

# Preconditioned Neural Posterior Estimation for Likelihood-free Inference

Anonymous authors

Paper under double-blind review

## Abstract

Simulation based inference (SBI) methods enable the estimation of posterior distributions when the likelihood function is intractable, but where model simulation is feasible. Popular neural approaches to SBI are the neural posterior estimator (NPE) and its sequential version (SNPE). These methods can outperform statistical SBI approaches such as approximate Bayesian computation (ABC), particularly for relatively small numbers of model simulations. However, we show in this paper that the NPE methods are not guaranteed to be highly accurate, even on problems with low dimension. In such settings the posterior cannot be accurately trained over the prior predictive space, and even the sequential extension remains sub-optimal. To overcome this, we propose preconditioned NPE (PNPE) and its sequential version (PSNPE), which uses a short run of ABC to effectively eliminate regions of parameter space that produce large discrepancy between simulations and data and allow the posterior emulator to be more accurately trained. We present comprehensive empirical evidence that this melding of neural and statistical SBI methods improves performance over a range of examples including a motivating example involving a complex agent-based models applied to real tumour growth data.

## 1 Introduction

Computational models, frequently termed as simulators, are typically governed by stochastic processes. When provided with a set of parameter values, these simulators output synthetic data that inherently capture the stochastic nature of the simulated phenomena. However, a substantial challenge arises in performing posterior inference for the parameters of these simulators, as the corresponding likelihood function is often intractable. An example is agent-based modelling of tumour growth (e.g. Jenner et al. (2020); Aylett-Bullock et al. (2021)), where cell proliferation, movement and invasion are governed by probabilistic rules, which depend on key biological parameters. Consequently, standard statistical inference methods that rely on a closed form expression for the likelihood function are inapplicable in this scenario.

To address this issue, simulation based inference (SBI) methods have been developed, which approximate the posterior based only on simulations from the model. The most thoroughly examined SBI method in the statistical literature is approximate Bayesian computation (ABC) (Sisson et al., 2018). Advancements in deep neural networks have led to the emergence of neural SBI methods (Cranmer et al., 2020). Their widespread application spans various fields, including biology (Wehenkel et al., 2023), neuroscience (Fengler et al., 2021; West et al., 2021), and astronomy (Mishra-Sharma, 2022; Dax et al., 2021).

Statistical SBI methods, such as ABC methods, are well-developed and boasts strong theoretical guarantees of convergence to the true posterior (Beaumont et al., 2009; Blum, 2010; Biau et al., 2015; Lintusaari et al., 2017; Sisson et al., 2018; Beaumont, 2019). ABC approaches compare observed to simulated data using a discrepancy function and prefer parameter values that generate discrepancies below a pre-defined threshold,  $\epsilon$ . However, often a small value of  $\epsilon$  is required to obtain accurate posterior approximations, which can significantly increase the number of model simulations required, and hence the computational cost (Csilléry et al., 2010; Turner & Van Zandt, 2012). Advanced ABC samplers such as adaptive sequential Monte Carlo ABC (SMC ABC, e.g. Sisson et al. (2007); Drovandi & Pettitt (2011)) have been developed to mitigate this

issue. However, the efficiency of ABC methods may decrease as  $\epsilon$  becomes smaller, requiring a significantly larger number of simulated datasets. Enhancing the efficiency of ABC algorithms remains an active area of research (Lintusaari et al., 2017; Sisson et al., 2018).

With the rapid advances in machine learning methods, more efficient approaches based on neural networks have been developed (Papamakarios et al., 2017; 2021). A popular approach is neural posterior estimator (NPE) and its sequential version, the sequential NPE (SNPE) (Papamakarios & Murray, 2016; Lueckmann et al., 2017; Greenberg et al., 2019). These methods use a set of training pairs of parameter values and simulated datasets to fit a neural conditional density estimator (NCDE), such as a conditional normalising flow (Rezende & Mohamed, 2015; Winkler et al., 2019; Durkan et al., 2019; Dolatabadi et al., 2020), to approximate the posterior. NPE uses parameter values drawn from the prior, whilst SNPE uses parameter values drawn from previous NPE approximations for a given number of rounds. The idea of the sequential approach is that a more accurate emulator of the posterior can be achieved when more (parameter value, simulated data) training pairs are generated in higher density regions of the posterior. Lueckmann et al. (2021) have shown that NPE can outperform ABC, particularly for a relatively small number of model simulations.

However, in real-world problems, there may be little known about the parameters *a priori*, so that a vague prior may be employed. For example, a uniform distribution with a wide constraint range may be used as a prior distribution. For some problems in this setting we find that it is difficult to construct an accurate NCDE across a wide parameter space, which leads to NPE producing an inaccurate posterior approximation. We find that even SNPE may not be able to recover from such an initially deficient approximation, even with a relatively large number of rounds, and hence model simulations. One way to mitigate this kind of unstable NCDE training is to clip the extreme simulated datasets (Shih et al., 2023; de Santi et al., 2023). However, this approach is ad-hoc and it is not clear how much clipping is required for a given problem and may require extensive experimentation, and each level of clipping requires refitting of the NCDE.

This paper contains three key contributions. Firstly, we explore several examples where NPE methods fail to produce highly accurate posterior approximations, even in relatively low dimensional problems. Our second contribution is the development of preconditioned NPE (PNPE), and its sequential extension (PSNPE), which combines the strengths of statistical and ML approaches to SBI. The preconditioning step involves applying an ABC algorithm for efficiently discarding parts of the parameter space that lead to large discrepancies, which then subsequently permits NPE methods to perform well. In a sense, our preconditioning step acts as a principled clipping method. Our third contribution shows via an extensive empirical study that our preconditioned NPE approaches outperform NPE approaches when the latter performs sub-optimally, and is competitive when it performs well. Our motivating example involves a complex agent-based model applied to real tumour growth data.

## 2 Simulation-based Inference

Consider a simulator that takes parameters  $\theta \in \mathbb{R}^d$  where  $d$  is the number of parameters and generates a simulated dataset  $x \in \mathbb{R}^D$  where  $D$  is the dimension of the data, but its density  $p(x|\theta)$  is intractable. The objective of SBI is to accurately estimate the posterior density of  $\theta$  conditional on the observed dataset  $x_o \in \mathbb{R}^D$  based only on simulating data from the model and not requiring evaluation of the intractable likelihood,  $p(x_o|\theta)$ . Two popular SBI methods are ABC and NPE, which are summarised below.

### 2.1 Approximate Bayesian computation

Statistical SBI (Sisson et al., 2018), such as the ABC rejection algorithm, is based on Monte Carlo rejection sampling. That is, it keeps only the parameter values simulated from the prior that generate simulated data  $x$  such that  $\rho(x, x_o) < \epsilon$ , where  $\rho(x, x_o)$  is a user-defined discrepancy function between the simulated and observed data, and  $\epsilon$  is a user-defined threshold often referred to as the ABC tolerance.

SMC ABC algorithms (e.g. Sisson et al. (2007); Drovandi & Pettitt (2011)) aim to be more efficient by sampling a sequence of ABC posteriors with decreasing  $\epsilon$ 's, updating the importance distribution at each iteration. More specifically, SMC ABC algorithms define a sequence of non-increasing ABC thresholds

$\epsilon_1 \geq \epsilon_2 \geq \dots \geq \epsilon_T$ , such that

$$p_{\epsilon_t}(\theta|x_o) \propto p(\theta) \int_{\mathbb{R}^D} \mathbb{I}(\rho(x_o, x) < \epsilon_t) p(x|\theta) dx, \quad \text{for } t = 1, \dots, T. \quad (1)$$

Here,  $\epsilon_T = \epsilon$  represents the target ABC posterior.

In many real-world applications,  $x, x_o \in \mathbb{R}^D$  are considered high-dimensional data, necessitating a mapping to a lower-dimensional space for computational efficiency. This is typically done using summary statistics  $S(\cdot)$ . If summary statistics are required, we use  $S(x)$  and  $S(x_o)$  instead of the full datasets  $x$  and  $x_o$ . The choice of appropriate summary statistics is a subject of ongoing research and is discussed extensively in the literature (see Sisson et al. (2018)).

However, even sophisticated ABC algorithms can require a significant number of model simulations to achieve a suitably small value of  $\epsilon$  (Biau et al., 2015; Csilléry et al., 2010; Beaumont et al., 2009; Blum, 2010).

## 2.2 Neural posterior estimation

NPE uses  $N$  training pairs of simulator parameter values and simulated datasets,  $\{\theta_i, x_i\}_{i=1}^N$ , to estimate the posterior distribution  $p(\theta|x)$  (Papamakarios et al., 2021). Once the NPE is trained on the simulated datasets, the posterior distribution  $p(\theta|x_o)$  can be computed by inputting the observed dataset  $x_o$ . A conditional neural density estimator  $q_{F(x, \psi)}(\theta)$ , utilizing a neural network  $F$  and its adjustable network weights  $\psi$ , is often used as an NPE. In order to train  $q_{F(x, \psi)}(\theta)$ , the following loss is minimized:

A neural density estimator  $q_{F(x, \phi)}(\theta)$  used neural network  $F$  with adjustable network weights  $\phi$  is used for NPE that trained by minimize the loss:

$$\psi^* = \arg \min_{\psi} - \sum_{i=1}^N \log q_{F(x_i, \psi)}(\theta_i), \quad (2)$$

over network weights  $\psi$ . For a sufficiently expressive  $q_F$ ,  $q_{F(x, \psi)}(\theta) = p(\theta|x)$ .

SNPE aims to improve the accuracy of the approximate posterior for a particular observed dataset  $x_o$  iteratively by sampling parameter values from a previous NPE approximation for a given number of rounds. The current NPE approximation is treated as a proposal distribution  $\tilde{p}(\theta)$  for the next round. However, training  $q_F$  using parameter values drawn from  $\tilde{p}(\theta)$  will not converge to the true posterior distribution, but rather to

$$\tilde{p}(\theta|x_o) \propto p(\theta|x_o) \frac{\tilde{p}(\theta)}{p(\theta)}. \quad (3)$$

Many approaches have been developed to overcome this limitation, such as Papamakarios & Murray (2016); Lueckmann et al. (2017); Greenberg et al. (2019). Among all these approaches, we use the automatic posterior transformation (APT, also known as SNPE-C), as proposed by Greenberg et al. (2019), which has been reported to significantly outperform the others (Lueckmann et al., 2021). However, APT can suffer from a ‘leakage’ issue, in which case we use truncated sequential neural posterior estimation (TSNPE) (Deistler et al., 2022a), which aims to overcome the leakage problem. For simplicity, we refer to our specific implementation as SNPE hereafter, but we note that other implementations of SNPE can be used with our method.

### 2.2.1 Illustrative Example

We consider a sparse vector autoregressive (SVAR) model that has been considered previously in the SBI literature (Thomas et al., 2020; Drovandi et al., 2023). The SVAR model is given by:

$$y_t = X y_{t-1} + \xi_t, \quad (4)$$

where  $y_t \in \mathbb{R}^k$  represents the  $k$ -dimensional observation of the time series at time  $t$ ,  $X \in \mathbb{R}^{k \times k}$  is the transition matrix, and  $\xi_t \sim \mathcal{N}(0, \sigma^2 \mathbb{I})$  is a  $k$ -dimensional noise vector with  $\sigma^2$  being the noise parameter.

The model considers a sparse transition matrix  $X$  where the only off-diagonal entries that are non-zero must satisfy the following conditions: (1) they are strictly off-diagonal entries, meaning  $i \neq j$  for  $1 \leq i, j \leq k$ ; (2) if variable  $i$  is coupled with variable  $j$ , then  $X_{i,j} \neq 0$  and  $X_{j,i} \neq 0$  (note that  $X_{i,j}$  is not necessarily equal to  $X_{j,i}$ ). To ensure the stability of the SVAR, the diagonal elements of  $X$  are set to -0.1. The parameter space of SVAR can easily scale to higher dimensions by increasing  $k$ . In this study, the model parameters  $\theta \in \mathbb{R}^{k+1}$  are the non-zero off-diagonal entries of  $X$  and its variance and we consider  $k = 6$ , which leads to 7 parameters. This choice is based on the assumption that if SNPE does not produce highly accurate approximations in this low-dimensional case, it is unlikely to be accurate in a higher dimensional parameter space. We generate an observed dataset of length  $T = 1000$  using the true parameter value  $\theta = (0.579, -0.143, 0.836, 0.745, -0.660, -0.254, 0.1)$ . We use summary statistics to reduce the dimension of the data. Following Thomas et al. (2020); Drovandi et al. (2023), we use the lag 1 autocovariance  $\frac{1}{T} \sum_{t=2}^T y_t^i y_{t-1}^j$  as the summary statistics, where  $y_t^i$  is the  $t$ th observation of the  $i$ th time series. We use the sample standard deviation of the  $k$  time series to inform  $\sigma$ . Thus there is a single summary statistic that is intended to be informative about each parameter.

We employ a uniform distribution as the prior, constrained between -1 and 1 for the  $k$  parameters and between 0 and 1 for  $\sigma$ . We find that extreme values of the summary statistics can be produced by parameter values away from the true parameter value. To stabilize the training, we clip simulated datasets with summary statistic outliers (any simulated values greater than 10, around 3% of training datasets). Note that some experimentation was required to obtain a clipping value that led to reasonable results for SNPE.

For illustrative purposes, we run three rounds of SNPE and compare the results with those from BSL, considering BSL results as the gold standard for this example. Ideally, we would expect performance to improve when increasing the number of SNPE rounds. However, as shown in Figures 1 and 2, even with datasets clipped for every round, SNPE does not improve the accuracy of the estimates as the number of rounds increases.

### 3 Method

We find that NPE can perform sub-optimally when the prior predictive distribution of the data is complex and has significant variability. When this occurs, the NCDE may not be sufficiently accurate, especially in regions of high posterior support. The sequential version of NPE was originally designed to overcome this issue, however if the initial NPE approximation is not substantially better than the prior, then subsequent rounds of SNPE may suffer from the same issue.

Other approaches to overcome this issue may be to increase the training sample size or to try different configurations of the neural network, but both of these may increase the computational cost substantially and may not address the issue. Instead, we propose the preconditioned NPE (PNPE) method, and its sequential extension below, in order to make NPE methods more reliable.

#### 3.1 Preconditioned NPE

For a vague prior distribution  $p(\theta)$ , parameter values drawn from it might be very far from the true posterior. We suggest using a short run of ABC to refine those parameter values so that they are closer to the true posterior  $p(\theta|x_o)$ , i.e.,

$$p_\epsilon(\theta|x_o) \propto p(\theta) \int_{\mathbb{R}^D} p(x|\theta) \mathbb{I}(\rho(x_o, x) < \epsilon) dx, \quad (5)$$

where  $\mathbb{I}(\cdot)$  is the indicator function and  $\epsilon$  can be chosen considerably larger than what might typically be used in an ABC algorithm. Then we fit a density estimator to those parameter samples. The key idea is to use an efficient ABC algorithm to quickly discard poor regions of the parameter space that generate unusual datasets relative to the observed data, which provides better quality training datasets for NPE. We note that any ABC algorithm could be employed here, but we use the SMC ABC algorithm of Drovandi & Pettitt (2011) in this paper (see Appendix A for a full description of this method). The SMC ABC algorithm generates  $n$  samples from a sequence of ABC posteriors based on decreasing ABC thresholds,  $\epsilon_1 > \dots > \epsilon_T$ , where  $\epsilon_T = \epsilon$  is the target ABC threshold. The sequence of tolerances is determined adaptively, by, at






Figure 1: Comparison of marginal posterior distributions between BSL (orange), NPE (dashed pink) and SNPE (red), with black dashed lines representing the true values. The SNPE results are based on three rounds.

each iteration of SMC, discarding a proportion of the samples,  $a \cdot n$ , with the highest discrepancy, where  $a$  is a tuning parameter. Then, the population of samples is rejuvenated through a resampling and move step. During the move step, a Markov chain Monte Carlo (MCMC) ABC kernel is employed to maintain the distribution of particles based on the current value of the tolerance. The number of MCMC steps  $R_t$  to apply to each particle is determined adaptively based on the overall MCMC acceptance rate, that is  $R_t = \left\lceil \frac{\log(c)}{\log(1-p_t^{\text{acc}})} \right\rceil$ , where  $p_t^{\text{acc}}$  is the estimated MCMC acceptance probability at the  $t$ th SMC iteration. A natural stopping rule for the algorithm is when the MCMC acceptance rate becomes intolerably small.

Based on  $n$  parameter samples from the ABC posterior, we fit an unconditional normalising flow  $q_G$  (note that other density estimators could be used). Then we can use  $q_G$  as the initial importance distribution for the (S)NPE process. We call this melding of ABC and (S)NPE as the preconditioned (S)NPE method. The method is summarised in Algorithm 1.

---

**Algorithm 1** Precondition SNPE

---

- 1: Choose preconditioning ABC algorithm and SNPE implementation
  - 2: Obtain  $\{\theta_i^*\}_{i=1}^n$  from the preconditioning ABC algorithm
  - 3: Set  $\phi^* \leftarrow \arg \min_{\phi} \sum_{i=1}^n -\log q_{G(\phi)}(\theta_i^*)$
  - 4: Perform SNPE using initial importance distribution  $q_{G(\phi^*)}(\theta)$
-




Figure 2: Comparison of posterior predictive distribution of the summary statistics of observation datasets between BSL (orange), NPE (dashed pink) and PNPE (red), with black dashed lines representing the true values. The SNPE results are based on three rounds.

If we obtain a well-trained unconditional normalizing flow, this unconditional normalizing flow can act as the initial importance distribution, i.e.,  $\tilde{p}(\theta)$  in Equation 3, and to draw samples for training the NPE. Following Papamakarios & Murray (2016), it is noted that given an expressive enough conditional normalizing flow, NPE converges to the true posterior  $p(\theta|x_o)$  as  $N \rightarrow \infty$ , with an appropriate importance re-weight if  $\tilde{p}(\theta) \neq p(\theta)$ .

Choosing a suitable value of  $\epsilon$  for our method requires some thought. A smaller value of  $\epsilon$  will focus in on more promising regions of the parameter space, but will increase the computational time of the preconditioning step. A larger value of  $\epsilon$  will lead to a fast preconditioning step, but may not eliminate enough of the poor parts of the space to improve the training of the NCDE. In this paper we use an MCMC acceptance rate of 10% (unless otherwise specified) as the stopping criteria for the SMC ABC algorithm in the preconditioning step. For our examples we find that this choice is effective at balancing the aforementioned objectives. We note that other choices are possible.

Furthermore, once these poor simulations have been removed, we find NPE to be more effective than ABC, since ABC requires an exponentially increasing number of simulations to drive  $\epsilon$  to 0. To avoid the scaling problem of ABC, the preconditioning step only takes a short run of ABC, and thus we are not interested in driving  $\epsilon$  to 0.

### 3.2 Computational cost

We now consider computational cost for P(S)NPE and compare it with SNPE. The preconditioning step can be considered as the initial round of NPE where the total number of simulated datasets generated during SMC ABC is denoted as  $n_{\text{ABC}}$ . Hence, it is worth noting that P(S)NPE, like SNPE, is not amortized since it requires running an ABC algorithm for each observation datasets  $x_o$ .

Furthermore, for complex real-world problems, the simulation time may depend on the parameter values, and parameter values with very low posterior support can produce substantially longer simulation times. For such problems, it is important from a computational perspective to quickly eliminate such regions from the parameter space, as is the motivation of the preconditioning ABC algorithm. Thus, for problems where SNPE does not perform well, we find PSNPE to be substantially more computationally efficient in terms of compute time.

### 3.3 Illustrative Example Revisited

We apply PNPE to the illustrative example shown in Section 2.2.1. For the preconditioning step, we use the adaptive SMC ABC algorithm proposed by Drovandi & Pettitt (2011), with tuning parameters  $n = 1k$ ,  $a = 0.5$ , and  $c = 0.01$ . We employ an unconditional normalizing flow as the unconditional density estimator. For this, we use the state-of-the-art neural spline flow implemented in the `Pyro` package.

In order to make a fair comparison between P(S)NPE and SNPE, we use the same number of simulations as in the SMC ABC algorithm, denoted  $n_{ABC}$ , to train the initial round of NPE. We run the SMC ABC algorithm ten times to obtain the average number of simulations it requires, which is  $n_{ABC} = 54k$ . For illustrative purposes, we only run two rounds of SNPE and compare it with PNPE.

The estimated marginal posterior plots are displayed in Figure 3, where the black solid lines represent the true parameter values. It is evident that our PNPE method (in a single round) produces a substantially sharper approximation of the posterior compared to that of SNPE.

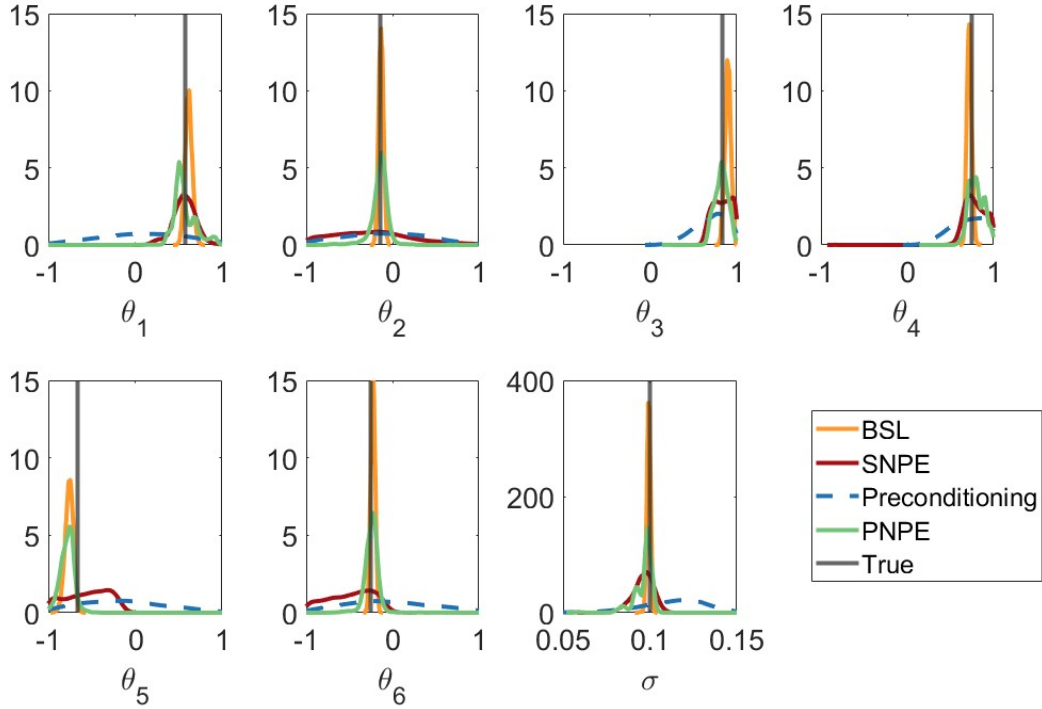


Figure 3: **Performance on SVAR model.** Comparison of marginal posterior distributions between BSL (orange), SNPE (red), preconditioning step (blue dash) and PNPE (green solid), with black dashed lines representing the true values.

It is evident that even a short run of the ABC algorithm gives a reasonable posterior approximation to train an unconditional normalizing flow as the unconditional density estimator, which then generates samples for

NPE training. Figures 3 and 4 show that the improved parameter posteriors leads to more accurate posterior predictive distributions of the summaries compared to SNPE. This indicates that with a good starting point, NPE can further improve accuracy.

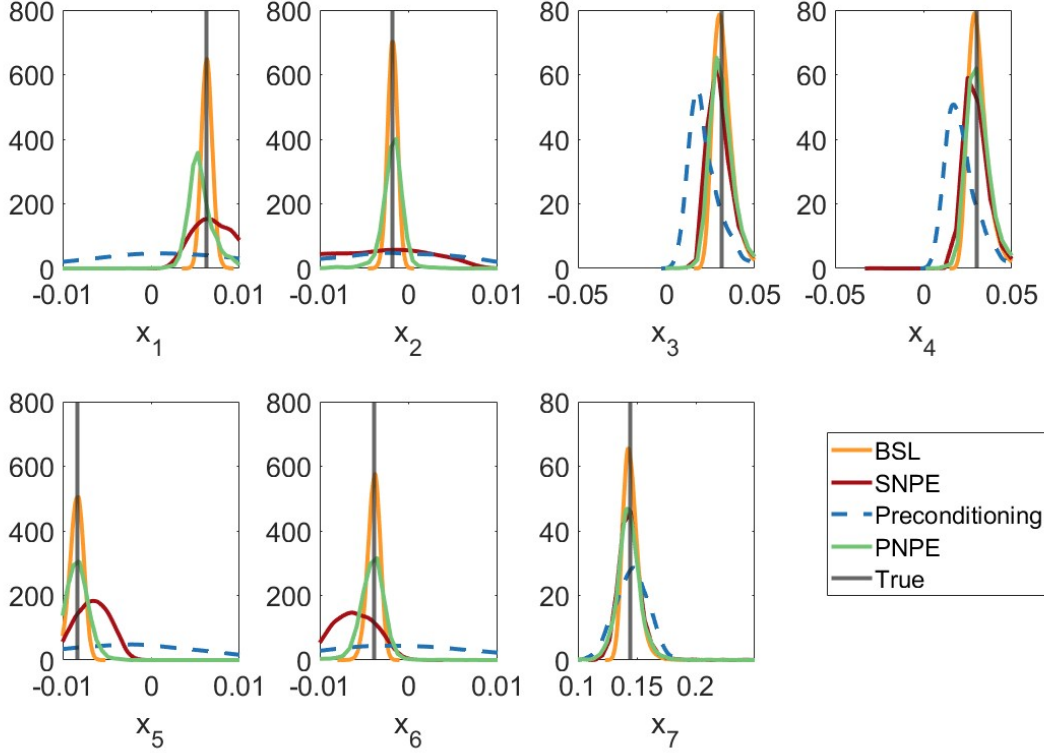


Figure 4: **Performance on SVAR model.** Comparison of posterior predictive distributions of the summary statistics of observation datasets between BSL (orange), SNPE (red), preconditioning step (blue dash) and PNPE (green solid), with black dashed lines representing the true values.

## 4 Further Experiments

We present two additional examples, including our motivating example, where SNPE, perhaps surprisingly, does not produce highly accurate posterior distributions. To fairly compare our method with vanilla SNPE and potentially TSNPE, we run the ABC algorithm 10 times and compute the average total number of simulations that the ABC algorithm requires. We then use this same number of simulations for the initial NPE in both SNPE and TSNPE. For all experiments, we utilize the adaptive SMC ABC algorithm proposed by Drovandi & Pettitt (2011) for the preconditioning part with tuning parameters  $n = 1k$ ,  $a = 0.5$ ,  $c = 0.01$ , and use an unconditional normalizing flows as the unconditional density estimator. For SNPE, we use 10k samples for each round of training.

We find that even with a single round of PNPE, there can be a significant improvement in performance. The code to reproduce the results has been included as supplementary material.

### 4.1 High-dimensional SVAR model

To investigate how our method scales to higher dimensional problems, we take the illustrative SVAR example from before and consider  $k = 20$ , which leads to 21 parameters. We detail the experimental settings in Section

B.1. To ensure a fair comparison, we run the SMC ABC algorithm ten times using a 10% acceptance rate as the stopping criterion and calculate the average number of simulations it takes. We then use the same number of simulations, approximately  $n_{\text{ABC}} \approx 45k$ , to train the initial NPE. Hence, the total number of simulations for both PNPE and SNPE is the same (55k in total). To stabilize the training, we apply the same clipping technique used in the previous low-dimensional case, which results in approximately 11% of the training samples being removed in the initial round of NPE and about 1% to 2% in the second rounds. Starting from the third round, SNPE is unable to sample any parameter values from the neural networks due to a severe leakage issue (Deistler et al., 2022a).

The estimated marginal posterior plots are displayed in Figure 5, where the black solid lines represent the true parameter values. It is evident that as the number of parameter dimensions increases, training the unconditional normalizing flows becomes more challenging. With well-trained unconditional normalizing flows, PNPE outperforms SNPE in high-dimensional cases (in this example, except for parameter  $\theta_3$ ).

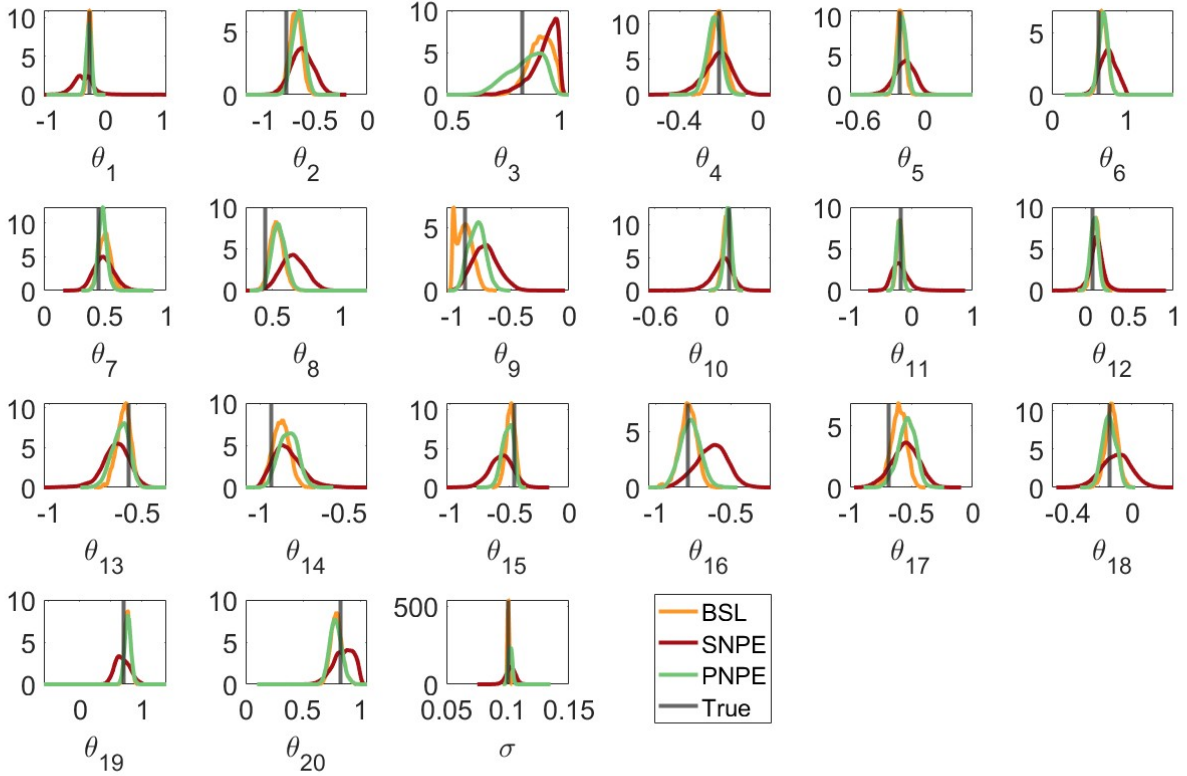


Figure 5: **Performance on SVAR model with 21 parameters.** Comparison of marginal posterior distributions between BSL (orange), SNPE (red) and PNPE (green), with black dashed lines representing the true values. The result of SNPE uses 2 rounds.

## 4.2 Biphasic Voronoi cell-based model

Finally, we consider a challenging real-world problem in cancer biology: calibrating the biphasic Voronoi cell-based model (BVCBM) (Wang et al., 2024) that models tumor growth. The model uses a parameter  $\tau$  to divide the tumor into two growth phases. Here, the term ‘growth phase’ refers to the different growth patterns of the tumor. There are four parameters that govern tumor growth during each phase, namely

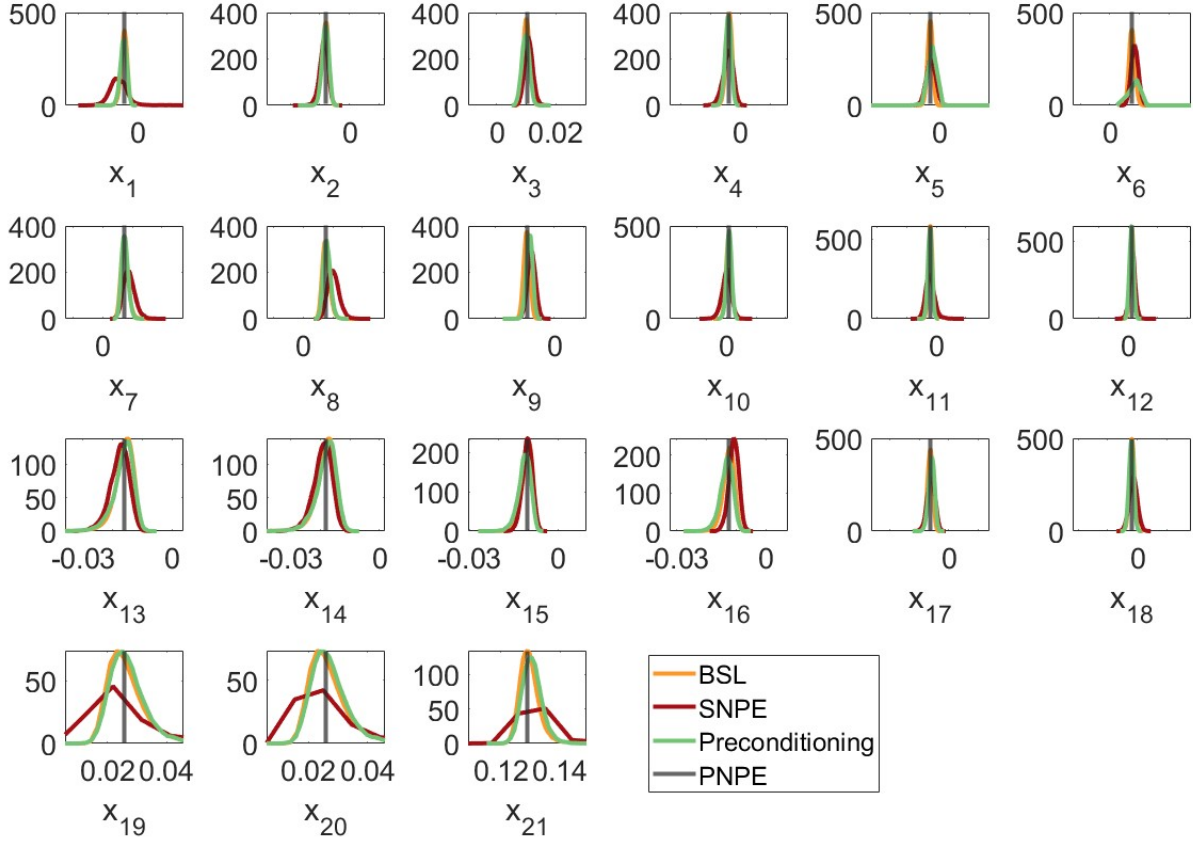


Figure 6: **Performance on SVAR model with 21 parameters.** Comparison of posterior predictive distributions of the summary statistics of observation datasets between BSL (orange), SNPE (red) and PNPE (green), with black dashed lines representing the true values. The result of SNPE uses 2 rounds.

$(p_0, p_{\text{psc}}, d_{\text{max}}, g_{\text{age}})$ , where  $p_0$  and  $p_{\text{psc}}$  are the probability of cell proliferation and invasion, respectively,  $d_{\text{max}}$  is the maximum distance between cell and nutrient, and  $g_{\text{age}}$  is the time taken for a cell to be able to divide. Thus there is nine parameters in total, four parameters each of two phases, and the parameter  $\tau$  at which the growth phase changes. In this paper, we calibrate to two real-world pancreatic cancer datasets Wade (2019), which describe tumor growth as time series data. The datasets span 26 and 32 days, respectively, with measurements taken each day. While the ground truth posteriors are unknown for those datasets, we compute posterior predictive distributions to assess if SNPE and PNPE can effectively calibrate the model to the data.

We employ vague prior distributions for all parameters. Specifically, we use a Uniform distribution constrained between 1 and 24 hours  $\times$  the number of days for  $g_{\text{age}}$  during both growth phases. Additionally, we use a Uniform distribution constrained between 2 and the number of days minus 1 for  $\tau$ . The prior distributions for the remaining parameters are detailed in Appendix B.2.

As reported by Wang et al. (2024), CPU times for model simulation range from 1.76 to 137.27 seconds per simulation when using samples from the prior distribution. This implies that 10k simulations for the first round of SNPE take approximately 2 hours. Consequently, the initial stages of SNPE are computationally expensive. In contrast, the ABC preconditioning step takes around 10-15 minutes. This is due to the fact



that the longer simulation times tend to also lead to large discrepancies with the observed data, and such samples are quickly rejected by ABC. For the ABC part, around 18k and 16k simulations are used for the 26-day and 32-day pancreatic cancer datasets, respectively.

We observed a leakage problem occurring in rounds 8 and 5 for the 26-day and 32-day datasets, respectively. At round 10, the acceptance rate of SNPE for the 26-day measurement dataset is above 50%, leading us to utilize rejection sampling for sample generation. While a 50% acceptance rate for rejection sampling is acceptable, it is less efficient compared to direct simulation from the trained conditional normalizing flows in the previous round. Hence, we employ TSNPE for this dataset. For the 32-day measurement dataset, only 0.000% of samples are accepted at round 8, making it computationally expensive. Consequently, we use TSNPE for this dataset.

To estimate the posterior predictive distributions, we sample 1k parameter values from (T)SNPE and PNPE, using them to simulate datasets. We then plot these data in the form of credible intervals. As a baseline, we show the prior predictive distributions in Appendix C. The top and bottom rows of Figure 7 displays the posterior predictive distribution for the 26-day and 32-day datasets, respectively (the same plots but on the log scale are shown in Appendix C). It is evident that both SNPE and TSNPE provide biased estimations, as the posterior predictive does not capture the observed data well. The posterior predictive distribution for the preconditioning step (middle column) shows that the ABC step can capture the data reasonably well since the observed data lie within the 90% posterior predictive interval. Our method (third column) provides a better fit, as the variance of the posterior predictive distribution is tighter than that of the preconditioning step and still captures the observed data. This demonstrates that even one round of PNPE can perform more accurate estimations based on our results.

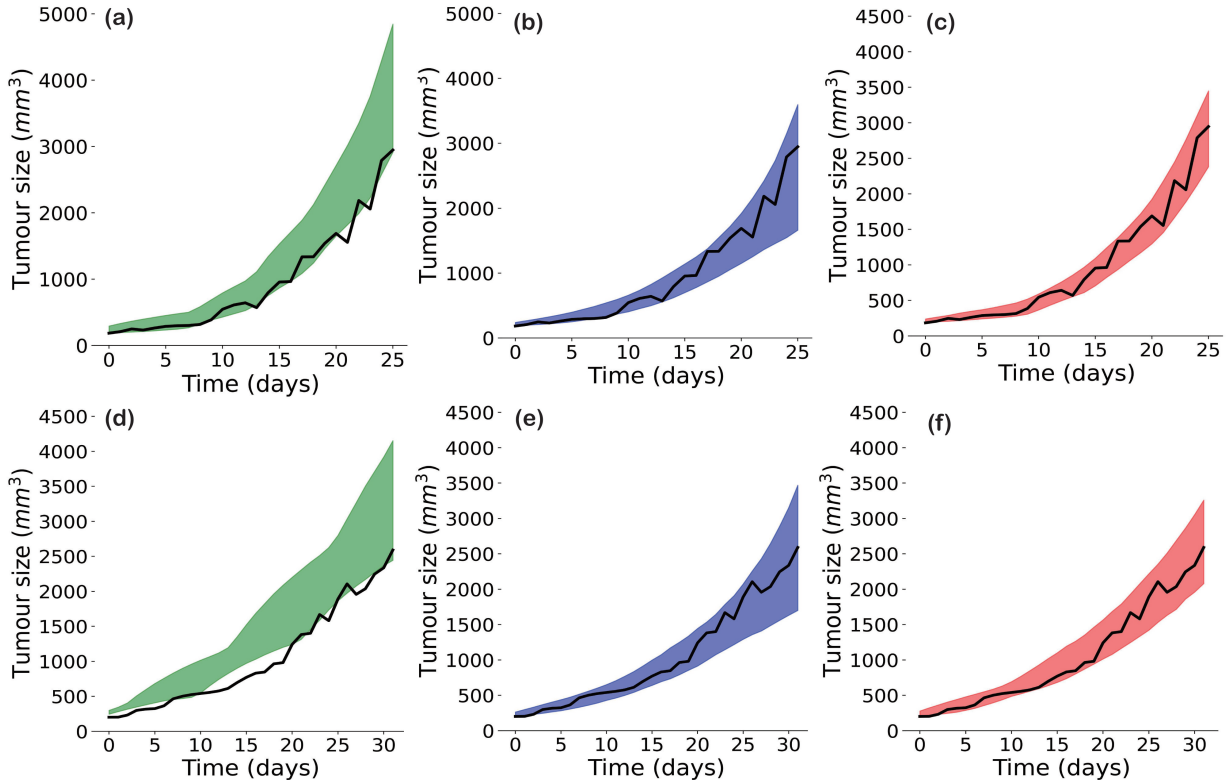


Figure 7: **Posterior predictive distributions for two pancreatic cancer datasets.** 90% Posterior predictive interval plots for SNPE (a), preconditioning step (b) and PNPE (c) for the 26-day dataset; 90% Posterior predictive interval plots for SNPE (d), preconditioning step (e) and PNPE (f) for the 32-day dataset.

## 5 Discussion

We present a neural SBI method that is both simple and easy to deploy, designed to enhance the accuracy of SNPE methods. Our method, termed preconditioned neural posterior estimation (PNPE) and its sequential version, PSNPE, employs an ABC algorithm for the initial step. This algorithm is used to efficiently filter out poor regions of the parameter space. Additionally, we use the ABC posterior samples to train an unconditional density estimator  $q_G$ , enabling  $q_G$  to serve as the initial proposal distribution for SNPE. The core concept is that an improved starting point can significantly enhance the accuracy of SNPE estimations. Indeed, we obtained very good results with PNPE. Lueckmann et al. (2017); Deistler et al. (2022b) have proposed similar ideas to improve accuracy of SNPE. They train a classifier to predict samples that fall into rejection criterion based on a certain distance metric. However, such a classifier can falsely reject samples. Our method can avoid this problem and hence is more principled compared to their method.

We showcase several examples where either SNPE failed to perform inference effectively, such as in the SVAR case, or produced biased results, as observed in the BVCBM. For the SVAR example, SNPE methods struggle due to the impact of low-quality samples from certain parameter space regions, adversely affecting the training process. The ABC method can efficiently eliminate these bad samples, thereby enhancing the training. For cases where SNPE results in biased estimations, our methods were effective at accurately fitting observed data (real data for BVCBM example). This is substantiated by our empirical results for the posterior predictive distribution discussed in experimental section.

Although our method demonstrates the capability to enhance estimation accuracy, it does have some limitations. Firstly, our method requires model simulations in the ABC preconditioning step, which may lead to greater computational demands in situations where SNPE methods perform well. However, by performing the preconditioning step, significantly fewer model simulations may be required in the SNPE part to achieve high accuracy. In this paper we used SMC ABC for the preconditioning step, but we note that other ABC algorithms or SBI methods could be used. We do not recycle the simulations performed in the ABC preconditioning step for the SNPE phase, but it could be possible to modify our method to exploit these model simulations. Here we used the acceptance rate as the stopping criterion for the preconditioning step, but another option could be to check after each iteration of SMC ABC and stop the preconditioning if a suitable NCDE is found. Secondly, the choice of an unconditional density estimator necessitates careful consideration. In scenarios involving low-dimensional parameter spaces, a kernel density estimator might be a preferable option compared to unconditional normalizing flows.

It is important to note that the preconditioning step of our method requires manual selection of the summary statistics or discrepancy function. However, this choice is not as critical as in a typical ABC application, since we aim to remove poor parts of the parameter space rather than achieve a highly accurate posterior approximation. In this paper we used the same summary statistics as in the preconditioning step for the subsequent rounds of NPE. However, we note it would be possible to use different summaries after the preconditioning step, and possibly use automated summary statistic selection methods (Fearnhead & Prangle, 2012; Chen et al., 2023).

In this paper we considered the well-specified scenario, where the model is either known to be correct or can provide a good fit to the data with a suitable choice of parameter values. Standard neural SBI methods are known to potentially perform poorly under model misspecification Bon et al. (2023); Cannon et al. (2022); Schmitt et al. (2023). Our preconditioning method may be useful in the misspecified scenario, since ABC are known to perform reasonably well under model misspecification. That is, ABC still converges onto the pseudo-true parameter value Frazier et al. (2020). The preconditioning step could be followed by a robust neural SBI method such as Kelly et al. (2023); Huang et al. (2024); Gloeckler et al. (2023); Ward et al. (2022). We plan to explore this in future research.

Overall, PNPE employs a preconditioning step to focus on important parts of the parameter space, thereby creating a good starting point for training SNPE and enhancing estimation accuracy. We have empirically demonstrated that PNPE is capable of producing more accurate estimations in complex real-world problems.



## References

- Joseph Aylett-Bullock, Carolina Cuesta-Lazaro, Arnau Quera-Bofarull, Miguel Icaza-Lizaola, Aidan Sedgewick, Henry Truong, Aoife Curran, Edward Elliott, Tristan Caulfield, Kevin Fong, et al. June: open-source individual-based epidemiology simulation. *Royal Society Open Science*, 8(7):210506, 2021.
- Mark A Beaumont. Approximate Bayesian computation. *Annual Review of Statistics and Its Application*, 6:379–403, 2019.
- Mark A Beaumont, Jean-Marie Cornuet, Jean-Michel Marin, and Christian P Robert. Adaptive approximate Bayesian computation. *Biometrika*, 96(4):983–990, 2009.
- G rard Biau, Fr d ric C rou, and Arnaud Guyader. New insights into approximate Bayesian computation. In *Annales de l’IHP Probabilit s et Statistiques*, volume 51, pp. 376–403, 2015.
- Eli Bingham, Jonathan P Chen, Martin Jankowiak, Fritz Obermeyer, Neeraj Pradhan, Theofanis Karaletsos, Rohit Singh, Paul Szerlip, Paul Horsfall, and Noah D Goodman. Pyro: Deep universal probabilistic programming. *The Journal of Machine Learning Research*, 20(1):973–978, 2019.
- Michael GB Blum. Approximate Bayesian computation: a nonparametric perspective. *Journal of the American Statistical Association*, 105(491):1178–1187, 2010.
- Joshua J Bon, Adam Bretherton, Katie Buchhorn, Susanna Cramb, Christopher Drovandi, Conor Hassan, Adrienne L Jenner, Helen J Mayfield, James M McGree, Kerrie Mengersen, et al. Being Bayesian in the 2020s: opportunities and challenges in the practice of modern applied Bayesian statistics. *Philosophical Transactions of the Royal Society A*, 381(2247):20220156, 2023.
- Patrick Cannon, Daniel Ward, and Sebastian M Schmon. Investigating the impact of model misspecification in neural simulation-based inference. *arXiv preprint arXiv:2209.01845*, 2022.
- Yanzhi Chen, Michael U Gutmann, and Adrian Weller. Is learning summary statistics necessary for likelihood-free inference? In *International Conference on Machine Learning*, pp. 4529–4544. PMLR, 2023.
- Kyle Cranmer, Johann Brehmer, and Gilles Louppe. The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences*, 117(48):30055–30062, 2020.
- Katalin Csill ry, Michael GB Blum, Oscar E Gaggiotti, and Olivier Fran ois. Approximate Bayesian computation (ABC) in practice. *Trends in Ecology & Evolution*, 25(7):410–418, 2010.
- Maximilian Dax, Stephen R Green, Jonathan Gair, Jakob H Macke, Alessandra Buonanno, and Bernhard Sch lkopf. Real-time gravitational wave science with neural posterior estimation. *Physical Review Letters*, 127(24):241103, 2021.
- Natal  SM de Santi, Francisco Villaescusa-Navarro, L Raul Abramo, Helen Shao, Lucia A Perez, Tiago Castro, Yueying Ni, Christopher C Lovell, Elena Hernandez-Martinez, Federico Marinacci, et al. Field-level simulation-based inference with Galaxy catalogs: the impact of systematic effects. *arXiv preprint arXiv:2310.15234*, 2023.
- Michael Deistler, Pedro J Goncalves, and Jakob H Macke. Truncated proposals for scalable and hassle-free simulation-based inference. *Advances in Neural Information Processing Systems*, 35:23135–23149, 2022a.
- Michael Deistler, Jakob H Macke, and Pedro J Gon alves. Energy-efficient network activity from disparate circuit parameters. *Proceedings of the National Academy of Sciences*, 119(44):e2207632119, 2022b.
- Hadi Mohaghegh Dolatabadi, Sarah Erfani, and Christopher Leckie. Invertible generative modeling using linear rational splines. In *International Conference on Artificial Intelligence and Statistics*, pp. 4236–4246. PMLR, 2020.

- Christopher Drovandi, David J Nott, and David T Frazier. Improving the accuracy of marginal approximations in likelihood-free inference via localisation. *Journal of Computational and Graphical Statistics*, pp. 1–19, 2023.
- Christopher C Drovandi and Anthony N Pettitt. Estimation of parameters for macroparasite population evolution using approximate Bayesian computation. *Biometrics*, 67(1):225–233, 2011.
- Conor Durkan, Artur Bekasov, Iain Murray, and George Papamakarios. Neural spline flows. *Advances in Neural Information Processing Systems*, 32, 2019.
- Paul Fearnhead and Dennis Prangle. Constructing summary statistics for approximate Bayesian computation: semi-automatic approximate Bayesian computation. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 74(3):419–474, 2012.
- Alexander Fengler, Lakshmi N Govindarajan, Tony Chen, and Michael J Frank. Likelihood approximation networks (LANs) for fast inference of simulation models in cognitive neuroscience. *Elife*, 10:e65074, 2021.
- David T Frazier, Christian P Robert, and Judith Rousseau. Model misspecification in approximate Bayesian computation: consequences and diagnostics. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 82(2):421–444, 2020.
- Manuel Gloeckler, Michael Deistler, and Jakob H Macke. Adversarial robustness of amortized Bayesian inference. pp. 11493–11524, 2023.
- David Greenberg, Marcel Nonnenmacher, and Jakob Macke. Automatic posterior transformation for likelihood-free inference. In *International Conference on Machine Learning*, pp. 2404–2414. PMLR, 2019.
- Daolang Huang, Ayush Bharti, Amauri Souza, Luigi Acerbi, and Samuel Kaski. Learning robust statistics for simulation-based inference under model misspecification. *Advances in Neural Information Processing Systems*, 36, 2024.
- Adrianne L Jenner, Federico Frascoli, Adelle CF Coster, and Peter S Kim. Enhancing oncolytic virotherapy: Observations from a Voronoi Cell-Based model. *Journal of Theoretical Biology*, 485:110052, 2020.
- Ryan P Kelly, David J Nott, David T Frazier, David J Warne, and Chris Drovandi. Misspecification-robust sequential neural likelihood. *arXiv preprint arXiv:2301.13368*, 2023.
- Pyung-Hwan Kim, Joo-Hyuk Sohn, Joung-Woo Choi, Yookyung Jung, Sung Wan Kim, Seungjoo Haam, and Chae-Ok Yun. Active targeting and safety profile of PEG-modified adenovirus conjugated with herceptin. *Biomaterials*, 32(9):2314–2326, 2011.
- Jarno Lintusaari, Michael U Gutmann, Ritabrata Dutta, Samuel Kaski, and Jukka Corander. Fundamentals and recent developments in approximate Bayesian computation. *Systematic Biology*, 66(1):e66–e82, 2017.
- Jan-Matthis Lueckmann, Pedro J Goncalves, Giacomo Bassetto, Kaan Öcal, Marcel Nonnenmacher, and Jakob H Macke. Flexible statistical inference for mechanistic models of neural dynamics. *Advances in Neural Information Processing Systems*, 30, 2017.
- Jan-Matthis Lueckmann, Jan Boelts, David Greenberg, Pedro Goncalves, and Jakob Macke. Benchmarking simulation-based inference. In *International Conference on Artificial Intelligence and Statistics*, pp. 343–351. PMLR, 2021.
- Frank A Meineke, Christopher S Potten, and Markus Loeffler. Cell migration and organization in the intestinal crypt using a lattice-free model. *Cell Proliferation*, 34(4):253–266, 2001.
- Siddharth Mishra-Sharma. Inferring dark matter substructure with astrometric lensing beyond the power spectrum. *Machine Learning: Science and Technology*, 3(1):01LT03, 2022.
- George Papamakarios and Iain Murray. Fast  $\varepsilon$ -free inference of simulation models with Bayesian conditional density estimation. *Advances in Neural Information Processing Systems*, 29, 2016.

- George Papamakarios, Theo Pavlakou, and Iain Murray. Masked autoregressive flow for density estimation. *Advances in Neural Information Processing Systems*, 30, 2017.
- George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. *The Journal of Machine Learning Research*, 22(1):2617–2680, 2021.
- Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in Python. *the Journal of Machine Learning Research*, 12:2825–2830, 2011.
- Leah F Price, Christopher C Drovandi, Anthony Lee, and David J Nott. Bayesian synthetic likelihood. *Journal of Computational and Graphical Statistics*, 27(1):1–11, 2018.
- Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *International Conference on Machine Learning*, pp. 1530–1538. PMLR, 2015.
- Marvin Schmitt, Paul-Christian Bürkner, Ullrich Köthe, and Stefan T Radev. Detecting model misspecification in amortized Bayesian inference with neural networks. pp. 541–557, 2023.
- David Shih, Marat Freytsis, Stephen R Taylor, Jeff A Dror, and Nolan Smyth. Fast parameter inference on pulsar timing arrays with normalizing flows. *arXiv preprint arXiv:2310.12209*, 2023.
- Scott A Sisson, Yanan Fan, and Mark M Tanaka. Sequential Monte Carlo without likelihoods. *Proceedings of the National Academy of Sciences*, 104(6):1760–1765, 2007.
- Scott A Sisson, Yanan Fan, and Mark Beaumont. *Handbook of approximate Bayesian computation*. CRC Press, 2018.
- Alvaro Tejero-Cantero, Jan Boelts, Michael Deistler, Jan-Matthis Lueckmann, Conor Durkan, Pedro J Gonçalves, David S Greenberg, and Jakob H Macke. SBI: A toolkit for simulation-based inference. *Journal of Open Source Software*, 5(52):2505, 2020.
- Owen Thomas, Henri Pesonen, Raquel Sá-Leão, Hermínia de Lencastre, Samuel Kaski, and Jukka Corander. Split-BOLFI for misspecification-robust likelihood free inference in high dimensions. *arXiv preprint arXiv:2002.09377*, 2020.
- Brandon M Turner and Trisha Van Zandt. A tutorial on approximate Bayesian computation. *Journal of Mathematical Psychology*, 56(2):69–85, 2012.
- Samantha Jane Wade. Fabrication and preclinical assessment of drug eluting wet spun fibres for pancreatic cancer treatment. 2019.
- Xiaoyu Wang, Adrienne L Jenner, Robert Salomone, David J Warne, and Christopher Drovandi. Calibration of agent based models for monophasic and biphasic tumour growth using approximate Bayesian computation. *Journal of Mathematical Biology*, 88(3):28, 2024.
- Daniel Ward, Patrick Cannon, Mark Beaumont, Matteo Fasiolo, and Sebastian Schmon. Robust neural posterior estimation and statistical model criticism. *Advances in Neural Information Processing Systems*, 35:33845–33859, 2022.
- Antoine Wehenkel, Jens Behrmann, Andrew C Miller, Guillermo Sapiro, Ozan Sener, Marco Cuturi, and Jörn-Henrik Jacobsen. Simulation-based inference for cardiovascular models. *arXiv preprint arXiv:2307.13918*, 2023.
- Timothy O West, Luc Berthouze, Simon F Farmer, Hayriye Cagnan, and Vladimir Litvak. Inference of brain networks with approximate Bayesian computation—assessing face validity with an example application in Parkinsonism. *Neuroimage*, 236:118020, 2021.
- Christina Winkler, Daniel Worrall, Emiel Hoogeboom, and Max Welling. Learning likelihoods with conditional normalizing flows. *arXiv preprint arXiv:1912.00042*, 2019.

## A Further background for SMC ABC

We provide a detailed description of the adaptive SMC ABC algorithm we used in this paper and provide pseudocode in Algorithm 2 for reference. This algorithm starts by drawing  $N$  independent samples from the prior distribution  $p(\theta)$ , represented as  $\{\theta_i\}_{i=1}^N$ . For each sample  $\theta_i$  (known as a particle), the algorithm simulates a dataset  $x_i$  from the stochastic model and calculates the corresponding discrepancy  $\rho_i = \rho(x_i, x_o)$ , resulting in the pair set  $\{\theta_i, \rho_i\}_{i=1}^N$ . These pairs are then arranged in order of increasing discrepancy such that  $\rho_1 < \rho_2 < \dots < \rho_N$ . The first tolerance threshold,  $\epsilon_1$ , is set as the largest discrepancy,  $\rho_N$ . To move through the target distributions, the algorithm adjusts the tolerance dynamically. The next tolerance,  $\epsilon_t$ , is set as  $\rho_{N-N_a}$ , where  $N_a = \lfloor Na \rfloor$ , and  $a$  is a tuning parameter. Essentially, in each step, the algorithm discards the top  $a \times 100\%$  of particles with the highest discrepancies. After discarding these particles, only  $N - N_a$  particles remain. To replenish the set back to  $N$  particles, the algorithm resamples  $N_a$  times from the ‘alive’ particles, copying both the parameter and discrepancy values. This process, however, leads to duplicates in the particle set. To add variety to the set, the algorithm applies an MCMC ABC kernel to each resampled particle. The parameters for the MCMC proposal distribution  $q_t(\cdot|\cdot)$  are derived from the current particle set. For instance, if using a multivariate normal random walk proposal, its covariance  $\Sigma_t$  is based on the particle set’s sample covariance. The acceptance of a proposed parameter (assuming a symmetric proposal) and simulated data is determined by the equation:

$$p_t = \min \left( 1, \frac{p(\tilde{\theta})}{p(\theta)} \mathbb{I}(\rho(x_o, \tilde{x}) < \epsilon_t) \right), \quad (6)$$

where  $\tilde{\theta} \sim q(\cdot|\theta)$  and  $\tilde{x} \sim p(\cdot|\tilde{\theta})$  are proposed parameter values and dataset, respectively. However, proposals may be rejected, leaving some particles unchanged. To address this, the algorithm performs  $R_t$  iterations of the MCMC kernel on each particle, where  $R_t = \left\lceil \frac{\log(c)}{\log(1-p_t^{\text{acc}})} \right\rceil$ , where  $c$  is a tuning parameter of the algorithm that can be interpreted as the probability that a particle is not moved in the  $R_t$  iterations. The acceptance probability  $p_t^{\text{acc}}$  is estimated from trial MCMC ABC iterations and used to compute  $R_t$  for the next set of MCMC ABC iterations. For this adaptive SMC ABC algorithm, two stopping rules can be used. The first stopping rule halts the ABC algorithm when the maximum discrepancy is below a set tolerance,  $\epsilon_T$ . The second stopping rule terminates the algorithm when the MCMC acceptance probability  $p_t^{\text{acc}}$  falls below a predefined threshold  $p_{\text{acc}}$ . Here we choose the second rule, and since we only require a short run of ABC, we set  $p_{\text{acc}}$  to be higher than what is typically used in an ABC analysis.

**Algorithm 2** Adaptive SMC ABC

**Input:** The observed data  $x_o$ , the stochastic model  $p(x|\theta)$ , distance function  $\rho(\cdot, \cdot)$ , prior distribution  $p(\theta)$ , number of particles  $N$ , tuning parameters  $a$  and  $c$  for adaptive selection of discrepancy thresholds and selecting the number of MCMC iterations in the move steps, target tolerance  $\epsilon_T$ , initial number of trial MCMC iterations  $S_{\text{init}}$ , minimum acceptable MCMC acceptance rate  $p_{\text{min}}$

**for**  $i = 1, \dots, N$  **do**

    Simulate  $x_i \sim p(x|\theta_i)$  where  $\theta_i \sim p(\theta)$

    Compute  $\rho_i = \rho(x_o, x_i)$

**end for**

Sort  $\{\theta_i\}_{i=1}^N$  by  $\{\rho_i\}_{i=1}^N$  such that  $\rho_1 \leq \rho_2 \leq \dots \leq \rho_N$

Set  $N_a = \lfloor aN \rfloor$ ,  $t = 2$ ,  $\epsilon_t = \rho_{N-N_a}$ ,  $\epsilon_1 = \rho_N$ ,  $S_t = S_{\text{init}}$ ,  $\tilde{p}_t^{\text{acc}} = 1$

**while**  $\epsilon_{t-1} > \epsilon_T$  **or**  $\tilde{p}_t^{\text{acc}} > p_{\text{min}}$  **do**

    Compute  $\Sigma$  as the sample covariance matrix of  $\{\theta_i\}_{i=1}^{N-N_a}$

    Generate  $\{\theta_i\}_{i=N-N_a+1}^N$  by resampling from  $\{\theta_i\}_{i=1}^{N-N_a}$  with replacement

**for**  $i = N - N_a + 1, \dots, N$  **do**

**for**  $j = 1, \dots, S_t$  **do**

            Simulate  $\tilde{x} \sim p(x|\tilde{\theta})$  based on proposal  $\tilde{\theta} \sim \mathcal{N}(\theta_i, \Sigma)$

            Compute  $\tilde{\rho} = \rho(x_o, \tilde{x})$

            Compute  $p_t^{i,j} = \min\left(1, \frac{p(\tilde{\theta})}{p(\theta)} \mathbb{I}(\tilde{\rho} < \epsilon_t)\right)$ .

            With probability  $p_t^{i,j}$ , set  $\theta_i = \tilde{\theta}$  and  $\rho_i = \tilde{\rho}$ ; otherwise, retain the current values of  $\theta_i$  and  $\rho_i$

**end for**

**end for**

$\tilde{p}_t = \sum_{i=N-N_a+1}^N \sum_{j=1}^{S_t} p_t^{i,j} / (S_t(N - N_a))$

$R_t = \lceil \log(c) / (1 + \log(1 - \tilde{p}_t)) \rceil$

**for**  $i = N - N_a + 1, \dots, N$  **do**

**for**  $j = R_t - S_t, \dots, R_t$  **do**

            Simulate  $\tilde{x} \sim p(x|\tilde{\theta})$  based on proposal  $\tilde{\theta} \sim \mathcal{N}(\theta_i, \Sigma)$

            Compute  $\tilde{\rho} = \rho(x_o, \tilde{x})$

            Compute  $p_t^{i,j} = \min\left(1, \frac{p(\tilde{\theta})}{p(\theta)} \mathbb{I}(\tilde{\rho} < \epsilon_t)\right)$ .

            With probability  $p_t^{i,j}$ , set  $\theta_i = \tilde{\theta}$  and  $\rho_i = \tilde{\rho}$ ; otherwise, retain the current values of  $\theta_i$  and  $\rho_i$

**end for**

**end for**

$\tilde{p}_t^{\text{acc}} = \sum_{i=N-N_a+1}^N \sum_{j=1}^{R_t} p_t^{i,j} / (R_t(N - N_a))$

$S_{t+1} = \lceil R_t/2 \rceil$

    Sort  $\{\theta_i\}_{i=1}^N$  by  $\{\rho_i\}_{i=1}^N$  such that  $\rho_1 \leq \rho_2 \leq \dots \leq \rho_N$

    Set  $\epsilon_{t+1} = \rho_{N-N_a}$ ,  $\epsilon_t = \rho_N$

$t = t + 1$

**end while**

**return** Samples  $\{\theta_i\}_{i=1}^N$  from ABC posterior

## B Experimental Details

We use the adaptive SMC ABC algorithm Drovandi & Pettitt (2011) in the preconditioning step for all experiments. We set the tuning parameters as  $a = 0.5$ ,  $c = 0.01$  and use 1k particles for the algorithm. As the stopping rule, we set the target MCMC acceptance rate at 10%, unless otherwise specified. For the unconditional density estimator, we employ unconditional normalizing flows using the `Pyro` package Bingham et al. (2019), with a spline coupling layer using the transformation:

$$Y_{1:d} = g_{\tilde{\theta}}(X_{1:d}) \quad (7)$$

$$Y_{(d+1):D} = h_{\phi}(X_{(d+1):D}; X_{1:d}) \quad (8)$$

where  $\mathbf{X}$  are the inputs,  $\mathbf{Y}$  are the outputs, e.g.,  $X_{1:d}$  represents the first  $d$  elements of the inputs,  $g_{\tilde{\theta}}$  is either the identity function or an elementwise rational monotonic spline with parameters  $\tilde{\theta}$ , and  $h_{\phi}$ , where  $\phi$  is element-wise bijection parameter, is a conditional elementwise spline, conditioning on the first  $d$  elements. Regarding the neural networks, we use four fully-connected layers and set the count bins to 16. Furthermore, if the dimensions of the parameter space are less than three, indicating a low-dimensional case, we also consider kernel density estimation with a Gaussian kernel as the unconditional density estimator, as implemented in the `Scikit-learn` package Pedregosa et al. (2011). For APT and TSNPE, we use the implementation of the `sbi` package Tejero-Cantero et al. (2020) with default settings.

For SNPE and TSNPE, we use conditional neural spline flows Durkan et al. (2019). We use five coupling layers, with each coupling layer using a multilayer perceptron of two layers with 50 hidden units. The flow is trained using the Adam optimiser with a learning rate of  $5 \times 10^{-4}$  and a batch size of 256. Flow training is stopped when either the validation loss, calculated on 10% of the samples, has not improved over 50 epochs or when the limit of 500 epochs is reached.

### B.1 High-dimensional SVAR model

We consider our illustrative example in a high-dimensional setting with  $k = 20$ , which leads to 21 parameters that need to be estimated. We set the true parameter values as follows:

$$\begin{aligned} \theta = & (-0.2764, -0.7765, 0.8231, -0.1972, -0.2254, 0.6334, 0.4495, \\ & 0.4465, -0.8961, 0.0647, -0.1791, 0.0795, -0.5464, -0.9354, \\ & -0.4639, -0.7851, -0.6833, -0.1408, 0.7032, 0.8321, 0.1000), \end{aligned}$$

and use these values to simulate the observation dataset. We employ the same summary statistics as in the low-dimensional example and use a uniform distribution as the prior, constrained between -1 and 1 for the  $k$  parameters and between 0 and 1 for  $\sigma$ .

Since the BSL method is the gold standard for this example, we use the standard BSL method proposed by Price et al. (2018), with the total number of simulations set to 20 million.

### B.2 BVCBM

The BVCBM simulation begins by initializing a square domain with cells arranged in a hexagonal lattice. The cell at the center of the domain is identified as a cancer cell, while the others are designated as healthy cells. The simulation proceeds until the tumor reaches a volume of  $100 \text{ mm}^2$ , in accordance with experimental measurements (Wade, 2019; Kim et al., 2011). When this volume is attained, the distribution of healthy and cancerous cells within the lattice is recorded. This configuration then serves as the starting point to simulate tumor growth over the desired number of days. The model progresses by determining whether a cancer cell proliferates at each timestep, as described by the equation:

$$p_d = p_0 \left( 1 - \frac{d}{d_{\max}} \right), \quad (9)$$

where  $p_d$  is the probability of cell division,  $p_0$  is the initial division rate,  $d$  is the current cell density, and  $d_{\max}$  is the maximum density. For cancer cells that do not proliferate, the model assesses their potential to

transition into invasive cells, governed by the probability  $p_{\text{psc}}$ . Subsequently, the positions of all cells, both healthy and cancerous, are updated using Hooke's law:

$$\mathbf{r}_i(t + \Delta t) = \mathbf{r}_i(t) + \frac{1}{\mu} \mathbf{F}_i(t) \Delta t = \mathbf{r}_i(t) + \lambda \sum_{\forall j} \frac{\mathbf{r}_{i,j}(t)}{\|\mathbf{r}_{i,j}(t)\|} (s_{i,j}(t) - \|\mathbf{r}_{i,j}(t)\|). \quad (10)$$

Here,  $\mathbf{r}_i(t + \Delta t)$  denotes the updated position of cell  $i$ ,  $\mu$  is the cell motility coefficient,  $\mathbf{F}_i(t)$  is the force on cell  $i$ ,  $\lambda$  is a mechanical interaction coefficient,  $\mathbf{r}_{i,j}(t)$  is the vector between cells  $i$  and  $j$ , and  $s_{i,j}(t)$  is the natural length of the spring connecting the two cells. The parameters for the mechanical interactions, such as  $\lambda$  and  $\mu$ , are sourced from prior studies in the literature Meineke et al. (2001). See Jenner et al. (2020); Wang et al. (2024) for more detailed model simulation.

Four parameters  $\theta = (p_0, p_{\text{psc}}, d_{\text{max}}, g_{\text{age}})$  control the tumor growth during a single phase, which is a period when the tumor grows based on fixed values for these four parameters. For the biphasic model, an additional parameter  $\tau$  is introduced, representing the time at which the tumor growth pattern changes, that is, the values for  $\theta$  change. Therefore, for BVCBM, we need to estimate nine parameters for two pancreatic cancer datasets, denoted as  $\theta_1 = (p_0^1, p_{\text{psc}}^1, d_{\text{max}}^1, g_{\text{age}}^1)$ ,  $\theta_2 = (p_0^2, p_{\text{psc}}^2, d_{\text{max}}^2, g_{\text{age}}^2)$ , and  $\tau$ , so that  $\theta = (\theta_1, \theta_2, \tau)$ .

The parameter  $p_{\text{psc}}$ , which is the probability of tumor cell invasion into healthy cells, significantly affects the simulation time. The value of  $p_{\text{psc}}$  should be around  $10^{-5}$ , indicating that an increase in probability will require more cells to be simulated. Moreover, a smaller value of  $p_{\text{psc}}$  results in simulation time. For PNPE, the total simulation time for 15k simulations (26-day dataset) and 17k simulations (32-day dataset) for the preconditioning step is approximately 13 and 15 minutes, respectively, whereas SNPE requires around 1 hour for the first round (i.e. based on samples from the prior) of 10k simulations. This is because the preconditioning step is effective at quickly eliminating values of  $p_{\text{psc}}$  that lead to longer model simulation times.

## C Further experimental results

In this section, we present prior predictive distributions of tumour volumes in (a) and (b) for two pancreatic datasets (in (c) and (d) we show the same plots but on the log scale). We also show the posterior predictive distributions on the log scale obtained with different methods in Figure 9. Then we present the bivariate posterior density and marginal posterior density plots for the BVCBM as additional results. It is evident from Figures 10 and 11 that PNPE provides more precise estimation than SNPE for both pancreatic cancer datasets.

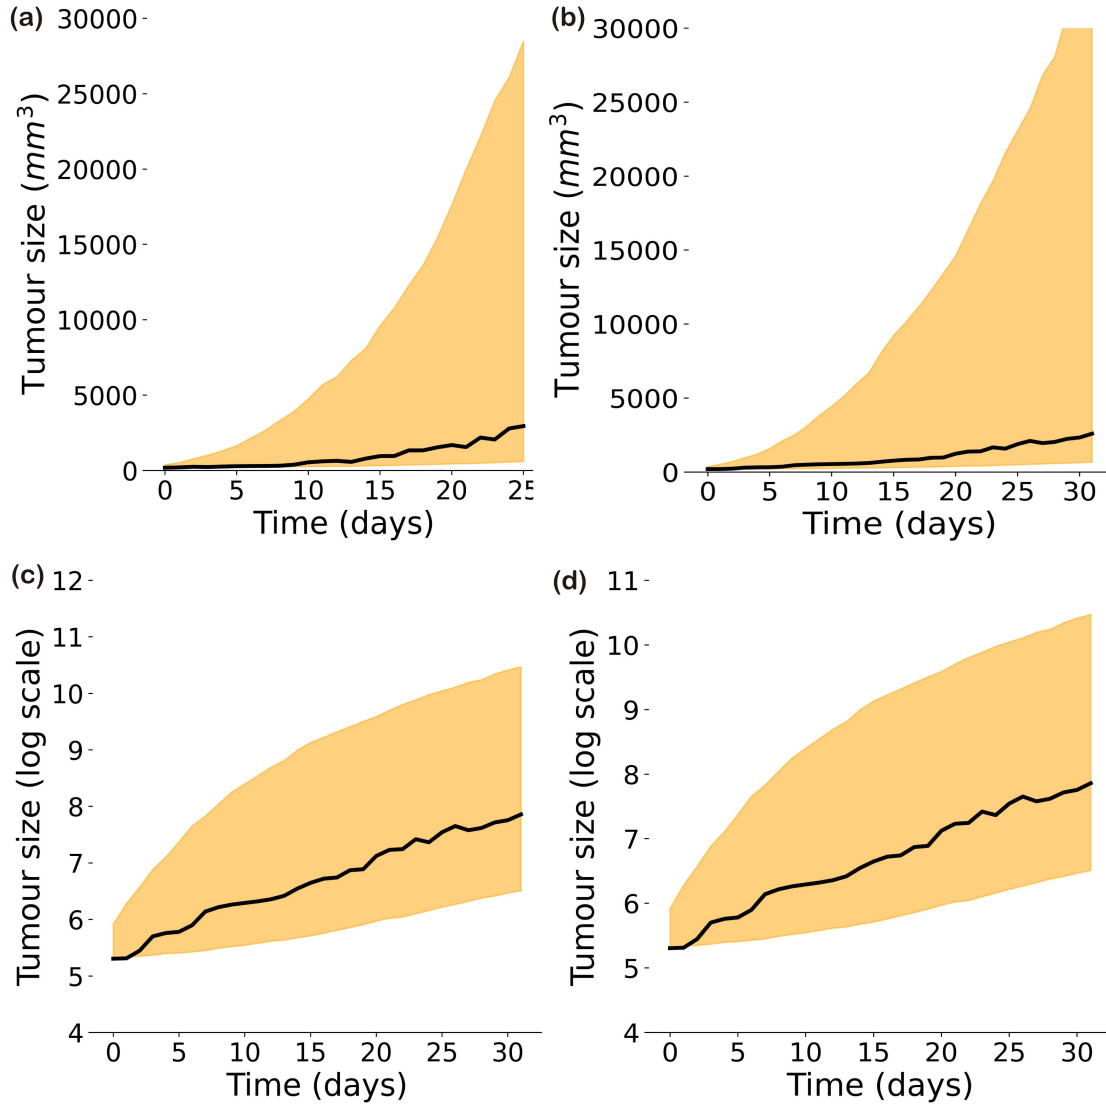


Figure 8: **Prior predictive distribution for BVCBM.** We sample 10k parameter values from prior distribution and plot the prior predictive distribution for two pancreatic datasets. In (a) and (b), the plots are in regular scale, and in (c) and (d), the plots are in log scale.



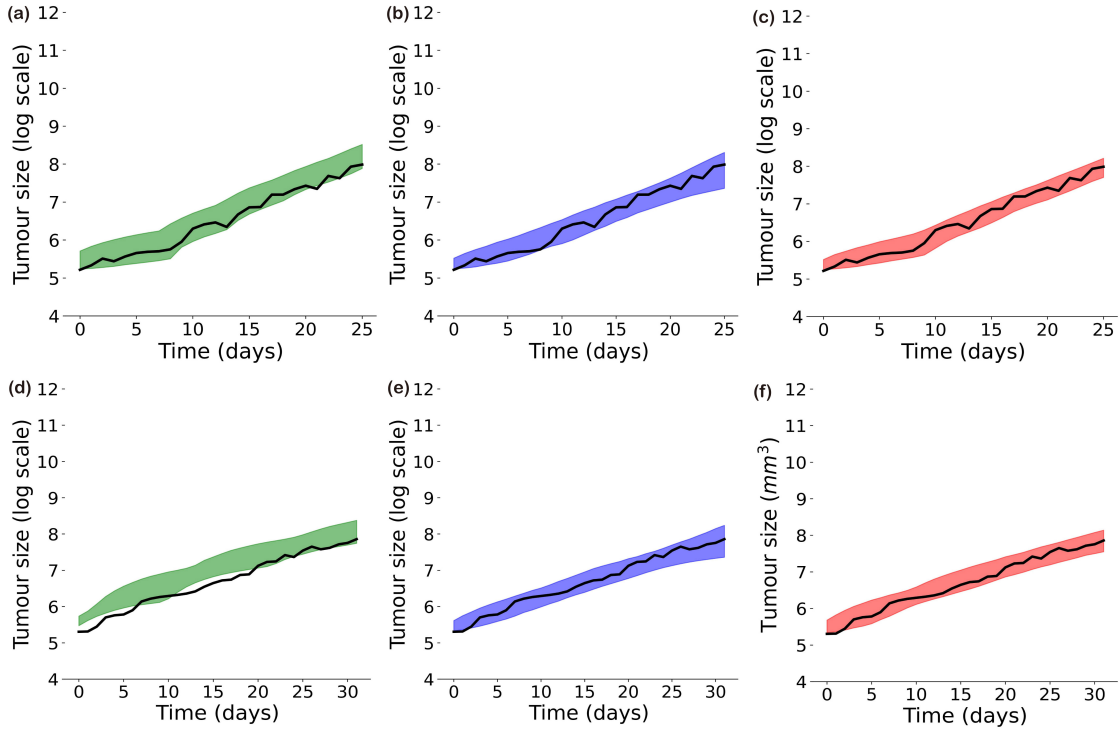


Figure 9: **Posterior predictive distributions for two pancreatic cancer datasets in log scale.** 90% Posterior predictive interval plots for SNPE (a), preconditioning step (b) and PNPE (c) for the 26-day dataset; 90% Posterior predictive interval plots for SNPE (d), preconditioning step (e) and PNPE (f) for the 32-day dataset.

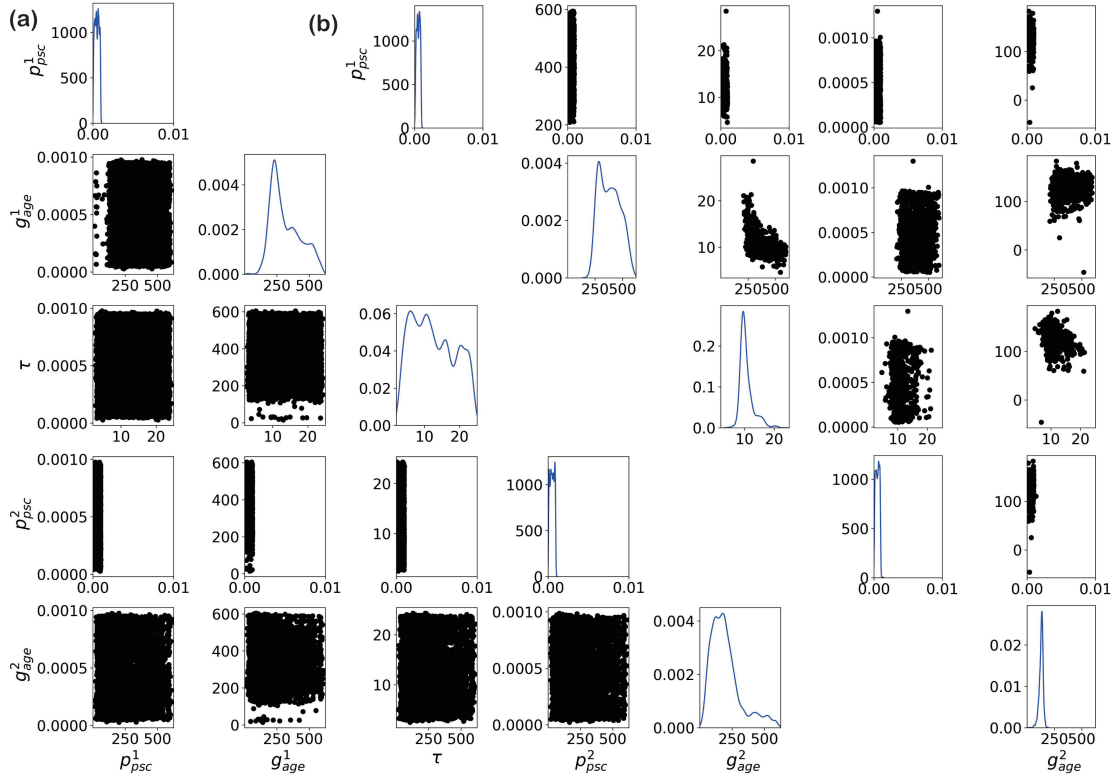


Figure 10: Bivariate density plots for the pancreatic dataset with 26-day measurements for (a) SNPE and (b) PNPE. The diagonal entries represent the marginal posterior densities for  $(p_{psc}^1, g_{age}^1, \tau, p_{psc}^2, g_{age}^2)$ .

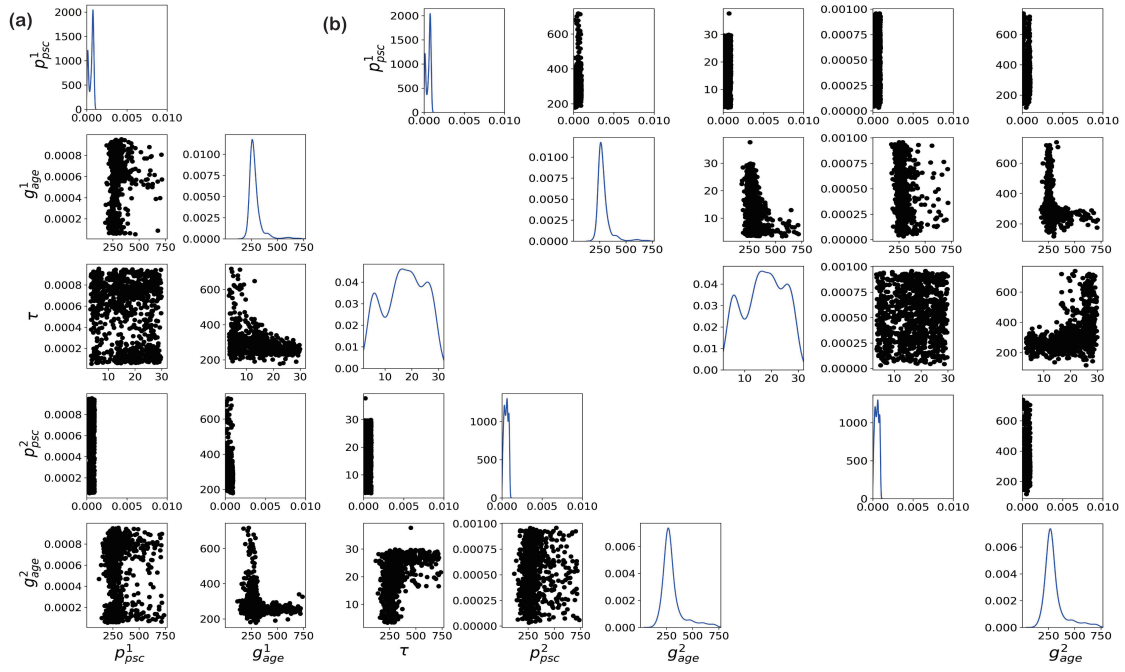


Figure 11: Bivariate density plots for the pancreatic dataset with 32-day measurements for (a) SNPE and (b) PNPE. The diagonal entries represent the marginal posterior densities for  $(p^1_{psc}, g^1_{age}, \tau, p^2_{psc}, g^2_{age})$ .