, the estimate $\hat{\theta}(\theta_t)$ is very poor when θ_t is still noisy, impeding the inference in the early stages. On the contrary, flows based on linear conditional transportation paths have empirically been shown to have trajectories with less curvature (Lipman et al., 2023) compared to, for example, diffusion models, thus enabling inference in fewer steps and providing better estimates for $\hat{\theta}_1$.

Controlled flow v_ϕ^C We pretrain the flow network $v_\phi(t, \theta, x_o)$ without any control signals to make sure that we can realize the best achievable performance possible based on learning alone. Then, in a second training phase, we introduce the control network $v_\phi^C(t, v, c)$ with pretrained flow v and control signal c as input. The control network is much smaller in size than the flow network, making up ca. 10% of the weights ϕ in our large-scale experiments. We freeze the network weights of v_ϕ and train with the conditional flow matching loss eq. 3 for a small number of additional steps. This reduces training time and compute since we do not need to backpropagate gradients through $v_\phi(t, \theta, x_o)$. We did not observe that freezing the weights of v_ϕ affects the performance negatively. We include algorithms for training in appendix A.

4.1 TYPES OF CONTROL SIGNALS

Aiming for high inference accuracy, we extend self-conditioning via physics-based control signals to include an additional feedback loop between the model output and an underlying physics-based prior. We distinguish between two types of control signals.

Gradient-based control signal In the first case, there is a differentiable cost function C and a deterministic differentiable simulator S as shown in fig. 2a. Given an observation x_o and the estimated prediction $\hat{\theta}_1$, the control signal relates to how well $\hat{\theta}_1$ explains x_o via some cost function C . The cost function can also depend directly on or be equal to the likelihood $p(x_o|\hat{\theta}_1)$. For a differentiable cost function C , we define the control signal via


$$c(\hat{\theta}_1, x_o) := [C(S(\hat{\theta}_1), x_o); \nabla_{\hat{\theta}_1} C(S(\hat{\theta}_1), x_o)]. \quad (8)$$

We can use any control that depends on $\hat{\theta}_1$ and x_o and is informative for the given task.


Learning-based control signal In the second case, the simulator is non-differentiable. To combine the simulator output with the observation x_o , we introduce a learnable encoder model Enc with parameters ϕ_E . The output of the encoder is small and of size $O(\dim(\theta))$. The control signal is then defined as

$$c(\hat{\theta}_1, x_o) := Enc(S(\hat{\theta}_1), x_o). \quad (9)$$

The gradient backpropagation is stopped at the simulator output, see fig. 2b.



(a) Gradient-based control signal



(b) Learning-based control signal

Figure 2: Control signals with simulator feedback.

4.2 ADDITIONAL CONSIDERATIONS FOR SIMULATOR FEEDBACK

Stochastic simulators Many Bayesian inference problems have a stochastic simulator. For simplicity, we assume that all stochasticity within such a simulator can be controlled via a variable $z \sim \mathcal{N}(0, I)$, which is an additional input. Motivated by the equivalence of exchanging expectation and gradient

$$\nabla_{\hat{\theta}_1} \mathbb{E}_{z \sim \mathcal{N}(0, I)} [C(S_z(\hat{\theta}_1), \mathbf{x}_o)] = \mathbb{E}_{z \sim \mathcal{N}(0, I)} [\nabla_{\hat{\theta}_1} C(S_z(\hat{\theta}_1), \mathbf{x}_o)], \quad (10)$$

when calling the simulator, we draw a random realization of z . During training, we randomly draw z for each sample and step while during inference we keep the value of z fixed for each trajectory.

Time-dependence If the estimate $\hat{\theta}_1$ is bad and the corresponding cost $C(\hat{\theta}_1, \mathbf{x}_o)$ is high, gradients and control signals can become unreliable. In appendix B, we empirically find that the estimates $\hat{\theta}_1$ become more reliable for $t \geq 0.8$. Therefore, we only train the control network v_ϕ^C in this range, which allows for focusing on control signals containing the most useful information. For $t < 0.8$, we directly output the pretrained flow $v_\phi(t, \theta, \mathbf{x}_o)$.

Theoretical correctness Contrary to likelihood-based guidance, which uses an approximation for $\nabla_{\theta_t} \log p(\mathbf{x}_o | \theta_t)$ as a guidance term during inference, the approximation $\hat{\theta}_1$ only influences the control signal, which is an input to the controlled flow network v_ϕ^C . In the case of a deterministic simulator, this makes the control signal a function of θ_t . The controlled flow network is trained with the same loss as vanilla flow matching (Lipman et al., 2023). Therefore all theoretical properties remain preserved.

5 SIMULATION-BASED INFERENCE

This section is organized as follows. First, in section 5.1, we introduce a set of SBI benchmark tasks and provide a comparison of popular neural posterior estimation (NPE) methods against a baseline of flow matching without simulator feedback. This comparison uses a similar training setup for all models and tasks. Then, in section 5.2, we focus on an optimal task-specific network with training hyperparameters based on an extensive grid search. We evaluate different variants of flow matching that are related to simulator feedback on the SBI tasks to push the performance as far as possible. In section 5.3, we pick the most challenging SBI task and improve it further by introducing simulator feedback via gradient-based and learned control signals. We carefully analyze the cost-accuracy trade-off for using simulators and show that improvements from simulator feedback cannot be replicated by increasing the training dataset size alone.

5.1 TASKS AND BASELINES

We consider the SBI tasks Lotka Volterra **LV**, a coupled ODE for the population dynamics of interacting species, **SIR**, an epidemiological model for the spread of diseases, **SLCP** and Two Moons (**TM**), two synthetic tasks having complicated multimodal posteriors. All tasks are part of the benchmark collection from Lueckmann et al. (2021). For each problem, the posterior distribution for a set of 10 observations is known, which allows for directly comparing it with the posterior predicted by the trained model. This is measured using the C2ST score (Lopez-Paz & Oquab, 2017), which trains a classifier to discriminate between samples from the true posterior and samples generated from the learned model. If the classifier cannot discriminate between two sets of samples, its test accuracy will be 0.5, whereas it increases when they become more dissimilar.

We include the following baseline methods for NPE: Continuous normalizing flows (Chen et al., 2018, CNF), Neural Spline Flows (Durkan et al., 2019, NSF), and FFIJORD (Grathwohl et al., 2019). Since we propose to include feedback from simulators, here we focus on the largest benchmark budget of 10^5 simulator calls for generating the training dataset. Table 1 highlights that flow matching yields a highly competitive performance in this setting. For details on the training setup, see appendix B.

Table 1: C2ST comparison with identical training setups and comparable number of network weights (ca. 300K).

Method	LV	SLCP	SIR	TM
CNF	0.99	<u>0.80</u>	0.99	0.60
NSF	0.99	-	0.75	0.54
FFJORD	<u>0.95</u>	0.82	<u>0.78</u>	0.59
<i>Flow-Mat.</i>	0.93	0.79	<u>0.79</u>	<u>0.58</u>




Figure 3: Evaluation of SBI tasks using different variants of flow matching training. Lower C2ST scores are better.

Flow matching has also been evaluated for the SBI benchmark tasks by Wildberger et al. (2023), who performed an extensive hyperparameter search for each task to find optimal hyperparameters. In the following, we focus on flow matching, and hence use the corresponding sets of optimal hyperparameters for each task.

5.2 TRAINING VARIANTS

There are several variants of training diffusion models that can be related to simulation feedback and which we consider promising in the context of SBI. Before we go on to evaluate the simulator feedback in section 5.3, we test if we can improve the performance using any of them. In particular, we assess the following modifications:

- **Self-conditioning:** conditioning a model on something that depends on its own output can be seen as a form of self-conditioning. We evaluate an adapted version of self-conditioning (Chen et al., 2023). Instead of providing θ_t to the flow network, the input is comprised of the concatenated vector $[\theta_t; \text{Dropout}(\hat{\theta}_1)]$, where $\hat{\theta}_1$ is the 1-step prediction eq. 6. For computing $\hat{\theta}_1$, we require one network evaluation with the input $[\theta_t; 0]$ and stop the gradient backpropagation at $\hat{\theta}_1$. This method is similar to our simulator feedback, as it introduces a feedback loop that conditions the model on its own output, but without any simulator.
- **Task-specific priors:** it is also possible to couple two non-Gaussian distributions by defining the coupling as $q(\mathbf{z}) = p_0(\theta_0)p_1(\theta_1)$ and setting the conditional probabilities to the linear paths defined by $p_t(\theta|\theta_0, \theta_1) = \mathcal{N}(\theta|t\theta_1 + (1-t)\theta_0, \sigma I)$ and $u_t(\theta|\theta_0, \theta_1) = \theta_1 - \theta_0$ with bandwidth $\sigma > 0$. We can choose p_0 as the prior distribution $p(\theta)$ which we know in the SBI setting. Obtaining information in the form of an observation changes our knowledge about θ from the prior distribution to the posterior, therefore resembling a transformation similar to the noise to data transformation in diffusion models. This also suggests that the prior distribution can be closer to the posterior than a noise distribution.
- **x -prediction:** the reliability of the control signal depends directly on the 1-step estimate $\hat{\theta}$. Instead of regressing the flow $u_t(\theta)$, we can directly predict the denoised estimate $\hat{\theta}$ and obtain the velocity by rearranging eq. 6, giving $v_\phi(t, \theta_t, \mathbf{x}_o) = \hat{\theta}_1 / (1 - t)$. We additionally weight the x -prediction loss with a time-dependent weighting $w_t := 1 / (1 - t)$ to account for the scaling in eq. 6. The x -prediction potentially produces better estimates for $\hat{\theta}$, thus allowing for obtaining more reliable feedback from control signals when $t < 0.8$.

Evaluation Figure 3 shows an evaluation of the different variants against vanilla flow matching (Gaussian sampling distribution, no self-conditioning and velocity prediction). Using task-specific priors produces outliers with better C2ST scores for SLCP but is consistently worse for LV and SIR. We conclude that normal Gaussian distributions are more suited as sampling distributions for most low-dimensional problems. Introducing self-conditioning does not show any improvements, so feedback loops without a simulator alone are not sufficient for better performance in this situation. Finally, the x -prediction loss consistently performs worse than the velocity prediction. Therefore, a potential improvement in the 1-step estimate is outweighed by a corresponding deterioration of the posterior correctness as indicated by the C2ST score.

5.3 SIMULATOR FEEDBACK: GRADIENT-BASED AND LEARNED

In this section, we focus on the Lotka-Volterra (LV) task for a more detailed analysis. It has the highest difficulty as seen by the C2ST score, and we use it to test the different types of feedback. We reimplement the LV simulator in JAX (Bradbury et al., 2018) to support differentiability and evaluate the gradient-based control signal as well as the learning-based control signal, using a small multilayer perceptron (MLP). In addition, to make sure that observed improvements are not due to the increased number of network parameters and finetuning with the control network, we also evaluate a variant where we finetune with the control network but set all simulator-dependent inputs to the control network to 0 (Zero Controls). We show an evaluation with C2ST in fig. 4. For both the learning and gradient-based control signals we see clear improvements with the gradient-based signal clearly ahead. The zero control signal improves only slightly, showing that the improvement can be directly attributed to the simulator.

While control signals are most useful for more high-dimensional problems with less sparse and noisy observations, this experiment demonstrates that they can also be used in low-dimensional settings. Moreover, while differentiable simulators can provide better control signals, feedback from non-differentiable simulators likewise shows clear improvements.

5.4 COMPUTATIONAL EFFICIENCY

A critical issue in SBI is that calls to the simulator are potentially expensive. This imposes the question of whether compute time is better spent on extending the training dataset or training with feedback from the simulator. We empirically verify that the latter is more efficient for the LV task in this setup by comparing our method to models with an increased training dataset from a larger simulator budget. Specifically, we train with dataset sizes of 10^6 and 10^7 . Training the gradient-based control signal took ca. 9×10^6 simulator calls. See fig. 5 for the evaluation. There is no improvement in the C2ST for models trained without simulator feedback beyond 10^5 data points, and the final train/validation loss for the 10^7 model indicates that there is no more overfitting. Nonetheless, the model trained with controls clearly outperforms the model trained with more data, indicating that the directed feedback of the simulator cannot be replaced by increased amounts of training data.

6 STRONG GRAVITATIONAL LENSING

We present our results for modeling strong gravitational lens systems, a challenging and highly relevant non-linear problem in astronomy. Strong gravitational lensing is a physical phenomenon whereby the light rays by a distant object, such as a galaxy, are deflected by an intervening massive object, such as another galaxy or a galaxy cluster. As a result, one observes multiple distorted images of the background source. We aim to recover both the lens and source light distribution as well as the lens mass density distribution with realistic simulated observations for which we know the ground truths. We evaluate flow matching as an NPE method with gradient-based control signals from a differentiable simulator with two MCMC methods.

Lens modeling The *lens equation* relates coordinates on the source plane β and the observed image plane Θ via the deflection angle α induced by the mass profile or gravitational potential of the lens galaxy. We use a Singular Isothermal Ellipsoid (SIE) to describe this lens mass and Sérsic profiles for both the source light and light emitted from the lens galaxy (full details are provided in appendix C). There are 9 parameters for the lens mass and 7 parameters for each Sérsic profile, giving 23 parameters in total. The likelihood is measured by the χ^2 -statistic, which is the modeled image plane Θ minus the observation x_o divided by the noise. To solve the lensing equation, we make use of




Figure 4: Evaluation of simulator feedback for LV.




Figure 5: Different simulator call budgets (training set sizes 10^5 , 10^6 , 10^7) compared with finetuning using simulator feedback (ours, ca. 9×10^6 simulator calls in total).




Figure 6: Reconstruction of observation. Flow matching is purely learning-based and shows noticeable residuals in the reconstruction. Including simulator feedback removes remaining residuals.

the publicly available raytracing code by (Galan et al., 2022). We want to stress that even small perturbations of the model parameters can cause the χ_2 to increase significantly; see fig. 14 in the appendix.

Datasets and pretraining Several instrument-specific measurement effects are included when simulating the observations. We include background and Poisson noise and smoothing by a point-spread function (PSF). The pixel size corresponds to 0.04 arc seconds. These directly affect the posterior, as more noise and a stronger PSF will widen the posterior distribution. We generate 250 000 data samples for training and 25 000 for validation. The flow network v_ϕ consists of a convolutional feature extraction neural network represented by a shallow CNN whose output is fed into a dense feed-forward neural network with residual blocks. Full details are in appendix C.

Finetuning with control signals The control network v_ϕ^C is represented by another dense feed-forward network, which accounts for 11% of all parameters in the combined model. The control signals are obtained from simulating an observation based on the predicted estimate $\hat{\theta}_1$ via ray-tracing (Galan et al., 2022) based on the parametric models, calculating the χ_2 -statistic and computing gradients with respect to the estimate $\hat{\theta}_1$. The χ_2 -statistic itself is also part of the control signal.

Reference posteriors As reference posteriors, we include Hamiltonian Monte Carlo (HMC) with No-U-Turn sampler (Hoffman et al., 2014, NUTS) and Affine-Invariant Ensemble Sampling (Goodman & Weare, 2010, AIES), which are both two popular MCMC-methods in astronomy. We adopt implementations of both methods using numpyro (Phan et al., 2019; Bingham et al., 2019). **Additionally, we compare to diffusion posterior sampling (Chung et al., 2023b, DPS), loss-guided diffusion (Song et al., 2023, LGD-MC) and twisted diffusion sampler (Wu et al., 2023, TDS).** Details on all baseline methods can be found in appendix C. We use Euler integration for both flow matching variants.

6.1 EVALUATION AND DISCUSSION

χ_2 -statistic We show an evaluation of all methods in table 2. The average χ_2 is computed over 1000 randomly chosen validation systems, where for each, we draw 1000 samples from the posterior. If we compute the χ_2 for the ground truth parameters, we obtain a value of 1.17 due to the noise in the observation. Since we cannot overfit to noise with the parametric models, this represents a lower bound for χ_2 in this experiment. Including the physics-based control improves the χ_2 from 1.83 to 1.48, representing an improvement of 53% relative to the best modeling. The improved χ_2 is even better than the best baseline method, AIES.

Modeling time We define the modeling time as the average compute time required to produce 1000 credible posterior samples. Both HMC and AIES require significant warmup times before producing the first samples from the posterior, which we include in the table. However, after warmup, it is relatively cheap to obtain new samples. On the other hand, flow matching does not require

Table 2: Evaluation with respect to average χ_2 and inference time for the posterior distribution.

Method	Avg. $\chi_2 \downarrow$	Modeling Time \downarrow
NUTS	1.83	$\sim 56x$ (564s)
AIES	1.74	$\sim 67x$ (672s)
DPS	9.98	$\sim 42x$ (427s)
LGD-MC(5)	21.62	$\sim 160x$ (1600s)
TDS (k=100)	20.94	$\sim 21x$ (210s)
<i>Flow-Mat.</i>	1.83	1x (10s)
+ Simulator	1.48	$\sim 2x$ (19s)

any warmup time and the modeling time increases linearly with the number of posterior samples. All methods were implemented in JAX (Bradbury et al., 2018) and used the same hardware. The measurements in table 2 show that DPS is faster than the classic baselines, but yields a very sub-optimal performance in terms of its distribution. The performance numbers also highlight that our method yields an accuracy that surpasses AIES, while being more than 30x faster.

This evaluation demonstrates that flow matching-based methods are highly competitive even in small to moderate-sized problems where established MCMC methods in terms of accuracy exist, clearly beating them in terms of inference time. Flow matching with our proposed control signals is especially interesting because it is not affected as much by the curse of dimensionality as traditional inference methods and allows for having non-trivial learnable high-dimensional priors. However, before these methods are widely trusted, they need to demonstrate their competitiveness with classical methods. Our results show that this is indeed the case, which opens up exciting avenues for applying and developing approaches targeting similar and adjacent inverse problems in science.

Simulation-based calibration Acquiring truthful posterior distributions for Bayesian inference problems is difficult, which makes it hard to robustly evaluate whether the predicted posterior distribution is correct. We use simulation-based calibration (Talts et al., 2018, SBC) as an additional evaluation tool. The data-averaged posterior obtained from averaging the posterior distribution over many problem instances has to be equal to the prior. This can be tested by considering a one-dimensional function $f : \theta \mapsto \mathbb{R}$ and L samples $\theta^1, \dots, \theta^L$ drawn from an inference method. If θ^* are the ground truth parameters, then the rank statistic $\sum_{l=1}^L \mathbf{1}_{f(\theta^l) < f(\theta^*)}$ has to be uniformly distributed over the integers $[0, L]$. If the distribution of the rank statistic is plotted as a histogram, systematic problems in the inference method can be identified visually, see fig. 7. We set $L = 1000$ and plot the histograms for all $n = 1000$ test problems and visualize the parameter x_{center} , which defines the position of the source in x -direction. The posteriors without simulator feedback are biased, as can be seen in the deviation from uniformity in the plots. Including simulator feedback improves the distribution of the rank statistic. For an extended analysis, see appendix C.4.

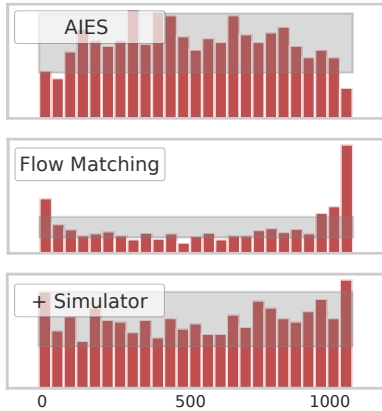


Figure 7: SBC for x_{center} of the source galaxy.

6.2 LIMITATIONS

While introducing additional control signals increases the quality of produced samples, it comes at the cost of slower inference and training times depending on the speed of the simulator. In general, using non-differentiable control signals is possible but removes the possibility of computing likelihoods via the instantaneous change of variables formula (Chen et al., 2018). Compared to MCMC approaches, inference with flow-based models requires a substantial upfront cost for training that needs to be amortized across many problems. Additionally, priors are encoded in the learned flow networks, so changing them would require retraining models with adjusted data sets.

7 CONCLUSION

We presented a method for improving flow-based models with simulator feedback using control signals. This allows us to refine an existing flow with only a few additional weights and little training time. We thereby efficiently bridge the gap between purely learning-based methods for simulation-based inference and optimization with hand-crafted cost functions within the framework of flow matching. This improvement is critical for scientific applications where high accuracy and trustworthiness in the methods are required. Purely learning-based methods face significant difficulties in producing very accurate samples, as there is usually no feedback during inference of how good samples are. In this paper, we demonstrated that we do not need large network sizes or tremendous amounts of data to train accurate models that are competitive with established MCMC methods if we include suitable control signals from simulators. We believe this work makes an important step towards making posterior inference in science more accurate, understandable, and reliable.

REFERENCES

- 540
541
542 Michael S. Albergo, Nicholas M. Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unifying
543 framework for flows and diffusions. *CoRR*, abs/2303.08797, 2023a. doi: 10.48550/ARXIV.2303.
544 08797. URL <https://doi.org/10.48550/arXiv.2303.08797>.
- 545 Michael S. Albergo, Mark Goldstein, Nicholas M. Boffi, Rajesh Ranganath, and Eric Vanden-Eijnden.
546 Stochastic interpolants with data-dependent couplings. *CoRR*, abs/2310.03725, 2023b. doi: 10.
547 48550/ARXIV.2310.03725. URL <https://doi.org/10.48550/arXiv.2310.03725>.
- 548
549 Atilim Gunes Baydin, Lei Shao, Wahid Bhimji, Lukas Heinrich, Saeid Naderiparizi, Andreas Munk,
550 Jialin Liu, Bradley Gram-Hansen, Gilles Louppe, Lawrence Meadows, Philip H. S. Torr, Victor W.
551 Lee, Kyle Cranmer, Prabhat, and Frank Wood. Efficient probabilistic inference in the quest for
552 physics beyond the standard model. In *Advances in Neural Information Processing Systems*
553 32, pp. 5460–5473, 2019. URL [https://proceedings.neurips.cc/paper/2019/
554 hash/6d19c113404cee55b4036fcela37c058-Abstract.html](https://proceedings.neurips.cc/paper/2019/hash/6d19c113404cee55b4036fcela37c058-Abstract.html).
- 555 Eli Bingham, Jonathan P. Chen, Martin Jankowiak, Fritz Obermeyer, Neeraj Pradhan, Theofanis
556 Karaletsos, Rohit Singh, Paul A. Szerlip, Paul Horsfall, and Noah D. Goodman. Pyro: Deep
557 universal probabilistic programming. *J. Mach. Learn. Res.*, 20:28:1–28:6, 2019. URL [http:
558 //jmlr.org/papers/v20/18-403.html](http://jmlr.org/papers/v20/18-403.html).
- 559
560 Kevin Black, Michael Janner, Yilun Du, Ilya Kostrikov, and Sergey Levine. Training diffusion
561 models with reinforcement learning. In *The Twelfth International Conference on Learning*
562 *Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024. URL
563 <https://openreview.net/forum?id=YCWjhGrJFD>.
- 564 James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal
565 Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and
566 Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL
567 <http://github.com/google/jax>.
- 568
569 Gabriel Cardoso, Yazid Janati El Idrissi, Sylvain Le Corff, and Eric Moulines. Monte carlo guided
570 diffusion for bayesian linear inverse problems. *CoRR*, abs/2308.07983, 2023. doi: 10.48550/
571 ARXIV.2308.07983. URL <https://doi.org/10.48550/arXiv.2308.07983>.
- 572 Tian Qi Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. Neural ordinary
573 differential equations. In *Advances in Neural Information Processing Systems 31*, pp.
574 6572–6583, 2018. URL [https://proceedings.neurips.cc/paper/2018/hash/
575 69386f6bb1dfed68692a24c8686939b9-Abstract.html](https://proceedings.neurips.cc/paper/2018/hash/69386f6bb1dfed68692a24c8686939b9-Abstract.html).
- 576
577 Ting Chen, Ruixiang Zhang, and Geoffrey E. Hinton. Analog bits: Generating discrete data using
578 diffusion models with self-conditioning. In *The Eleventh International Conference on Learning*
579 *Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023. URL
580 <https://openreview.net/pdf?id=3itjR9QxFw>.
- 581 Hyungjin Chung, Byeongsu Sim, and Jong Chul Ye. Come-closer-diffuse-faster: Accelerating
582 conditional diffusion models for inverse problems through stochastic contraction. In *IEEE/CVF*
583 *Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA,*
584 *June 18-24, 2022*, pp. 12403–12412. IEEE, 2022. doi: 10.1109/CVPR52688.2022.01209. URL
585 <https://doi.org/10.1109/CVPR52688.2022.01209>.
- 586
587 Hyungjin Chung, Jeongsol Kim, Michael Thompson McCann, Marc Louis Klasky, and Jong Chul
588 Ye. Diffusion posterior sampling for general noisy inverse problems. In *The Eleventh Interna-*
589 *tional Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*.
590 OpenReview.net, 2023a. URL <https://openreview.net/pdf?id=OnD9zGAGT0k>.
- 591 Hyungjin Chung, Jeongsol Kim, and Jong Chul Ye. Direct diffusion bridge using data
592 consistency for inverse problems. In *Advances in Neural Information Processing Sys-*
593 *tems 36*, 2023b. URL [http://papers.nips.cc/paper_files/paper/2023/hash/
165b0e600b1721bd59526131eb061092-Abstract-Conference.html](http://papers.nips.cc/paper_files/paper/2023/hash/165b0e600b1721bd59526131eb061092-Abstract-Conference.html).

- 594 Kyle Cranmer, Johann Brehmer, and Gilles Louppe. The frontier of simulation-based inference.
595 *Proceedings of the National Academy of Sciences*, 117(48):30055–30062, 2020.
- 596
- 597 Pedro VP Cunha and Carlos AR Herdeiro. Shadows and strong gravitational lensing: a brief review.
598 *General Relativity and Gravitation*, 50:1–27, 2018.
- 599 Maximilian Dax, Stephen R Green, Jonathan Gair, Jakob H Macke, Alessandra Buonanno, and
600 Bernhard Schölkopf. Real-time gravitational wave science with neural posterior estimation.
601 *Physical review letters*, 127(24):241103, 2021.
- 602
- 603 Prafulla Dhariwal and Alexander Quinn Nichol. Diffusion models beat gans on image synthesis. In
604 Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman
605 Vaughan (eds.), *Advances in Neural Information Processing Systems 34: Annual Conference on
606 Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pp.
607 8780–8794, 2021. URL [https://proceedings.neurips.cc/paper/2021/hash/
608 49ad23d1ec9fa4bd8d77d02681df5cfa-Abstract.html](https://proceedings.neurips.cc/paper/2021/hash/49ad23d1ec9fa4bd8d77d02681df5cfa-Abstract.html).
- 609 Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real NVP. In *5th
610 International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26,
611 2017, Conference Track Proceedings*. OpenReview.net, 2017. URL [https://openreview.
612 net/forum?id=HkpbnH9lx](https://openreview.net/forum?id=HkpbnH9lx).
- 613 Conor Durkan, Artur Bekasov, Iain Murray, and George Papamakarios. Neural
614 spline flows. In *Advances in Neural Information Processing Systems 32*, pp. 7509–
615 7520, 2019. URL [https://proceedings.neurips.cc/paper/2019/hash/
616 7ac71d433f282034e088473244df8c02-Abstract.html](https://proceedings.neurips.cc/paper/2019/hash/7ac71d433f282034e088473244df8c02-Abstract.html).
- 617
- 618 Ying Fan, Olivia Watkins, Yuqing Du, Hao Liu, Moonkyung Ryu, Craig Boutilier, Pieter Abbeel,
619 Mohammad Ghavamzadeh, Kangwook Lee, and Kimin Lee. DPOK: reinforcement learning for
620 fine-tuning text-to-image diffusion models. *CoRR*, abs/2305.16381, 2023. doi: 10.48550/ARXIV.
621 2305.16381. URL <https://doi.org/10.48550/arXiv.2305.16381>.
- 622 A. Galan, G. Vernardos, A. Peel, F. Courbin, and J. L. Starck. Using wavelets to capture deviations
623 from smoothness in galaxy-scale strong lenses. *Astronomy and Astrophysics*, 668:A155, December
624 2022. doi: 10.1051/0004-6361/202244464.
- 625 Tilmann Gneiting and Adrian E Raftery. Weather forecasting with ensemble methods. *Science*, 310
626 (5746):248–249, 2005.
- 627
- 628 Jonathan Goodman and Jonathan Weare. Ensemble samplers with affine invariance. *Communications
629 in applied mathematics and computational science*, 5(1):65–80, 2010.
- 630 Alexandros Graikos, Nikolay Malkin, Nebojsa Jojic, and Dimitris Samaras. Diffusion models as plug-
631 and-play priors. In Sanmi Koyejo, S. Mohamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh
632 (eds.), *Advances in Neural Information Processing Systems 35: Annual Conference on Neural
633 Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - De-
634 cember 9, 2022*, 2022. URL [http://papers.nips.cc/paper_files/paper/2022/
635 hash/5e6cec2a9520708381fe520246018e8b-Abstract-Conference.html](http://papers.nips.cc/paper_files/paper/2022/hash/5e6cec2a9520708381fe520246018e8b-Abstract-Conference.html).
- 636
- 637 Will Grathwohl, Ricky T. Q. Chen, Jesse Bettencourt, Ilya Sutskever, and David Duvenaud. FFJORD:
638 free-form continuous dynamics for scalable reversible generative models. In *7th International
639 Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*.
640 OpenReview.net, 2019. URL <https://openreview.net/forum?id=rJxgknCcK7>.
- 641 Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola. A
642 kernel two-sample test. *The Journal of Machine Learning Research*, 13(1):723–773, 2012.
- 643 Yashar D Hezaveh, Laurence Perreault L’Evesque, and Philip J Marshall. Fast automated analysis of
644 strong gravitational lenses with convolutional neural networks. *Nature*, 548(7669):555–557, 2017.
- 645
- 646 Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *CoRR*, abs/2207.12598, 2022.
647 doi: 10.48550/ARXIV.2207.12598. URL [https://doi.org/10.48550/arXiv.2207.
12598](https://doi.org/10.48550/arXiv.2207.12598).

- 648 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in*
649 *Neural Information Processing Systems 33*, 2020. URL [https://proceedings.neurips.](https://proceedings.neurips.cc/paper/2020/hash/4c5bcfec8584af0d967f1ab10179ca4b-Abstract.html)
650 [cc/paper/2020/hash/4c5bcfec8584af0d967f1ab10179ca4b-Abstract.](https://proceedings.neurips.cc/paper/2020/hash/4c5bcfec8584af0d967f1ab10179ca4b-Abstract.html)
651 [html](https://proceedings.neurips.cc/paper/2020/hash/4c5bcfec8584af0d967f1ab10179ca4b-Abstract.html).
- 652 Matthew D Hoffman, Andrew Gelman, et al. The no-u-turn sampler: adaptively setting path lengths
653 in hamiltonian monte carlo. *J. Mach. Learn. Res.*, 15(1):1593–1623, 2014.
- 654 Benjamin J. Holzschuh, Simona Vegetti, and Nils Thuerey. Solving inverse physics
655 problems with score matching. In *Advances in Neural Information Processing Systems*
656 *36*, 2023. URL [http://papers.nips.cc/paper_files/paper/2023/hash/](http://papers.nips.cc/paper_files/paper/2023/hash/c2f2230abc7ccf669f403be881d3ffb7-Abstract-Conference.html)
657 [c2f2230abc7ccf669f403be881d3ffb7-Abstract-Conference.html](http://papers.nips.cc/paper_files/paper/2023/hash/c2f2230abc7ccf669f403be881d3ffb7-Abstract-Conference.html).
- 658 Bahjat Kawar, Gregory Vaksman, and Michael Elad. SNIPS: solving noisy inverse prob-
659 lems stochastically. In *Advances in Neural Information Processing Systems 34*, pp.
660 21757–21769, 2021. URL [https://proceedings.neurips.cc/paper/2021/hash/](https://proceedings.neurips.cc/paper/2021/hash/b5c01503041b70d41d80e3dbe31bbd8c-Abstract.html)
661 [b5c01503041b70d41d80e3dbe31bbd8c-Abstract.html](https://proceedings.neurips.cc/paper/2021/hash/b5c01503041b70d41d80e3dbe31bbd8c-Abstract.html).
- 662 Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising diffu-
663 sion restoration models. In *Advances in Neural Information Processing Systems*
664 *35*, 2022. URL [http://papers.nips.cc/paper_files/paper/2022/hash/](http://papers.nips.cc/paper_files/paper/2022/hash/95504595b6169131b6ed6cd72eb05616-Abstract-Conference.html)
665 [95504595b6169131b6ed6cd72eb05616-Abstract-Conference.html](http://papers.nips.cc/paper_files/paper/2022/hash/95504595b6169131b6ed6cd72eb05616-Abstract-Conference.html).
- 666 Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua
667 Bengio and Yann LeCun (eds.), *3rd International Conference on Learning Representations, ICLR*
668 *2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL [http://](http://arxiv.org/abs/1412.6980)
669 arxiv.org/abs/1412.6980.
- 670 Diederik P. Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1
671 convolutions. In *Advances in Neural Information Processing Systems 31*, pp. 10236–
672 10245, 2018. URL [https://proceedings.neurips.cc/paper/2018/hash/](https://proceedings.neurips.cc/paper/2018/hash/d139db6a236200b21cc7f752979132d0-Abstract.html)
673 [d139db6a236200b21cc7f752979132d0-Abstract.html](https://proceedings.neurips.cc/paper/2018/hash/d139db6a236200b21cc7f752979132d0-Abstract.html).
- 674 Rene Laureijs, J Amiaux, S Arduini, J-L Augueres, J Brinchmann, R Cole, M Cropper, C Dabin,
675 L Duvet, A Ealet, et al. Euclid definition study report. *arXiv preprint arXiv:1110.3193*, 2011.
- 676 Ronan Legin, Yashar Hezaveh, Laurence Perreault Levasseur, and Benjamin Wandelt. Simulation-
677 based inference of strong gravitational lensing parameters. *arXiv preprint arXiv:2112.05278*,
678 2021.
- 679 Ronan Legin, Yashar Hezaveh, Laurence Perreault-Levasseur, and Benjamin Wandelt. A framework
680 for obtaining accurate posteriors of strong gravitational lensing parameters with flexible priors and
681 implicit likelihoods using density estimation. *The Astrophysical Journal*, 943(1):4, 2023.
- 682 Pablo Lemos, Adam Coogan, Yashar Hezaveh, and Laurence Perreault Levasseur. Sampling-based
683 accuracy testing of posterior estimators for general inference. In Andreas Krause, Emma Brunskill,
684 Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *International*
685 *Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume
686 202 of *Proceedings of Machine Learning Research*, pp. 19256–19273. PMLR, 2023. URL
687 <https://proceedings.mlr.press/v202/lemos23a.html>.
- 688 Laurence Perreault Levasseur, Yashar D Hezaveh, and Risa H Wechsler. Uncertainties in parameters
689 estimated with neural networks: Application to strong gravitational lensing. *The Astrophysical*
690 *Journal Letters*, 850(1):L7, 2017.
- 691 Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow
692 matching for generative modeling. In *The Eleventh International Conference on Learning*
693 *Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023. URL
694 <https://openreview.net/pdf?id=PqvMRDCJT9t>.
- 695 Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate
696 and transfer data with rectified flow. In *The Eleventh International Conference on Learning*
697 *Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023. URL
698 <https://openreview.net/pdf?id=XVjTT1nw5z>.

- 702 David Lopez-Paz and Maxime Oquab. Revisiting classifier two-sample tests. In *5th International*
703 *Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Confer-*
704 *ence Track Proceedings*. OpenReview.net, 2017. URL [https://openreview.net/forum?](https://openreview.net/forum?id=SJkXfE5xx)
705 [id=SJkXfE5xx](https://openreview.net/forum?id=SJkXfE5xx).
706
- 707 Jan-Matthis Lueckmann, Jan Boelts, David S. Greenberg, Pedro J. Gonçalves, and Jakob H.
708 Macke. Benchmarking simulation-based inference. In *The 24th International Conference*
709 *on Artificial Intelligence and Statistics, AISTATS 2021, April 13-15, 2021, Virtual Event*, vol-
710 *ume 130 of Proceedings of Machine Learning Research*, pp. 343–351. PMLR, 2021. URL
711 <http://proceedings.mlr.press/v130/lueckmann21a.html>.
- 712 Morteza Mardani, Jiaming Song, Jan Kautz, and Arash Vahdat. A variational perspective on solving
713 inverse problems with diffusion models. In *The Twelfth International Conference on Learning*
714 *Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024. URL
715 <https://openreview.net/forum?id=1Y04EE3SPB>.
716
- 717 Hunter Nisonoff, Junhao Xiong, Stephan Allenspach, and Jennifer Listgarten. Unlocking guidance
718 for discrete state-space diffusion and flow models. *CoRR*, abs/2406.01572, 2024. doi: 10.48550/
719 [ARXIV.2406.01572](https://doi.org/10.48550/arXiv.2406.01572). URL <https://doi.org/10.48550/arXiv.2406.01572>.
- 720 George Papamakarios, Iain Murray, and Theo Pavlakou. Masked autoregressive flow for
721 density estimation. In *Advances in Neural Information Processing Systems 30*, pp.
722 2338–2347, 2017. URL [https://proceedings.neurips.cc/paper/2017/hash/](https://proceedings.neurips.cc/paper/2017/hash/6c1da886822c67822bcf3679d04369fa-Abstract.html)
723 [6c1da886822c67822bcf3679d04369fa-Abstract.html](https://proceedings.neurips.cc/paper/2017/hash/6c1da886822c67822bcf3679d04369fa-Abstract.html).
724
- 725 George Papamakarios, David C. Sterratt, and Iain Murray. Sequential neural likelihood: Fast
726 likelihood-free inference with autoregressive flows. In Kamalika Chaudhuri and Masashi Sugiyama
727 (eds.), *The 22nd International Conference on Artificial Intelligence and Statistics, AISTATS*,
728 *volume 89 of Proceedings of Machine Learning Research*, pp. 837–848. PMLR, 2019. URL
729 <http://proceedings.mlr.press/v89/papamakarios19a.html>.
- 730 Du Phan, Neeraj Pradhan, and Martin Jankowiak. Composable effects for flexible and accelerated
731 probabilistic programming in numpyro. *arXiv preprint arXiv:1912.11554*, 2019.
732
- 733 Jason Poh, Ashwin Samudre, Aleksandra Ćiprijanović, Brian Nord, Gourav Khullar, Dimitrios
734 Tanoglidis, and Joshua A Frieman. Strong lensing parameter estimation on ground-based imaging
735 data using simulation-based inference. *arXiv preprint arXiv:2211.05836*, 2022.
- 736 Aram-Alexandre Pooladian, Heli Ben-Hamu, Carles Domingo-Enrich, Brandon Amos, Yaron Lip-
737 man, and Ricky T. Q. Chen. Multisample flow matching: Straightening flows with minibatch
738 couplings. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan
739 Sabato, and Jonathan Scarlett (eds.), *International Conference on Machine Learning, ICML 2023,*
740 *23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning*
741 *Research*, pp. 28100–28127. PMLR, 2023. URL [https://proceedings.mlr.press/](https://proceedings.mlr.press/v202/pooladian23a.html)
742 [v202/pooladian23a.html](https://proceedings.mlr.press/v202/pooladian23a.html).
- 743 Danilo Jimenez Rezende and Shakir Mohamed. Variational inference with normalizing flows. In
744 *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France,*
745 *6-11 July 2015*, volume 37 of *JMLR Workshop and Conference Proceedings*, pp. 1530–1538.
746 *JMLR.org*, 2015. URL <http://proceedings.mlr.press/v37/rezende15.html>.
747
- 748 Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L. Denton, Seyed
749 Kamyar Seyed Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans,
750 Jonathan Ho, David J. Fleet, and Mohammad Norouzi. Photorealistic text-to-image diffusion
751 models with deep language understanding. In *Advances in Neural Information Processing Sys-*
752 *tems 35*, 2022. URL [http://papers.nips.cc/paper_files/paper/2022/hash/](http://papers.nips.cc/paper_files/paper/2022/hash/ec795aeadae0b7d230fa35cbaf04c041-Abstract-Conference.html)
753 [ec795aeadae0b7d230fa35cbaf04c041-Abstract-Conference.html](http://papers.nips.cc/paper_files/paper/2022/hash/ec795aeadae0b7d230fa35cbaf04c041-Abstract-Conference.html).
- 754 S Schuldts, SH Suyu, T Meinhardt, L Leal-Taixé, R Cañameras, S Taubenberger, and A Halkola.
755 Holismokes-iv. efficient mass modeling of strong lenses through deep learning. *Astronomy &*
Astrophysics, 646:A126, 2021.

- 756 Louis Sharrock, Jack Simons, Song Liu, and Mark Beaumont. Sequential neural score estimation:
757 Likelihood-free inference with conditional score based diffusion models. *CoRR*, abs/2210.04872,
758 2022. doi: 10.48550/ARXIV.2210.04872. URL [https://doi.org/10.48550/arXiv.
759 2210.04872](https://doi.org/10.48550/arXiv.2210.04872).
- 760 Jiaming Song, Qinsheng Zhang, Hongxu Yin, Morteza Mardani, Ming-Yu Liu, Jan Kautz, Yongxin
761 Chen, and Arash Vahdat. Loss-guided diffusion models for plug-and-play controllable gen-
762 eration. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan
763 Sabato, and Jonathan Scarlett (eds.), *International Conference on Machine Learning, ICML 2023,
764 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning
765 Research*, pp. 32483–32498. PMLR, 2023. URL [https://proceedings.mlr.press/
766 v202/song23k.html](https://proceedings.mlr.press/v202/song23k.html).
- 767 Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben
768 Poole. Score-based generative modeling through stochastic differential equations. In *9th Interna-
769 tional Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*.
770 OpenReview.net, 2021. URL <https://openreview.net/forum?id=PxtIG12RRHS>.
- 771 Sean Talts, Michael Betancourt, Daniel Simpson, Aki Vehtari, and Andrew Gelman. Validating
772 bayesian inference algorithms with simulation-based calibration. *arXiv preprint arXiv:1804.06788*,
773 2018.
- 774 Alexander Tong, Nikolay Malkin, Guillaume Hugué, Yanlei Zhang, Jarrid Rector-Brooks, Kilian
775 Fatras, Guy Wolf, and Yoshua Bengio. Conditional flow matching: Simulation-free dynamic
776 optimal transport. *CoRR*, abs/2302.00482, 2023. doi: 10.48550/ARXIV.2302.00482. URL
777 <https://doi.org/10.48550/arXiv.2302.00482>.
- 778 Jonas Wildberger, Maximilian Dax, Simon Buchholz, Stephen R. Green, Jakob H.
779 Macke, and Bernhard Schölkopf. Flow matching for scalable simulation-
780 based inference. In *Advances in Neural Information Processing Systems 36*,
781 2023. URL [http://papers.nips.cc/paper_files/paper/2023/hash/
782 3663ae53ec078860bb0b9c6606e092a0-Abstract-Conference.html](http://papers.nips.cc/paper_files/paper/2023/hash/3663ae53ec078860bb0b9c6606e092a0-Abstract-Conference.html).
- 783 Luhuan Wu, Brian L. Trippe, Christian A. Naesseth, David M. Blei, and John P.
784 Cunningham. Practical and asymptotically exact conditional sampling in dif-
785 fusion models. In *Advances in Neural Information Processing Systems 36*,
786 2023. URL [http://papers.nips.cc/paper_files/paper/2023/hash/
787 63e8bc7bbf1cfea36d1db6538aecce5-Abstract-Conference.html](http://papers.nips.cc/paper_files/paper/2023/hash/63e8bc7bbf1cfea36d1db6538aecce5-Abstract-Conference.html).
- 788 Qinqing Zheng, Matt Le, Neta Shaul, Yaron Lipman, Aditya Grover, and Ricky T. Q. Chen. Guided
789 flows for generative modeling and decision making. *CoRR*, abs/2311.13443, 2023. doi: 10.48550/
790 ARXIV.2311.13443. URL <https://doi.org/10.48550/arXiv.2311.13443>.
- 791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809

APPENDIX

A ALGORITHMS

We include algorithms for training using flow matching and control signals, see algorithm 1. For flow matching with self-conditioning, see algorithm 2.

Algorithm 1 FM with Control Signals

Input: Training distribution q_1 , pretrained network v_ϕ , control network v_ϕ^C , σ_{\min}

while Training **do**

$(\theta_1, \mathbf{x}_o) \sim q_1; z \leftarrow \mathcal{N}(0, I)$

$\theta \leftarrow t\theta_1 + (1-t)z$

$\mathbf{v} \leftarrow \text{stopgrad}(v_\phi(t, \theta, \mathbf{x}_o))$

$\hat{\theta}_1 \leftarrow \theta + (1-t)\mathbf{v}$

$\mathbf{c} \leftarrow \text{control}(\hat{\theta}_1, \mathbf{x}_o)$

$\tilde{\mathbf{v}} \leftarrow v_\phi^C(t, \mathbf{v}, \mathbf{c}) + \mathbf{v}$

$u_t(\theta|\theta_1, \mathbf{x}_o) \leftarrow \frac{\theta_1 - (1 - \sigma_{\min})\theta}{1 - (1 - \sigma_{\min})t}$

$\mathcal{L}_{\text{CFM}} \leftarrow \|\tilde{\mathbf{v}} - u_t(\theta|\theta_1, \mathbf{x}_o)\|_2^2$

$\theta \leftarrow \text{Update}(\phi, \nabla_\phi \mathcal{L}_{\text{CFM}}(\phi))$

return: v_ϕ, v_ϕ^C

Algorithm 2 FM with Self-conditioning

Input: Training distribution q_1 , flow network v_ϕ, σ_{\min}

while Training **do**

$(\theta_1, \mathbf{x}_o) \sim q_1; z \leftarrow \mathcal{N}(0, I); s \leftarrow \mathcal{U}(0, 1)$

$\theta \leftarrow t\theta_1 + (1-t)z; \hat{\theta}_1 \leftarrow 0$

$\mathbf{v} \leftarrow \text{stopgrad}(v_\theta(t, [\theta, \hat{\theta}_1], \mathbf{x}_o))$

if $s > 0.5$ **then**

$\hat{\theta}_1 \leftarrow \theta + (1-t)\mathbf{v}$

$\tilde{\mathbf{v}} \leftarrow v_\phi(t, [\theta, \hat{\theta}_1], \mathbf{x}_o)$

$u_t(\theta|\theta_1, \mathbf{x}_o) \leftarrow \frac{\theta_1 - (1 - \sigma_{\min})\theta}{1 - (1 - \sigma_{\min})t}$

$\mathcal{L}_{\text{CFM}} \leftarrow \|\tilde{\mathbf{v}} - u_t(\theta|\theta_1, \mathbf{x}_o)\|_2^2$

$\theta \leftarrow \text{Update}(\phi, \nabla_\phi \mathcal{L}_{\text{CFM}}(\phi))$

return: v_ϕ

B SIMULATION-BASED INFERENCE

Baselines comparison in section 5.1 For a fairer comparison, we set up all baseline methods with a similar number of network weights and available compute time.

We train all baselines and flow matching with a batch size of 512 on the largest 10^5 simulation budgets for all tasks. For optimization, we apply Adam (Kingma & Ba, 2015) with default settings and constant learning rate of 10^{-4} and weight decay 2×10^{-5} .

All network architectures are chosen to have a similar number of ca. $3 \cdot 10^5$ parameters. For flow matching and continuous normalizing flows (CNFs), we use the same architecture based on a dense feed-forward neural net with skip connections using 8 residual blocks with each 128 neurons and *elu* activation. As input, we concatenate time t and θ_t . For Neural Spline Flow (Durkan et al., 2019) and FFJORD (Grathwohl et al., 2019), we adopt the released implementation by the authors.

Depending on the time per epoch for each method, we modify the number of epochs and steps per epoch to allow all methods to train for a similar amount of time, ensuring a sufficient window for convergence. For NSF, we train for 1 000 epochs, for flow matching for 2 000 epochs, and for FFJORD and CNF 100 epochs.

Flow matching with optimized hyperparameters For the experiments in section 5.2 and section 5.3, we adopt the hyperparameters and network architecture from Wildberger et al. (2023), which is based on a hyperparameter grid search. The hyperparameters for each task are listed in table 3. Otherwise, we follow the implementation as provided by the authors.

Table 3: Hyperparameters for SBI from Wildberger et al. (2023).

Task	Time α	Batch size	Learning rate	Residual blocks
LV	1	32	10^{-3}	[32, 64, 128, 256, 5×512 , 256, 128, 64, 32]
SLCP	-0.5	256	$5 \cdot 10^{-4}$	[32, 64, 128, 256, 5×512 , 256, 128, 64, 32]
SIR	4	256	$5 \cdot 10^{-4}$	[32, 64, 128, 256, 7×512 , 256, 128, 64, 32]
TM	4	64	$2 \cdot 10^{-4}$	[32, 64, 128, 256, 512, 3×1024 , 512, 128, 64, 32]

Analyzing the 1-step estimate We simulate the flow ODE from the sampling distribution at $t = 0$ until t^* (x -axis). Then, we compute the posterior in a single step by linearly extrapolating the flow, see eq. 6, to obtain the estimate $\hat{\theta}_1$. Results are shown in fig. 8.




Figure 8: (a) and (b): C2ST score and MMD for predictive samples $\hat{\theta}_1$. The x -axis shows from which we compute the predictive sample.

Analyzing step size We analyze the influence of the step size of the ODE solver on the quality of the posterior distribution as shown in fig. 9.




Figure 9: (a) and (b): C2ST score and MMD vs. step size during inference.

Additional results for maximum mean discrepancy For the evaluation in section 5.2, we show additional results for the maximum mean discrepancy (Gretton et al., 2012, MMD) in fig. 10.




Figure 10: Evaluation of SBI tasks using different variants of flow matching training. Lower MMD scores are better.

B.1 RECTIFIED FLOWS

The 1-step estimate $\hat{\theta}_1$ becomes more accurate and closer to the end point of the trajectory θ_1 as paths become straighter. Rectified flows (Liu et al., 2023) have been proposed to learn a coupling between two distributions by solving a nonlinear least squares optimization problem. Flows can be recursively rectified, leading to increasingly straighter paths. We have trained the k -th rectified flow up to $k = 3$ following Algorithm 1 from Liu et al. (2023) for the Lotka Volterra SBI task. Networks, optimizers and learning rates are the same as for the flow matching experiments. Results for the rectified flows are show in table 4. The C2ST score gets worse for the 2- and 3-Rectified flow. We also finetune with gradient-based control signals. A difference compared to the finetuning experiments in section 5.3 is that we train the network starting at $t \geq 0$, whereas we have used $t \geq 0.8$ before. As flows become straighter, the 1-step estimates should become more reliable. This is why we consider finetuning with the control signal on the entire trajectory in this experiment, instead of only focusing on the last part ($t \geq 0.8$). When adding the finetuning, the C2ST score becomes better for the 1-Rectified flow compared to the 2-Rectified flow, indicating that the 2-Rectified flow produces more reliable 1-step estimates. However, the results for the rectified flows are not as good as the flow matching setup with $t \geq 0.8$ which we have used in section 5.3.

	ODE solution	+ gradient-based controls
1-Rectified Flow	0.94	0.91
2-Rectified Flow	0.98	0.89
3-Rectified Flow	0.98	0.89

Table 4: C2ST score of the Lotka Volterra task for the k -th rectified flow.

C STRONG GRAVITATIONAL LENSING

We consider the following models:

- For modeling the lens we use an SIE model with 6 parameters: the Einstein radius θ_E , the ellipticities e_1 and e_2 and x_{center} and y_{center} . There is shear, for which we only consider γ_1 and γ_2 as free parameters.
- The source is modeled by a Sersic profile with free parameters being the amplitude, the half-light radius, the Sersic index n , the ellipticities e_1 and e_2 as well as the positions x_{center} and y_{center} .
- The lens light is modeled in the same way as the source, although when generating the mock data, we fix the position as well as ellipticities to be the same as the lens mass model. For training and inference, we infer positions for both lens mass and lens light model, so the model could produce different values for them. The MCMC methods use the same parameter for both lens light and lens mass position.

We list all priors in table 5, table 6 and table 7. We do not have priors on the ellipticities e_1 and e_2 directly, but we obtain them from priors on the position angle and axis ratio. Also, we obtain the shear parameters from γ_1 and γ_2 from ϕ_{ext} and γ_{ext} by converting them polar to cartesian coordinates. For SBI, we also include the two parameters ra_0 and dec_0 for the shear, which are always set to 0 when generating the training data sets, but in general our network could infer other values. Overall, there are 23 parameters for v_θ , which fully describe the simulation setup. However, in our dataset there are only 17 free parameters. The MCMC methods only infer the reduced set of parameters, making use of the dependencies between them.

Measurement instruments Observations have 160 times 160 pixels. The pixel size is 0.04 arc seconds. We use a Gaussian points spread function (PSF) with full width at half maximum (FWHM) of 0.3. There is Gaussian background noise with a root mean-squared values of 0.01 and an exposure time of 1000s.

Setup of MCMC-based methods We setup both baselines methods as follows:

1. Hamiltonian Monte Carlo: we use the No-U-Turn samples with a maximum tree depth of 10 and 5 000 warmup steps.
2. Affine-Invariant Ensemble Sampling: we use DEMove and StretchMove both with probability 0.5. There are 400 chains and we warm up for 20 000 steps before starting sampling.

Both methods are implemented in numpyro and optimized with JAX, so their runtimes are comparable with each other.

Network architectures and training

- Our flow network v_ϕ comprises a lightweight feature extraction network, represented by a CNN, which consists of 6 downsampling blocks with 1 layer each a 32 channels and kernel size 3. As postprocessing of the output, we apply GroupNorm, silu and an additional 2DConv layer with kernel size 3 and a single channel. We reshape the output and feed it through a final dense layer. The output of the feature extraction has the same dimensionality as the parameters θ .
- An additional dense feed-forward neural network receives the concatenated the time t , θ_t and extracted features as input. The feed-forward neural neural networks consists of 8 residual blocks with hidden dimension 128 and elu activation.
- The control network v_ϕ^C is represented by a small feed-forward neural network, consisting of 3 residual blocks with 64 hidden layers and 3 residual blocks with 32 hidden layers. We condition each block on the time via gated linear units and use a 16 dimensional time embedding.

For training, we use a batch size of 256 for the flow network v_ϕ . When training v_ϕ^C , we decrease the batch size to 16. We use the Adam optimizer with a learning rate of 10^{-4} and weight decay of 10^{-5} .

1026 Table 5: Priors for lens mass model parameters
1027

1028 Parameter	1029 Prior
1030 x_{center}	$\mathcal{U}(-0.2, 0.2)$
1031 y_{center}	$\mathcal{U}(-0.2, 0.2)$
1032 position angle ϕ	$\mathcal{U}(0, 180)$
1033 axis ratio q	$\mathcal{U}(0.25, 1)$
1034 external shear orientation ϕ_{ext}	$\mathcal{U}(0, 180)$
1035 external shear strength γ_{ext}	$\mathcal{U}(0, 0.1)$
1036 Einstein radius θ_E	$\mathcal{U}(0.5, 2.0)$

Table 6: Priors for the source light

Parameter	Prior
amplitude	$\mathcal{U}(5.0, 10.0)$
half-light radius	$\mathcal{U}(0.5, 2.0)$
Sersic index n	$\mathcal{U}(1.5, 4.0)$
position angle ϕ	$\mathcal{U}(0, 180)$
axis ratio q	$\mathcal{U}(0.25, 1)$
x_{center}	$\mathcal{U}(-0.2, 0.2)$
y_{center}	$\mathcal{U}(-0.2, 0.2)$

1037 Table 7: Priors for the lens light
1038

1039 Parameter	1040 Prior
1041 amplitude	$\mathcal{U}(5.0, 10.0)$
1042 half-light radius	$\mathcal{U}(0.5, 2.0)$
1043 Sersic index n	$\mathcal{U}(1.5, 4.0)$

1044
1045 Training v_ϕ was done on a single NVIDIA Ampere A100 GPU for ca. 45 hours. We trained v_ϕ^C for
1046 an additional 24 hours. A lot of the training time was spent on running evaluation metrics, so it can
1047 be substantially improved.
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079

C.1 DIFFUSION POSTERIOR SAMPLING (DPS)

We setup diffusion posterior sampling Chung et al. (2023a) as an additional baseline. The training dataset is the same as in 6, however since the diffusion model is unconditional, we drop any conditioning information.

Network architecture and training The neural network architecture is a multilayer perceptron MLP with 8 residual blocks and 128 neurons each. The activation function is *elu*. As optimizer, we use Adam with weight decay (10^{-5}). We train for 2000 epochs and for each epoch we sample 1000 batches from the dataset using a batch size of 4. We train the network as a denoising diffusion probabilistic model (DDPM) following Ho et al. (2020).

Unconditional generation Below, in figure 11, we visualize three samples generated by unconditionally sampling from the model. The observations are created using the lensing simulation code with the generated samples as input.




Figure 11: Visualization of unconditionally generated lensing systems.

Inference We directly follow Chung et al. (2023a) Algorithm 1 for inference, where the measurement operator \mathcal{A} is replaced by the lensing simulation code. The step size in the algorithm is defined via a hyperparameter ζ , which needs to be finetunes depending on the problem. We empirically test different values for ζ to find an optimal choice. Our results are shown in table 8. In this evaluation, we only model a smaller number of systems ($n = 25$).

ζ	0.0	0.0005	0.001	0.005	0.01	0.05
Avg. χ^2	28.15	16.20	9.98	10.07	12.98	12.64
Min. χ^2	15.14	3.84	3.07	1.58	1.40	1.53

Table 8: Evaluation of DPS and choosing ζ .

C.2 LOSS-GUIDED DIFFUSION

We consider loss-guided diffusion (LGD) with a Monte Carlo-based estimate of the guidance term (Song et al., 2023, LGD-MC). LGD-MC can be seen as an extension of DPS, which uses m points for estimating the guidance term, whereas DPS only uses a single point. We have evaluated LGD-MC using a different number of points m using 100 steps for each sample. Because multiple points are used for the calculation of the guidance term at each step, the number of simulator calls grows by a factor of m . Results are shown in table 9 below. Similar to the DPS evaluation, we only consider a smaller subset of systems ($n = 25$). Interestingly, even though LGD can be seen as an extension to DPS, it performs worse. Ther performance of DPS critically depends on the hyperparameter ζ that corresponds to the step size and needs to be finetunes. LGD does not have this hyperparameter.

C.3 TWISTED DIFFUSION SAMPLER

Twisted diffusion sampler (TDS) is a sequential Monte Carlo algorithm for asymptotically exact conditional sampling from diffusion models that has been proposed by Wu et al. (2023). We evaluate

1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187

m	5
Avg. χ^2	21.62
Min. χ^2	2.52
Modeling time	$\sim 1600s$

Table 9: Evaluation of LGD-MC for number of points m .

TDS using a different number of particles K with the unconditional diffusion model from section C.1. We follow Algorithm 1 from Wu et al. (2023) using 100 steps ($T = 100$). In the paper, the algorithm is described for a variance exploding (VE) noise schedule. We adjust the algorithm for the variance preserving (VP) as described in Wu et al. (2023) Appendix A. We give results below in table 10. Similar to the DPS evaluation, we only consider a smaller subset of systems ($n = 25$).

k	100
Avg. χ^2	20.94
Min. χ^2	2.47
Modeling time	$\sim 210s$

Table 10: Evaluation of TDS for number of particles.

C.4 SIMULATION-BASED CALIBRATION

We use simulation-based calibration (Talts et al., 2018, SBC) as an additional evaluation method to assess the correctness of the posterior distributions. We adopt the SBC implementation from the Python package *sbi*. Below, in fig. 12, we show histograms for 8 parameters based on $n = 1000$ lens systems with $L = 1000$ posterior samples each.

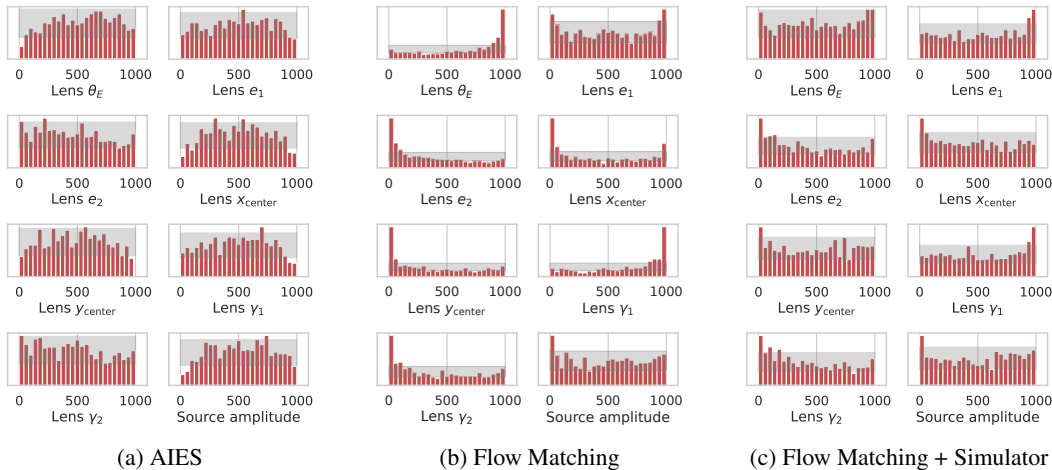


Figure 12: Simulation-based calibration histograms for different inference methods.

C.5 TESTS OF ACCURACY WITH RANDOM POINTS

We have included an additional evaluation using sampling-based accuracy testing of posterior estimators (Lemos et al., 2023, TARP), see figure 13. We included HMC initialized with the ground truth values as a reference, which shows perfect coverage. If not initialized with the ground truth parameters, HMC and AIES produce biased samples. Flow matching more closely covers the posterior and shows visible improvements when including simulator feedback.

1188
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241

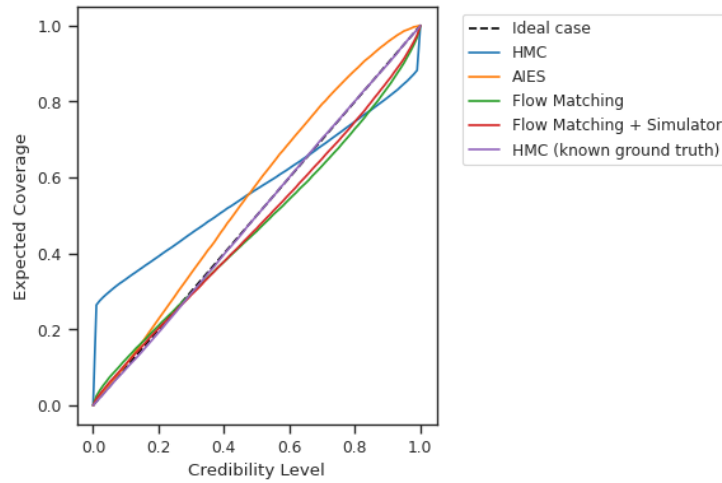


Figure 13: Evaluation of posterior coverage using TARP based on $n = 1000$ lens systems with $L = 1000$ posterior samples each.

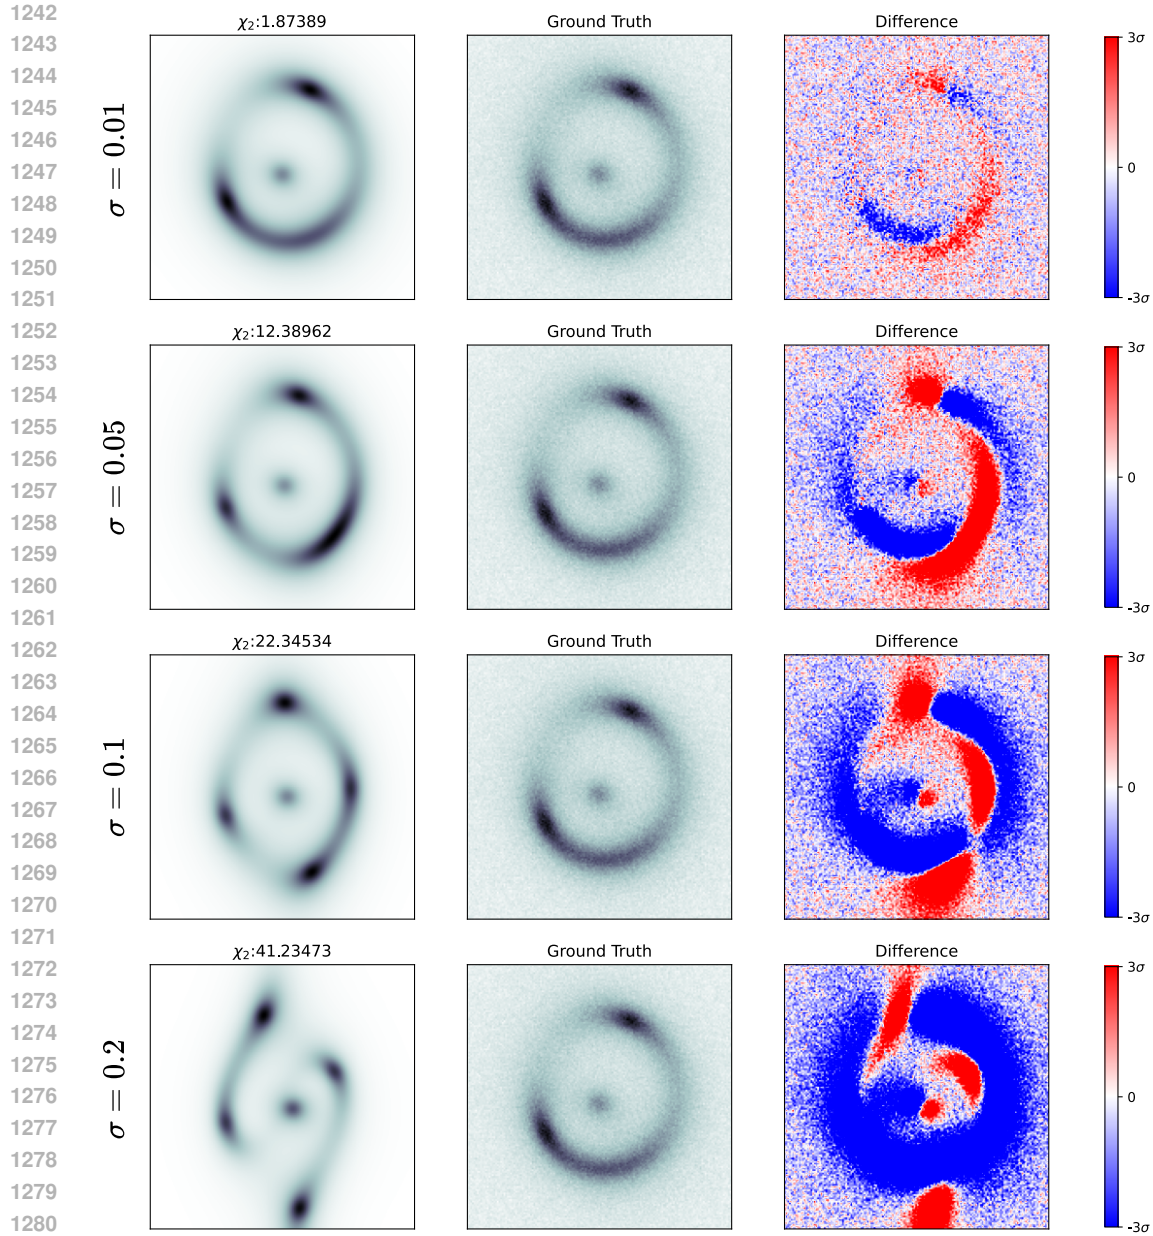


Figure 14: We show how a small noise σ affects the simulated observation. We add a normal Gaussian with mean 0 and standard deviation σ to a lens system’s ground truth parameters x . Then, we plot the simulated observation based on the noised parameters and show the residuals.

D POSTERIORIS AND RECONSTRUCTIONS FOR LENS MODELING

We show how small perturbations in the lens system’s parameters affect the simulated observation in figure 14. We show extended plots of the posteriors in fig. 15 for lens system 1 and fig. 16 for lens system 6. Additionally, we show reconstructions based on flow matching with and without simulator feedback of lens systems 1 to 6 in fig. 17

1296
 1297
 1298
 1299
 1300
 1301
 1302
 1303
 1304
 1305
 1306
 1307
 1308
 1309
 1310
 1311
 1312
 1313
 1314
 1315
 1316
 1317
 1318
 1319
 1320
 1321
 1322
 1323
 1324
 1325
 1326
 1327
 1328
 1329
 1330
 1331
 1332
 1333
 1334
 1335
 1336
 1337
 1338
 1339
 1340
 1341
 1342
 1343
 1344
 1345
 1346
 1347
 1348
 1349

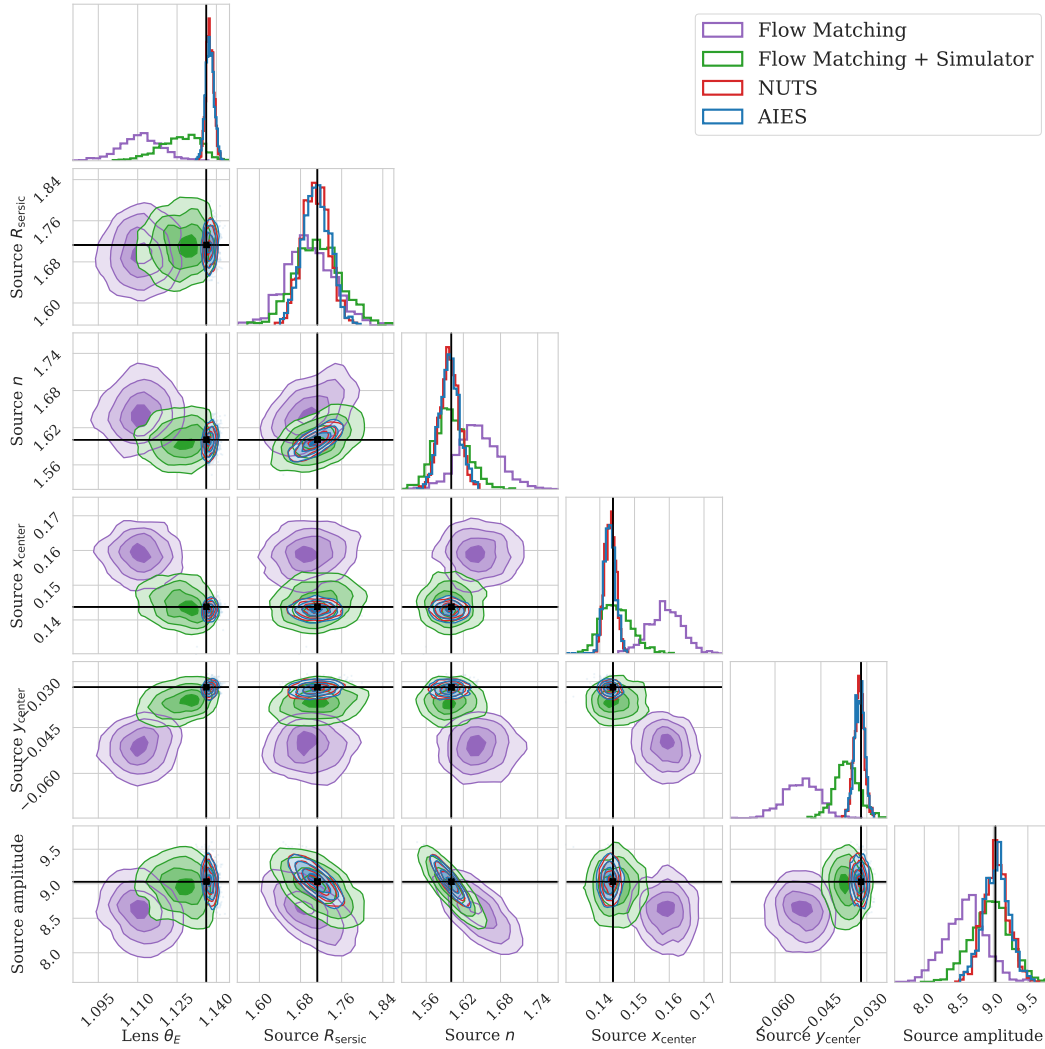


Figure 15: Posterior plot for system 1.

1350
 1351
 1352
 1353
 1354
 1355
 1356
 1357
 1358
 1359
 1360
 1361
 1362
 1363
 1364
 1365
 1366
 1367
 1368
 1369
 1370
 1371
 1372
 1373
 1374
 1375
 1376
 1377
 1378
 1379
 1380
 1381
 1382
 1383
 1384
 1385
 1386
 1387
 1388
 1389
 1390
 1391
 1392
 1393
 1394
 1395
 1396
 1397
 1398
 1399
 1400
 1401
 1402
 1403

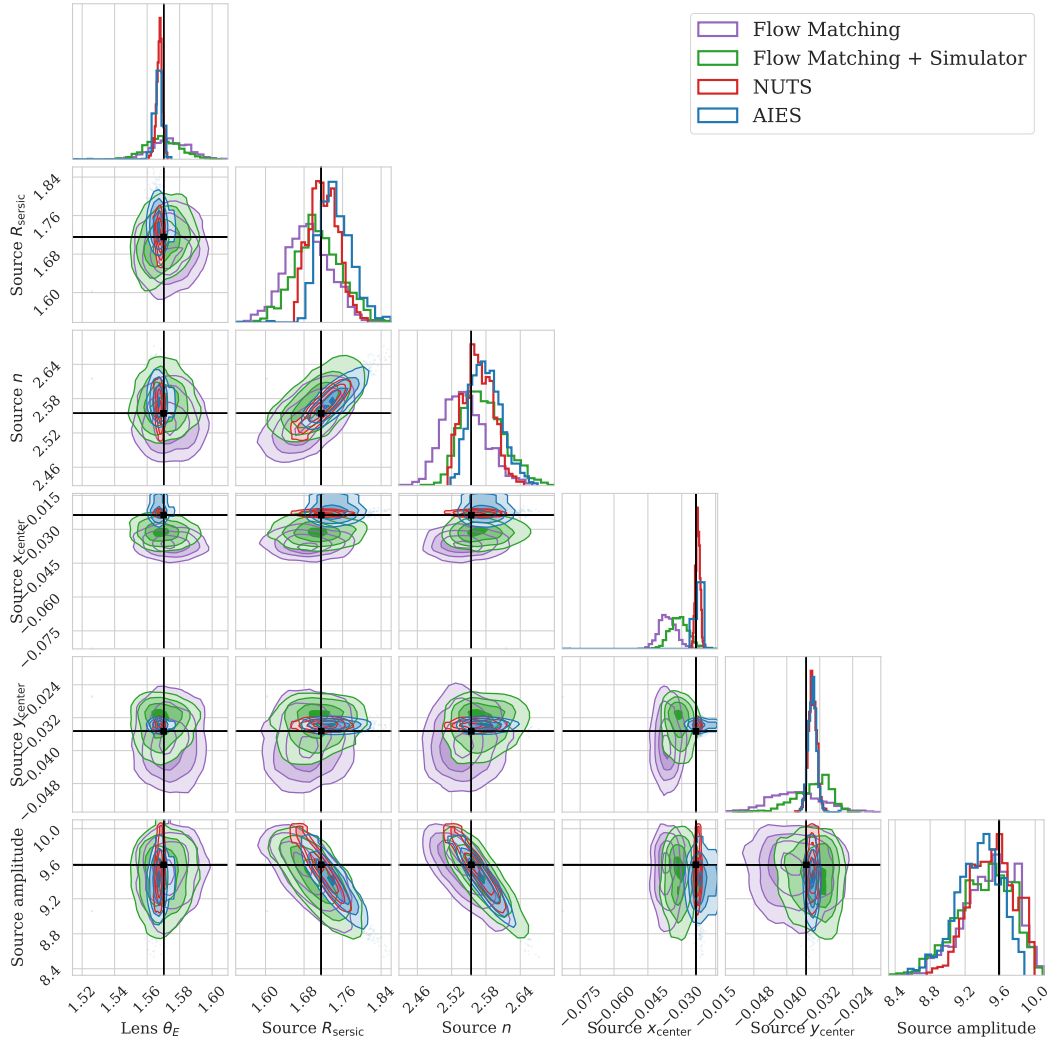


Figure 16: Posterior plot for system 6.

1404
 1405
 1406
 1407
 1408
 1409
 1410
 1411
 1412
 1413
 1414
 1415
 1416
 1417
 1418
 1419
 1420
 1421
 1422
 1423
 1424
 1425
 1426
 1427
 1428
 1429
 1430
 1431
 1432
 1433
 1434
 1435
 1436
 1437
 1438
 1439
 1440
 1441
 1442
 1443
 1444
 1445
 1446
 1447
 1448
 1449
 1450
 1451
 1452
 1453
 1454
 1455
 1456
 1457




Figure 17: Modeling of different lens systems: system 1 (top) to system 6 (bottom).