

# Score Matched Neural Exponential Families for Likelihood-Free Inference

**Lorenzo Pacchiardi**

*Department of Statistics  
University of Oxford  
Oxford, OX1 3LB  
United Kingdom*

LORENZO.PACCHIARDI@STATS.OX.AC.UK

**Ritabrata Dutta**

*Department of Statistics  
University of Warwick  
Coventry, CV4 7AL  
United Kingdom*

RITABRATA.DUTTA@WARWICK.AC.UK

**Editor:** Mohammad Emtiyaz Khan

## Abstract

Bayesian Likelihood-Free Inference (LFI) approaches allow to obtain posterior distributions for stochastic models with intractable likelihood, by relying on model simulations. In Approximate Bayesian Computation (ABC), a popular LFI method, summary statistics are used to reduce data dimensionality. ABC algorithms adaptively tailor simulations to the observation in order to sample from an approximate posterior, whose form depends on the chosen statistics. In this work, we introduce a new way to learn ABC statistics: we first generate parameter-simulation pairs from the model independently on the observation; then, we use Score Matching to train a neural conditional exponential family to approximate the likelihood. The exponential family is the largest class of distributions with fixed-size sufficient statistics; thus, we use them in ABC, which is intuitively appealing and has state-of-the-art performance. In parallel, we insert our likelihood approximation in an MCMC for doubly intractable distributions to draw posterior samples. We can repeat that for any number of observations with no additional model simulations, with performance comparable to related approaches. We validate our methods on toy models with known likelihood and a large-dimensional time-series model.

Code for reproducing the experiments is available at <https://github.com/LoryPack/SM-ExpFam-LFI>.

**Keywords:** Likelihood-Free Inference, Score Matching, Approximate Bayesian Computation, MCMC for doubly intractable distributions, Exponential Family.

## 1. Introduction

Stochastic simulator models are used to simulate realizations of physical phenomena; usually, a set of parameters  $\theta$  govern the simulation output. As the model is stochastic, repeated simulations with fixed  $\theta$  yield different outputs; their distribution is the model's *likelihood*. Simple models provide an analytic expression of the likelihood, which is however unavailable for more complex ones.

Upon observing a real-world realization of the phenomenon the model is describing, researchers may want to obtain a *posterior distribution* over parameters. If the likelihood is known, standard Bayesian inference tools (such as MCMC or variational inference) allow to get posterior samples; if otherwise the likelihood is missing, Likelihood-Free Inference (LFI) techniques are the only viable solution. LFI has been applied in several domains, including genomics (Tavaré et al., 1997), biological science (Dutta et al., 2018), meteorology (Hakkarainen et al., 2012), geological science (Pacchiardi et al., 2020), genomics (Tavaré et al., 1997; Toni et al., 2009; Marttinen et al., 2015), and epidemiology (McKinley et al., 2018; Minter and Retkute, 2019; Dutta et al., 2021a).

In some LFI techniques (Wood, 2010; Thomas et al., 2020; Price et al., 2018), the missing likelihood is replaced with an explicit approximation built from model simulations at each parameter value. Alternatively, Approximate Bayesian Computation (ABC) (Marin et al., 2012; Lintusaari et al., 2017) circumvents the unavailability of the likelihood by comparing model simulations with the observation according to some notion of distance. To reduce data dimensionality, ABC usually relies on summary statistics, whose choice is not straightforward. Recently, the expressive capabilities of Neural Networks (NNs) have been leveraged to learn ABC statistics (Jiang et al., 2017; Wiqvist et al., 2019; Pacchiardi et al., 2020; Åkesson et al., 2020). These techniques train NNs parametrizing the statistics by minimizing a suitable loss on a training set of parameter-simulation pairs generated from the model.

In the present work, we propose a new way to learn ABC statistics with NNs. Most previous works (Jiang et al., 2017; Wiqvist et al., 2019; Åkesson et al., 2020) trained a single NN with the regression loss introduced in Fearnhead and Prangle (2012). Instead, we consider an exponential family with two NNs parametrizing respectively the sufficient statistics and natural parameters and fit this to the model. As the exponential family is the most general class of distributions with sufficient statistics of fixed size (Appendix A), it makes intuitive sense to use the learned sufficient statistics in ABC. Indeed, this approach empirically yields superior or equivalent performance with respect to state-of-the-art approaches.

As in previous approaches, we consider a training set of parameter-simulation pairs; however, we train the NNs with Score Matching (SM) or its Sliced approximation (SSM), which do not require evaluating the normalizing constant of the exponential family. We extend the SM and SSM objectives to the setting of conditional densities, thus fitting the likelihood approximation for all values of data and parameters.

In contrast to related approaches (Jiang et al., 2017; Wiqvist et al., 2019; Pacchiardi et al., 2020; Åkesson et al., 2020), our method provides a full likelihood approximation. We test therefore direct sampling from the corresponding posterior (in place of ABC) with an MCMC algorithm for doubly intractable distributions. This approach achieves satisfactory performance with no additional model simulations and is therefore a valid alternative to standard LFI schemes for expensive simulator models. The computational gain is even larger when inference is performed for several observations, as the same likelihood approximation can be used.

The rest of our paper is organized as follows. In Section 2, we briefly review some LFI methods. Next, in Section 3 we introduce the neural conditional exponential family and show how to fit it with SM or SSM. In Section 4 we discuss how to exploit the exponential family to extract ABC statistics or for direct sampling. We extensively validate our proposed

approaches in Section 5. Finally, we discuss related works in Section 6 and conclude in Section 7, where we also highlight directions for future research.

### 1.1 Notation

We set here notation for the rest of our manuscript. We will denote respectively by  $\mathcal{X} \subseteq \mathbb{R}^d$  and  $\Theta \subseteq \mathbb{R}^p$  the data and parameter space. Upper case letters will denote random variables while lower case ones will denote observed (fixed) values. Subscripts will denote vector components and superscript in brackets sample indices, while  $\|\cdot\|$  will denote the  $\ell_2$  norm.

## 2. Likelihood-Free Inference

Let us consider a model which allows to generate a simulation  $x \in \mathcal{X}$  at any parameter value  $\theta \in \Theta$ , but for which it is not possible to evaluate the likelihood  $p_0(x|\theta)$ . Given an observation  $x^0$  and a prior on the parameters  $\pi(\theta)$ , Bayesian inference obtains the posterior distribution  $\pi_0(\theta|x^0) = \frac{\pi(\theta)p_0(x^0|\theta)}{p_0(x^0)}$ . However, obtaining that explicitly (or even sampling from it with Markov Chain Monte Carlo, MCMC) is impossible without having access to the likelihood, at least up to a normalizing constant that is independent on  $x$ .

Likelihood-Free Inference techniques yield approximations of the posterior distribution by replacing likelihood evaluations with model simulations. Broadly, they can be split into two kinds of approaches: Surrogate Likelihood, which explicitly builds a likelihood function, and Approximate Bayesian Computation (ABC), which instead uses discrepancy between simulated and observed data. We detail those two approaches in the following.

### 2.1 Approximate Bayesian Computation (ABC)

Approximate Bayesian Computation (ABC, Marin et al. 2012; Lintusaari et al. 2017) algorithms sample from an approximate posterior distribution by finding parameter values that yield simulated data resembling the observations to a sufficient degree. Usually, ABC methods rely on statistics function to reduce the dimensionality of simulated and observed data and thus improve the computational efficiency of the algorithm. Therefore, the similarity between simulated and observed data is normally defined via a distance function between the corresponding statistics. ABC algorithms sample from:

$$\pi^\epsilon(\theta|s(x^0)) \propto \pi(\theta) \int_{\mathcal{X}} \mathbb{1}[d(s(x^0), s(x^{\text{sim}})) \leq \epsilon] p_0(x^{\text{sim}}|\theta) dx^{\text{sim}},$$

where  $s$  and  $d$  denote respectively the statistics function and the discrepancy measure, while  $\mathbb{1}[\cdot]$  is an indicator function and  $\epsilon$  is an acceptance threshold.<sup>1</sup>

It is easy to sample from  $\pi^\epsilon(\theta|s(x^0))$  with rejection, by generating  $(\theta^{(j)}, x^{(j)}) \sim \pi(\theta)p_0(x|\theta)$  and accepting  $\theta^{(j)}$  if  $d(s(x^0), s(x^{(j)})) \leq \epsilon$ . This is inefficient when  $\epsilon$  is small, as most simulations end up being rejected; more efficient ABC algorithms have been developed, based for instance on Markov Chain Monte Carlo (Marjoram et al., 2003), Population Monte Carlo

---

1. In general,  $\mathbb{1}(d(s(x^0), s(x^{\text{sim}})))$  can be replaced with  $\mathbb{K}_\epsilon(d(s(x^0), s(x^{\text{sim}})))$ , where  $\mathbb{K}_\epsilon$  is a rejection kernel which concentrates around 0 and goes to 0 for argument going to  $\pm\infty$ , and whose width decreases when  $\epsilon$  decreases. We restrict here on  $\mathbb{1}$ , which is the most common choice.

(Beaumont, 2010) or Sequential Monte Carlo (Del Moral et al., 2012), which reach lower values of  $\epsilon$  with the same computational budget.

If  $d$  is a distance measure and the statistic  $s$  is sufficient for  $p(x|\theta)$ , then the ABC posterior  $\pi^\epsilon(\theta|s(x^0))$  converges to the true one as  $\epsilon \rightarrow 0$  (Lintusaari et al. 2017; see Appendix A for definition of sufficient statistics).

In general, however, sufficient statistics are unavailable so that ABC can at best give an approximation of the exact posterior even with  $\epsilon \rightarrow 0$ . Hand-picking the best summary statistics for a given model limits the applicability of ABC. Therefore, some approaches to automatically build them have been developed (Fearnhead and Prangle, 2012; Nunes and Balding, 2010; Joyce and Marjoram, 2008; Aeschbacher et al., 2012; Blum et al., 2013; Pacchiardi et al., 2020). We describe here the approach that is most commonly used.

**ABC with semi-automatic statistics.** The summary statistic  $s^*(x) = \mathbb{E}[\theta|x]$  yields an approximate posterior  $\pi^\epsilon(\theta|s^*(x^0))$  whose mean minimizes the expected squared error loss from the true parameter value, as  $\epsilon \rightarrow 0$  (Fearnhead and Prangle, 2012). Of course,  $\mathbb{E}[\theta|x]$  is unknown; therefore, Fearnhead and Prangle (2012) suggested to consider a set of functions  $s_\beta(x)$  and find the one that best approximates  $\mathbb{E}[\theta|x]$ . Empirically, this corresponds to finding:

$$\hat{\beta} = \arg \min_{\beta} \frac{1}{N} \sum_{j=1}^N \left\| s_\beta(x^{(j)}) - \theta^{(j)} \right\|_2^2, \quad (1)$$

where  $(\theta^{(j)}, x^{(j)})_{j=1}^N$ ,  $\theta^{(j)} \sim \pi(\theta)$ ,  $x^{(j)} \sim p_0(\cdot|\theta^{(j)})$  are  $N$  data-parameter pairs generated before performing ABC.

In practice, Fearnhead and Prangle (2012) considered  $s_\beta(x)$  to be a linear function of the weights  $\beta$ , such that solving the minimization problem in Eq. (1) amounts to linear regression. Jiang et al. (2017); Wiqvist et al. (2019) and Åkesson et al. (2020) instead used Neural Networks to parametrize  $s_\beta$  and train them to minimize the loss in Eq. (1).

## 2.2 Surrogate Likelihood

Surrogate likelihood approaches exploit simulations to build an explicit likelihood approximation, which is then inserted in standard likelihood-based sampling schemes (say, MCMC; Wood 2010; Thomas et al. 2020; Fasiolo et al. 2018; An et al. 2019, 2020; Price et al. 2018). Here, we discuss two methods that fall under this framework.

**Synthetic Likelihood.** Synthetic Likelihood (SL, Wood 2010) replaces the exact likelihood with a normal distribution for the summary statistics  $s = s(x)$ ; specifically, it assumes  $S|\theta \sim \mathcal{N}(\mu_\theta, \Sigma_\theta)$ , where the mean vector  $\mu_\theta$  and covariance matrix  $\Sigma_\theta$  depend on  $\theta$ . For each  $\theta$ , an estimate of  $\mu_\theta$  and  $\Sigma_\theta$  is built with model simulations, and the likelihood of the observed statistics  $s(x^0)$  is evaluated.

Using the statistics likelihood to define a posterior distribution yields a Bayesian SL (BSL). In Price et al. (2018), MCMC is used to sample from the BSL posterior; as in ABC, this approach requires generating simulations for each considered  $\theta$ . However, one single simulation per parameter value is usually sufficient in ABC, while estimating  $\mu_\theta, \Sigma_\theta$  in BSL requires multiple simulations. Nevertheless, a parametric likelihood approximation seemingly allows scaling to larger dimensional statistics space (Price et al., 2018).

**Ratio Estimation.** Thomas et al. (2020) tackle LFI by estimating the ratio between likelihood and data marginal:  $r(x, \theta) = p(x|\theta)/p(x)$ . As  $r(x, \theta) \propto p(x|\theta)$  with respect to  $\theta$ , this can be inserted in a likelihood-based sampling scheme in order to sample from the posterior, similarly to BSL. This method is termed Ratio Estimation (RE).

In practice, for each value of  $\theta$ , an estimate  $\hat{r}(x, \theta)$  is built by generating simulations from the model  $p_0(x|\theta)$  and from the marginal<sup>2</sup>  $p(x)$ , and then fitting a logistic regression discriminating between the two sets of samples. In fact, logistic regression attempts to find a function  $\hat{h}(x; \theta)$  for which  $e^{\hat{h}(x; \theta)} \approx r(x; \theta)$ .

In Thomas et al. (2020),  $\hat{h}$  is chosen in the class of functions  $h_\beta(x; \theta) = \sum_{i=1}^k \beta_i(\theta) \psi_i(x) = \beta(\theta)^T \psi(x)$ , where  $\psi$  is a vector of summary statistics and  $\beta(\theta)$  are coefficients determined independently for each  $\theta$ ; this therefore boils down to a linear logistic regression. Viewed differently, this approach approximates the likelihood with an exponential family, as in fact it assumes  $p(x|\theta) \propto \hat{r}(x, \theta) = \exp(\hat{\beta}(\theta)^T \psi(x))$ , for some coefficients  $\hat{\beta}$  determined by data; it is thus a more general likelihood assumption than SL (as the normal distribution belongs to the exponential family).

**Remark 1 (Reducing the number of simulations.)** *Several approaches have attempted to reduce the number of simulations required in SL, RE and ABC. Specifically, Gaussian Processes can be exploited to replace evaluations of simulation-based quantities with emulated values, while at the same time providing an uncertainty quantification used to guide the next model simulations; that has been done for both ABC (Meeds and Welling, 2014; Wilkinson, 2014; Gutmann et al., 2016; Järvenpää et al., 2019; Jarvenpaa et al., 2020) and SL (Meeds and Welling, 2014; Wilkinson, 2014; Moores et al., 2015; Järvenpää et al., 2021). These approaches all consider a fixed observation  $x$  and emulate over parameter values.*

*In Hermans et al. (2020), a classifier more powerful than logistic regression is used to learn the likelihood ratio estimate for all  $(x, \theta)$  using parameter-simulation pairs (similarly to what we propose in Section. 3.3), amortizing therefore RE with respect to the observation.*

### 3. Likelihood approximation with the exponential family

In this Section, we first introduce the parametric family which we use to approximate the likelihood. Then, we describe Score Matching (SM) and Sliced Score Matching (SSM), which allow us to fit distributions with unknown normalizing constants to data. Finally, we discuss how to extend SM and SSM to the setting of conditional densities, to obtain a likelihood approximation valid for all data and parameter values.

#### 3.1 Conditional exponential family

A probability distribution belongs to the exponential family if it has a density of the following form:<sup>3</sup>

- 
2. Simulating from the marginal can be done by drawing  $\theta^j \sim p(\theta)$ ,  $x^j \sim p(x|\theta^j)$  and discarding  $\theta^j$ .
  3. As we are concerned with continuous random variables, across the work we use the Lebesgue measure as a base measure, without explicitly referring to it.

$$p(x|\theta) = \frac{e^{\eta(\theta)^T f(x)} h(x)}{Z(\theta)}, \quad (2)$$

where  $x \in \mathcal{X}$ ,  $\theta \in \Theta$ . Here,  $f : \mathcal{X} \rightarrow \mathbb{R}^{d_s}$  is a function of the data (sufficient statistics),  $\eta : \Theta \rightarrow \mathbb{R}^{d_s}$  is a function of the parameters (natural parameters) and  $h(x) : \mathcal{X} \rightarrow \mathbb{R}$  is a scalar function of data (base measure). The normalizing constant  $Z(\theta) = \int_{\mathcal{X}} \exp\{\eta(\theta)^T f(x)\} h(x) dx$  is intractable and assumed to be finite  $\forall \theta \in \Theta$ ; we will discuss later (Sec. 3.3) why this assumption is not an issue.

Across this work, we refer to Eq. (2) as *conditional exponential family* to stress that we learn an approximation valid for all  $x$ 's and  $\theta$ 's by selecting functions  $f$ ,  $\eta$  and  $h$ .

In the following, we rewrite Eq. (2) as:  $p(x|\theta) = \exp(\bar{\eta}(\theta)^T \bar{f}(x)) / Z(\theta)$ , where  $\bar{\eta}(\theta) = (\eta(\theta), 1) \in \mathbb{R}^{d_s+1}$  and  $\bar{f}(x) = (f(x), \log h(x)) \in \mathbb{R}^{d_s+1}$ . For simplicity, we will drop the bar notation, using as convention that  $\eta(\theta)$  contains the natural parameters of the exponential family plus an additional constant term, and that the last component of  $f(x)$  is  $\log h(x)$ .

**Neural conditional exponential family.** Let now  $f_w$  and  $\eta_w$  denote two Neural Networks (NNs) with collated weights  $w$  (in practice the two NNs do not share parameters); then, the neural conditional exponential family is defined as:

$$p_w(x|\theta) = \frac{e^{\eta_w(\theta)^T f_w(x)}}{Z_w(\theta)}. \quad (3)$$

**Remark 2** *Ratio Estimation (Sec. 2.2) parametrizes the likelihood as an exponential family with user-specified statistics and natural parameters  $\beta$  independently learned for each  $\theta \in \Theta$ . In contrast, here we learn both the  $f_w(x)$  and the  $\eta_w(\theta)$  transformations over all  $\mathcal{X}$  and  $\Theta$  at once by selecting the best  $w$ .*

**Remark 3 (Identifiability)** *For a family of distributions indexed by a parameter  $\phi$ , identifiability means that  $p_\phi(x|\theta) = p_{\phi'}(x|\theta) \forall x, \theta \implies \phi = \phi'$ . The weights  $w$  in the neural conditional exponential family are not identifiable for two reasons: first, NNs have many intrinsic symmetries. Secondly, replacing  $f_w$  in Eq. (3) with  $A \cdot f_w$  and  $\eta_w$  with  $(A^T)^{-1} \cdot \eta_w$  does not change the probability density. In Khemakhem et al. (2020), two suitable concepts of function identifiability up to some linear transformations are defined.  $f_w$  and  $\eta_w$  in Eq. (3) are identifiable according to those definitions, under strict conditions on the architectures of the Neural Networks parametrizing them (Khemakhem et al., 2020). Moreover, they empirically verify that NNs not satisfying the above assumptions result in approximately identifiable  $f_w$  and  $\eta_w$ , according to their definition. More details are given in Appendix B.1.*

**Remark 4 (Universal approximation)** *Using larger  $d_s$  in Eq. (2) increases the expressive power of the approximating family. Khemakhem et al. (2020) proved that, by considering a freely varying  $d_s$  and generic  $f$  and  $\eta$ , the conditional exponential family has universal approximation capabilities for the set of conditional probability densities; we give more details in Appendix B.2. This result does not consider the practicality of fitting the approximating family to data, which arguably becomes harder when  $d_s$  increases.*

### 3.2 Score Matching

In order to fit a parametric density  $p_w$  to data, the standard approach is finding the Maximum Likelihood Estimator (MLE) for  $w$ ; in the limit of infinite data, that corresponds to minimizing the Kullback-Leibler divergence from the data distribution. MLE requires however the normalizing constant of  $p_w$  to be known, which is not the case for our approximating family (Eq. 3).

Score Matching (SM, Hyvärinen 2005) is a possible way to bypass the intractability of the normalizing constant. In this Subsection, we review SM for unconditional densities, discuss a faster version and provide an extension for distributions with bounded domain.

**The original Score Matching (SM)** Let us discard now the conditional dependency on  $\theta$  and, following Hyvärinen (2005), consider a random variable  $X$  distributed according to  $p_0(x)$ . We want to use samples from  $p_0$  to fit a generic model  $p_w(x) = \tilde{p}_w(x)/Z_w$ , where  $\tilde{p}_w$  is unnormalized and  $Z_w$  is intractable.

**Definition 5** *Score Matching (SM) corresponds to finding:*

$$\arg \min_w D_F(p_0 \| p_w),$$

where the Fisher Divergence  $D_F$  is defined as:

$$D_F(p_0 \| p_w) = \frac{1}{2} \int_{\mathcal{X}} p_0(x) \|\nabla_x \log p_0(x) - \nabla_x \log p_w(x)\|^2 dx. \quad (4)$$

The Fisher divergence depends only on the logarithmic derivatives  $\nabla_x \log p_0(x)$  and  $\nabla_x \log p_w(x)$ , which are termed *scores*<sup>4</sup>; computing  $D_F$  does not require therefore knowing the normalizing constant, as in fact:

$$\nabla_x \log p_w(x) = \nabla_x (\log \tilde{p}_w(x) - \log Z_w) = \nabla_x \log \tilde{p}_w(x).$$

Nevertheless, a Monte Carlo estimate of  $D_F$  in Eq. (4) using samples from  $p_0$  would require knowing the logarithmic gradient of  $p_0$ ; for this reason, Eq. (4) is usually termed *implicit* Fisher divergence. If  $\mathcal{X} = \mathbb{R}^d$  and under mild conditions, Theorem 1 in Hyvärinen (2005) obtains an equivalent form for  $D_F$  in which  $p_0$  only appears as the distribution over which expectation needs to be computed. We give here a similar result holding for more general  $\mathcal{X} = \bigotimes_{i=1}^d (a_i, b_i)$ , where  $a_i, b_i \in \mathbb{R} \cup \{\pm\infty\}$ ; specifically, we consider the following assumptions:

**A1**  $p_0(x) \frac{\partial}{\partial x_i} \log p_w(x) \rightarrow 0$  when  $x_i \searrow a_i$  and  $x_i \nearrow b_i, \forall w, i,$

**A2**  $\mathbb{E}_{p_0} [\|\nabla_x \log p_0(X)\|^2] < \infty, \mathbb{E}_{p_0} [\|\nabla_x \log p_w(X)\|^2] < \infty, \forall w,$

**A3**  $\mathbb{E}_{p_0} \left| \frac{\partial^2}{\partial x_i \partial x_j} \log p_w(X) \right| < \infty, \forall w, \forall i, j = 1, \dots, d.$ <sup>5</sup>

4. In most of the statistics literature, *score* usually refer to the derivative of the log-likelihood with respect to the parameter; here, the nomenclature is slightly different.

5. We remark that this assumption was not present in Hyvärinen (2005), but it is necessary to apply Fubini-Tonelli theorem in the proof (see Appendix C.1), as already discussed in Yu et al. (2019).

With the above, we can state the following:

**Theorem 6** *Let  $\mathcal{X} = \bigotimes_{i=1}^d (a_i, b_i)$ , where  $a_i, b_i \in \mathbb{R} \cup \{\pm\infty\}$ . If  $p_0(x)$  is differentiable and  $p_w(x)$  doubly differentiable over  $\mathcal{X}$ , then, under assumptions **A1**, **A2**, **A3**, the objective in Eq. (4) can be rewritten as:*

$$D_F(p_0||p_w) = \int_{\mathcal{X}} p_0(x) \sum_{i=1}^d \left[ \frac{1}{2} \left( \frac{\partial \log p_w(x)}{\partial x_i} \right)^2 + \left( \frac{\partial^2 \log p_w(x)}{\partial x_i^2} \right) \right] dx + C. \quad (5)$$

Here,  $C$  is a constant w.r.t.  $p_w$  and  $x_i$  is the  $i$ -th coordinate of  $x$ .

A proof is given in Appendix C.1. We refer to Eq. (5) as the *explicit* Fisher Divergence, as it is now immediate to build an unbiased Monte Carlo estimate using samples from  $p_0$ .

From Eq. (4), it is clear that  $D_F(p_0||p_w) \geq 0 \forall w$ , and that choosing  $p_w = p_0$  implies  $D_F(p_0||p_w) = 0$ ; however, the converse can not be said in general unless  $p_0$  is supported over all  $\mathcal{X}$ , as stated in the following theorem:

**Theorem 7 (Theorem 2 in Hyvärinen, 2005)** *Assume  $p_0(x) > 0 \forall x \in \mathcal{X}$ . Then:*

$$D_F(p_0||p_w) = 0 \iff p_0(x) = p_w(x) \forall x \in \mathcal{X}.$$

We prove Theorem 7 in Appendix C.2. If  $p_0(x)$  is zero for some  $x \in \mathcal{X}$ , there could be a distribution  $p_w$  such that  $D_F(p_0||p_w) = 0$  even if  $p_w \neq p_0$  (as discussed in Appendix C.6).

**Sliced Score Matching (SSM)** In the explicit Fisher Divergence (Eq. 5), the second derivatives of the log-density with respect to all components of  $x$  are required. When using the neural conditional exponential family (Eq. 3), this amounts to evaluating the second derivatives of  $f_w$  with respect to its input (see Appendix C.10). Practically, automatic differentiation libraries to obtain the derivatives effortlessly; however, the computational cost of the second derivatives is substantial, as it requires a number of forward and backward passes proportional to the dimension of the data  $x$  (see Appendix E.1)<sup>6</sup>.

Some approaches to reducing the computational burden have been proposed; we review some in Section 6. Here, we consider Sliced Score Matching (SSM, Song et al. 2020), which considers projections of the scores on random directions; matching the projections on *all* random directions ensures the two distributions are identical (under the same conditions as the original SM). More precisely, let  $v \in \mathcal{V} \subseteq \mathbb{R}^d$  be a noise vector with a distribution  $q$ ; SSM is defined as follows:

**Definition 8** *Sliced Score Matching (SSM) corresponds to finding:*

$$\arg \min_w D_{FS}(p_0||p_w),$$

where the Sliced Fisher Divergence is defined as:

$$D_{FS}(p_0||p_w) = \frac{1}{2} \int_{\mathcal{V}} q(v) \int_{\mathcal{X}} p_0(x) (v^T \nabla_x \log p_0(x) - v^T \nabla_x \log p_w(x))^2 dx dv. \quad (6)$$

6. Computing the derivatives during the forward pass offers some speedup, as automatic differentiation repeats some computations several times. However, this approach requires custom NN implementation (see Appendix E.1.1) and is thus not scalable to complex architectures. More importantly, the computational gain is limited with respect to what is achieved by the Sliced SM version introduced next.



We will require the noise distribution  $q(v)$  to satisfy the following Assumption:

**A4** For the random vector  $V \sim q$ , the covariance matrix  $\mathbb{E}[VV^T] \succ 0$  is positive definite and  $\mathbb{E}\|V\|_2^2 < \infty$ .

As for SM, we can obtain an *explicit* formulation from the *implicit* one in Eq. (6). This is done in the following Theorem, which we prove for convenience in Appendix C.3:

**Theorem 9 (Theorem 1 in Song et al. (2020))** *Let  $\mathcal{X} = \bigotimes_{i=1}^d (a_i, b_i)$ , where  $a_i, b_i \in \mathbb{R} \cup \{\pm\infty\}$ . If  $p_0(x)$  is differentiable and  $p_w(x)$  doubly differentiable over  $\mathcal{X}$ , then, under assumptions **A1**, **A2**, **A3**, **A4**, the objective in Eq. (6) can be rewritten as:*

$$D_{FS}(p_0\|p_w) = \int_{\mathcal{V}} q(v) \int_{\mathcal{X}} p_0(x) \left[ v^T (\mathbf{H}_x \log p_w(x)) v + \frac{1}{2} (v^T \nabla_x \log p_w(x))^2 \right] dx dv + C, \quad (7)$$

where  $\mathbf{H}_x \log p_w(x)$  denotes the Hessian matrix of  $\log p_w(x)$  with respect to components of  $x$  and  $C$  is a constant w.r.t.  $p_w$ .

Assumption **A4** is satisfied by, among others, Gaussian and Rademacher random variables (Song et al., 2020). Additionally, these two distributions allow to compute explicitly the expectation with respect to  $v$  of  $(v^T \nabla_x \log p_w(x))^2$  in Eq. (7), which leads to:

$$D_{FS}(p_0\|p_w) = \int_{\mathcal{X}} p_0(x) \left\{ \int_{\mathcal{V}} q(v) [v^T (\mathbf{H}_x \log p_w(x)) v] dv + \frac{1}{2} \|\nabla_x \log p_w(x)\|_2^2 \right\} dx + C; \quad (8)$$

in Song et al. (2020), the Monte Carlo estimate of the latter expression is found to perform better than the one for Eq. (7); across this work, we will therefore consider Eq. (8) with the Rademacher noise when using SSM.

Analogously to SM, a non-negative  $p_0$  ensures that the sliced Fisher divergence is zero if and only if  $p_w = p_0$ :

**Theorem 10 (Lemma 1 in Song et al. (2020))** *Assume Assumption **A4** holds and that  $p_0(x) > 0 \forall x \in \mathcal{X}$ . Then:*

$$D_{FS}(p_0\|p_w) = 0 \iff p_0(x) = p_w(x) \forall x \in \mathcal{X}.$$

The above result is proven in Appendix C.4.

SSM has a reduced computational cost with respect to SM when automatic differentiation is used. In fact, computing the second derivatives in Eq. (5) requires  $d$  backward propagations, while the quadratic form involving the Hessian in Eq. (8) only requires two backward propagations independently on the dimension of  $x$  (see Appendix E.1).

**SM and SSM over transformed domain.** If the domain  $\mathcal{X}$  is unbounded ( $a_i = -\infty$ ,  $b_i = +\infty \forall i$ ), **A1** is usually satisfied. Instead, it is easy for that assumption to be violated if  $\mathcal{X}$  is bounded: for instance, if  $p_0(x)$  converges to a constant at the boundary of  $\mathcal{X}$ , **A1** requires  $\nabla_x \log p_w(x)$  to go to 0. Further, if  $p_0(x)$  diverges,  $\nabla_x \log p_w(x)$  has to converge to 0 faster than some rate, which is in general a strong requirement.

To apply SM to distributions with bounded domain  $\mathcal{X}$  under looser conditions, we then transform  $\mathcal{X}$  to  $\mathcal{Y} = \mathbb{R}^d$  with a suitable bijection mapping  $t$ ; this creates a new random variable  $Y = t(X)$ , with distributions  $p_0^Y(y)$  and  $p_w^Y(y)$  induced by  $p_0$  and  $p_w$  on  $\mathcal{X}$ .

**Definition 11** We define respectively as *Transformed Score Matching (TransSM)* and *Transformed Sliced Score Matching (TransSSM)* the procedures:

$$\arg \min_w D_F(p_0^Y \| p_w^Y) \quad \text{and} \quad \arg \min_w D_{FS}(p_0^Y \| p_w^Y).$$

TransSM and TransSSM enjoy similar properties as SM and SSM, as stated in the following Theorem, which mirrors Theorem 7 and Theorem 10:

**Theorem 12** Let  $Y = t(X) \in \mathcal{Y}$  for a bijection  $t$ . Assume Assumption **A4** is satisfied,  $p_0(x) > 0 \forall x \in \mathcal{X}$ , and let  $D$  denote either  $D_F$  or  $D_{FS}$ ; then:

$$D(p_0^Y \| p_w^Y) = 0 \iff p_w(x) = p_0(x) \forall x \in \mathcal{X}.$$

We prove this Theorem (in Appendix C.5) by relying on the equivalence between distributions for the random variables  $Y$  and  $X$ .

Motivated by Theorem 12, across this work we will apply TransSM and TransSSM whenever  $\mathcal{X}$  is bounded; precisely, we adopt the same bijections as in the `Stan` package for statistical modeling (Appendix C.8, Carpenter et al. 2017).

**Remark 13** Another way to extend SM to distributions with bounded support involves multiplying the scores in the implicit Fisher divergence by a correction factor that allows to obtain an explicit form under looser assumptions. This *Corrected SM (CorrSM)* approach was first proposed for non-negative random variables in Hyvärinen (2007), which also remarked how CorrSM and TransSM are equivalent (Appendix C.7). TransSM is however more practically viable, as a single SM implementation is needed, while CorrSM requires separate implementations for different kinds of domains. Additionally, the transformations can be straightforwardly applied to SSM, while (to the best of our knowledge) a correction approach for SSM has not yet been proposed.

**Remark 14** Across this work we restrict to domains  $\mathcal{X}$  defined by independent constraints across the coordinates, i.e.  $\mathcal{X} = \bigotimes_{i=1}^d (a_i, b_i)$ . However, TransSM, TransSSM and CorrSM can be potentially applied to distributions with more general supports. We briefly review this and other extensions in Appendix C.9.

### 3.3 Score Matching for conditional densities

We now go back to conditional densities  $p_0(\cdot|\theta)$  and define an expected Fisher divergence (Arbel and Gretton, 2018) by considering  $\theta \sim \pi(\theta)$ :

$$\begin{aligned} D_F^E(p_0 \| p_w) &= \int_{\Theta} \pi(\theta) D_F(p_0(\cdot|\theta) \| p_w(\cdot|\theta)) d\theta \\ &= \frac{1}{2} \int_{\Theta} \int_{\mathcal{X}} p_0(x|\theta) \pi(\theta) \|\nabla_x \log p_0(x|\theta) - \nabla_x \log p_w(x|\theta)\|^2 dx d\theta. \end{aligned} \tag{9}$$

Analogously, we define an expected sliced Fisher divergence  $D_{FS}^E(p_0 \| p_w)$  from  $D_{FS}$ .

Note that  $D_F^E(p_0 \| p_w) \geq 0$ ; moreover, under the assumption that  $p_0(x|\theta)$  is supported on the whole domain  $\mathcal{X} \forall \theta$ , the above objective is equal to 0 if and only if  $p_0(\cdot|\theta) = p_w(\cdot|\theta)$

$\pi(\theta)$ -almost everywhere (Arbel and Gretton, 2018). In fact,  $D_F(p_0(\cdot|\theta)||p_w(\cdot|\theta)) \geq 0 \forall \theta$ , with equality holding if and only if the two conditional distributions are the same, as long as they are both supported on the whole  $\mathcal{X}$  (by Theorem 7). Exploiting Theorem 10, a similar result can be obtained for  $D_{FS}^E$ .

Requiring  $p_0(x|\theta) > 0 \forall x \in \mathcal{X}, \forall \theta \in \Theta$  means that the model is capable of generating all possible values of  $x \in \mathcal{X}$  with non-zero probability for all  $\theta \in \Theta$ ; otherwise, there may be distributions  $p_w \neq p_0$  minimizing the objective (see Appendix C.6).

Analogously to Eqs. (5) and (7), we can obtain explicit formulations<sup>7</sup> of  $D_F^E$  and  $D_{FS}^E$ :

$$D_F^E(p_0||p_w) = \underbrace{\int_{\Theta} \int_{\mathcal{X}} p_0(x|\theta) \pi(\theta) \sum_{i=1}^d \left[ \frac{1}{2} \left( \frac{\partial \log p_w(x|\theta)}{\partial x_i} \right)^2 + \left( \frac{\partial^2 \log p_w(x|\theta)}{\partial x_i^2} \right) \right]}_{J(w)} dx d\theta + C, \quad (10)$$

$$D_{FS}^E(p_0||p_w) = \underbrace{\int_{\mathcal{V}} \int_{\Theta} \int_{\mathcal{X}} q(v) p_0(x|\theta) \pi(\theta) \left[ v^T (\mathbf{H}_x \log p_w(x|\theta)) v + \frac{1}{2} \|\nabla_x \log p_w(x|\theta)\|_2^2 \right]}_{J_S(w)} dx d\theta dv + C,$$

where  $C$  denotes constants with respect to  $w$ . For these two expressions, it is immediate to obtain Monte Carlo estimates using samples  $(\theta^{(j)}, x^j)_{j=1}^N, \theta^{(j)} \sim \pi$  and  $x^{(j)} \sim p(\cdot|\theta^{(j)})$ , and draws from the noise model  $\{v^{(j,k)}\}_{1 \leq j \leq N, 1 \leq k \leq M}$ :

$$\begin{aligned} \hat{J}(w) &= \frac{1}{N} \sum_{j=1}^N \left[ \sum_{i=1}^d \left( \frac{1}{2} \left( \frac{\partial \log p_w(x^{(j)}|\theta^{(j)})}{\partial x_i} \right)^2 + \left( \frac{\partial^2 \log p_w(x^{(j)}|\theta^{(j)})}{\partial x_i^2} \right) \right) \right], \\ \hat{J}_S(w) &= \frac{1}{NM} \sum_{j=1}^N \sum_{k=1}^M \left[ v^{(j,k),T} (\mathbf{H}_x \log p_w(x^{(j)}|\theta^{(j)})) v^{(j,k)} + \frac{1}{2} \|\nabla_x \log p_w(x^{(j)}|\theta^{(j)})\|_2^2 \right]. \end{aligned} \quad (11)$$

As we discussed before, computing the square bracket term in  $\hat{J}_S(w)$  only requires two backward propagations, while the square bracket term in  $\hat{J}(w)$  requires  $d$ . However,  $\hat{J}_S(w)$  sums over  $M \cdot N$  terms, while  $\hat{J}(w)$  sums over  $N$ . Nevertheless, when using  $\hat{J}_S(w)$  to train a neural network, good results can be obtained by using a single different noise realization for each training sample  $(\theta^{(j)}, x^j)$  at each epoch, thus leading to smaller computational cost with respect to  $\hat{J}(w)$  (Song et al., 2020).

7. Requiring  $p_0(x|\theta)$  to be differentiable and  $p_w(x|\theta)$  doubly differentiable over  $\mathcal{X}$  for all  $\theta$  and:

- A1.**  $p_0(x|\theta) \frac{\partial}{\partial x_i} \log p_w(x|\theta) \rightarrow 0$  for  $x_i \searrow a_i$  and  $x_i \nearrow b_i, \forall w, i, \theta$ ,
- A2.**  $\mathbb{E}_{p_0} [\|\nabla_x \log p_0(X|\theta)\|^2] < \infty, \mathbb{E}_{p_0} [\|\nabla_x \log p_w(X|\theta)\|^2] < \infty, \forall w, \theta$ ,
- A3.**  $\mathbb{E}_{p_0} \left| \frac{\partial^2}{\partial x_i \partial x_j} \log p_w(X|\theta) \right| < \infty, \forall w, \theta, \forall i, j = 1, \dots, d$ .

**Score Matching for conditional exponential family.** Both the implicit and explicit versions  $D_F^E(p_0||p_w)$  and  $D_{FS}^E(p_0||p_w)$  are well defined only if  $p_w$  is a proper density, which requires  $Z_w(\theta) < \infty$ . In practice, we are interested in:

$$\hat{w} = \arg \min \hat{J}(w) \quad \text{or} \quad \hat{w} = \arg \min \hat{J}_S(w). \quad (12)$$

When we compute  $\hat{J}(w)$  or  $\hat{J}_S(w)$ , we only need to evaluate  $f_w$  in a finite set of points  $\{x^{(j)}\}_{j=1}^N$  which belong to a bounded subset of  $\mathcal{X}$ , say  $A \subset \mathcal{X}$ . We can thus redefine the approximating family as follows:

$$p'_w(x|\theta) = \frac{\tilde{p}'_w(x|\theta)}{Z'_w(\theta)}, \quad \tilde{p}'_w(x|\theta) = \begin{cases} \exp(\eta_w(\theta)^T f_w(x)) & \text{if } x \in A, \\ g_w(x, \theta) & \text{otherwise,} \end{cases}$$

where  $g_w$  is such that  $Z'_w(\theta) < \infty$ , and can always be chosen such that  $\tilde{p}'_w(x|\theta)$  is a smooth and continuous function of  $x$ .

Replacing  $p_w$  in Eq. (11) with  $p'_w$  does not change the value of  $\hat{J}(w)$  and  $\hat{J}_S(w)$ ; however, as  $p'_w$  is normalized, Eqs. (9) and (10) are now well defined for all  $w$ . Additionally, we do not need to specify  $A$  or  $g_w$  explicitly as we never evaluate the normalizing constant  $Z'_w(\theta)$  (which depends on them). In the following we will thus use interchangeably  $p_w$  and  $p'_w$ .

**Remark 15 (Notation)** *For notational simplicity, in the rest of the work we drop the hat in  $\hat{w}$ , denoting by  $p_w$  the approximation obtained by one of the strategies in Eq. (12), and by  $f_w$  and  $\eta_w$  the corresponding sufficient statistics and natural parameters networks.*

## 4. Inference using the likelihood approximation

By fitting the neural conditional exponential family (Eq. 3) with SM or SSM to parameter-simulation pairs generated from the model, we obtain an approximation of the likelihood up to the normalization constant  $Z_w(\theta)$ : for each fixed  $\theta$ , the function  $x \mapsto \exp(\eta_w(\theta)^T f_w(x))$  is approximately proportional to  $p_0(x|\theta)$ . We exploit  $p_w$  in two ways: first, by using  $f_w$  as summaries in ABC; secondly, using the full approximation to draw samples from the posterior with MCMC for doubly intractable distributions. Both approaches are illustrated in Figure 1 and discussed next.

### 4.1 ABC with neural conditional exponential family statistics

The exponential family is the most general set of distributions with sufficient statistics of a given size (see Appendix A). Using  $f_w$  as summaries in ABC is therefore intuitively appealing:  $f_w$  represents in fact the sufficient statistics of the best exponential family approximation to  $p_0$ , according to the (sliced) Fisher divergence. If  $p_0$  belongs to the exponential family,  $D_F^E(p_0||p_w) = D_{FS}^E(p_0||p_w) = 0$  if  $f_w$  and  $\eta_w$  are sufficient statistics and natural parameters for  $p_0$ .

To use  $f_w$  in ABC, some practicalities are needed: first, we discard the last component of  $f_w$ , which represents the base measure  $h_w(x)$ . Secondly, as discussed in Section 3.1,  $f_w$  is identifiable only up to a scale parameter, so that the magnitude of the different components of  $f_w$  can vary significantly. Before using ABC, then, we rescale the different components

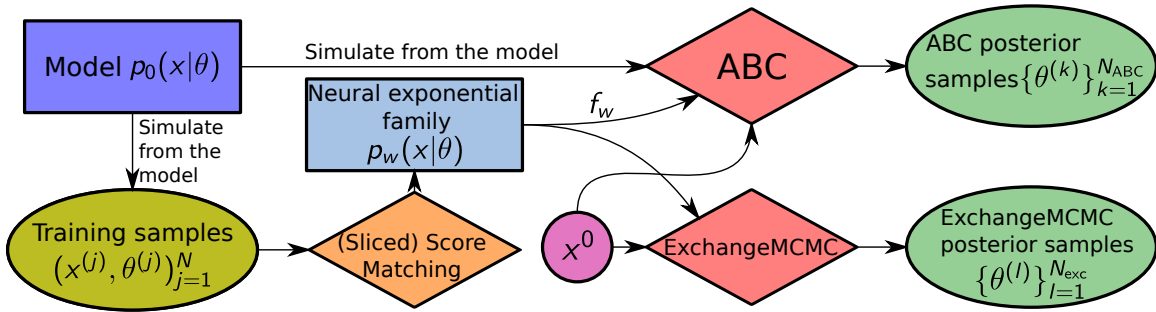


Figure 1: **Diagram showing the roles of the different components in the inference routines.** SSM or SM are used to train the neural exponential family on model simulations. Then, the sufficient statistics  $f_w$  are used in ABC or, alternatively, the neural exponential family is inserted in ExchangeMCMC. Notice how additional model simulation are needed for ABC but not for ExchangeMCMC.

by their standard deviation on new samples generated from the model, to prevent the ones with larger magnitude from dominating the ABC distance.

In the following, we will call the approach described here ABC with Score Matching statistics, for short ABC-SM, or ABC-SSM in the Sliced Score Matching case.

## 4.2 ExchangeMCMC with neural conditional exponential family

In contrast to other statistics learning methods (Jiang et al., 2017; Wiqvist et al., 2019; Pacchiardi et al., 2020; Åkesson et al., 2020), our technique provides a full likelihood approximation; it is therefore tempting to sample from the corresponding posterior directly, bypassing in this way ABC (and the additional model simulations required) altogether.

Unfortunately,  $p_w$  is known only up to the normalizing constant  $Z_w(\theta)$ ; therefore, standard MCMC cannot be directly exploited, and methods for *doubly-intractable* distributions are required (see Park and Haran, 2018 for a review). Here, we use ExchangeMCMC (Murray et al. (2012), Algorithm 1 in Appendix E.4), an MCMC where, for each proposed  $\theta'$ , a simulation from the distribution  $p_w(x|\theta')$  is used within a Metropolis-Hastings rejection step.

In our case, we cannot generate samples from  $p_w(\cdot|\theta')$  directly; to circumvent such issue, Murray et al. (2012) suggested to run an MCMC targeting  $p_w(\cdot|\theta')$  for each ExchangeMCMC step; if the chain is long enough, the last step can be considered as a draw from  $p_w$ . Empirical results (Caimo and Friel, 2011; Alquier et al., 2016; Everitt, 2012; Liang, 2010) show that a relatively small number of inner MCMC steps are enough for good performance and that initializing the inner chain at the observation improves convergence; we employ therefore this strategy in our work.

Further, a variant of ExchangeMCMC employs Annealed Importance Sampling to improve the mixing of the chain (Murray et al., 2012). Specifically, a sequence of  $K$  auxiliary variables are drawn from Metropolis-Hastings kernels targeting some intermediate distributions, and the ExchangeMCMC acceptance rate is modified accordingly. This approach, termed *bridging*, decreases the number of rejections in ExchangeMCMC due to poor auxiliary variables. See Appendix E.4 for more details.

In the following, we will refer to using ExchangeMCMC with our likelihood approximation as Exc-SM or Exc-SSM, according to whether SM or SSM are used to obtain the likelihood approximation. Exc-SM and Exc-SSM avoid additional model simulations (beyond the ones required to train  $p_w$ ) at the cost of running a nested MCMC. However, the number of steps in the inner MCMC required to obtain good performance increases with the dimension of  $\mathcal{X}$ . Exc-SM and Exc-SSM are therefore ideally applied to models with moderate dimension  $x$  (up to a few hundred), for which simulation is expensive. The same likelihood approximation can be used to perform inference on several observations, which makes the computational gain even greater in this case.

**Remark 16 (Model misspecification)** *The neural exponential family approximation is only valid close to where training data was distributed. Specifically, if  $x^0$  is far from that region,  $p_{\hat{w}}(x^0|\theta)$  may be assigned a large value rather than a (correct) small one. This could happen when the model  $p_0$  is unable to generate data resembling the observation for any parameter value, such that standard Bayesian inference in presence of the likelihood would also perform poorly. To get better inference in such scenarios, we could use the generalized posterior introduced in Matsubara et al. (2021), which is robust to outliers and suitable for doubly-intractable distributions (as the exponential family).*

*In Exc-SM and Exc-SSM, we may wonder whether running MCMC over  $x$  targeting  $p_w(x|\theta)$  for a fixed  $\theta$  can fail due to what we discussed above. We believe this is not the case: if the MCMC is initialized close to training data and  $p_w(\cdot|\theta)$  is a good representation of  $p_0(\cdot|\theta)$  in that region,  $p_w(\cdot|\theta)$  is small for values of  $x$  close to the boundary of the region where training data was distributed. Then, the MCMC is “trapped” inside that region and has no way of reaching regions of  $\mathcal{X}$  where  $p_w$  may behave irregularly.*

## 5. Simulation Studies

We perform simulation studies with our proposed approaches (Exc-SM, Exc-SSM, ABC-SM and ABC-SSM) and we compare with:

- ABC with statistics learned with the rejection approach discussed in Section 2.1 (Fearnhead and Prangle, 2012; Jiang et al., 2017), which we term ABC-FP.
- Population Monte Carlo (PMC, Cappé et al. 2004) with Ratio Estimation (PMC-RE).
- PMC with Synthetic Likelihood, using the robust covariance matrix estimator developed in Ledoit and Wolf (2004) to estimate  $\Sigma_\theta$ . We will denote this as PMC-SL.

Specifically, we consider three exponential family models and two time-series models (AR(2) and MA(2)) for which the exact likelihood is available, as well as a large-dimensional model with unknown likelihood (the Lorenz96 model, Lorenz 1996; Wilks 2005).

Exc-SM and Exc-SSM do not require additional simulations and run an MCMC, in contrast to sequential algorithms for the other methods, which we run with parallel computing. Comparing the computational cost is therefore not easy; in Appendix G.3, we report the number of simulations needed by the different methods to reach the same performance achieved by Exc-SM for the models with known likelihood.

**Choice of neural network architecture** In our exponential family approximation, we fix  $d_s$  to the number of parameters in each model. We added a Batch Normalization layer (see Appendix E.2) on top of the neural network representing  $\eta_w$  to reduce the unidentifiability introduced by the dot product between  $f_w(x)$  and  $\eta_w(\theta)$  (as discussed in Appendix C.10). For all techniques, we use  $10^4$  training samples. All NNs use SoftPlus nonlinearity (NNs using the more common ReLU nonlinearity cannot be used with SM and SSM as they have null second derivative with respect to the input).

For all models,  $\eta_w$  is represented by fully connected NNs. For the exponential family models, fully connected NNs are also used for  $f_w$  and  $s_\beta$ . For the time series and Lorenz96 models, we parametrize  $f_w$  and  $s_\beta$  with Partially Exchangeable Networks (PENs, Wiqvist et al. 2019). The output of an  $r$ -PEN is invariant to input permutations which do not change the probability density of data distributed according to a Markovian model of order  $r$  (see Appendix E.3). As AR(2) is a 2-Markovian model, we use a 2-PEN; similarly, a 1-PEN is used for the Lorenz96 model, which is 1-Markovian. Finally, we use a 10-PEN for the MA(2) model; albeit not being Markovian, Wiqvist et al. (2019) argued that MA(2) can be considered as “almost” Markovian so that the loss of information in imposing a PEN structure of high enough order is negligible. Further details are given in Appendix F.

**Choice of inferential parameters.** For ABC, we employ Simulated annealing ABC (SABC, Albert et al. 2015), which considers a set of particles and updates their position in parameter space across several iterations. We use 100 iterations and 1000 particles (posterior samples), corresponding to 1000 model simulations per iteration.

In Exc-SM and Exc-SSM, we run Exchange MCMC for 20000 steps, of which the first 10000 are burned-in. During burn-in, at intervals of 100 outer steps, we tune the proposal sizes according to the acceptance rate in the previous 100 steps. Our implementation of the Exchange algorithm is detailed in Appendix E.4.1. For the exponential family and time-series models, we test different numbers of inner MCMC steps, and eventually use 30 for the former and 100 for the latter, above which there was no noticeable performance improvement (more details in Appendix G.1.3 and Appendix G.2).

In PMC-SL and PMC-RE, we run the PMC algorithm with 10 iterations, 1000 posterior samples and respectively 100 (with SL) and 1000 (with RE) simulations per parameter value in order to estimate the approximate likelihood; such a large number of simulations (respectively  $10^5$  and  $10^6$  for each iteration) is required for the likelihood estimate to be numerically stable. For the exponential family models, we use the true sufficient statistics; for AR(2) and MA(2), we instead use autocovariances with lag 1 and 2 (as for instance in Marin et al. 2012). For PMC-RE, the cross-product of the statistics is also added to the list of statistics, as PMC-SL implicitly uses it.

## 5.1 Exponential family models

First, we consider three models for which a sample is defined as a 10-dimensional independent and identical distributed (i.i.d.) draw from either a Gaussian, Gamma or Beta distribution (all belonging to the exponential family). We put uniform priors on the parameters, with bounds given in Table 1. These models have different data ranges: unbounded for Gaussian, and respectively lower bounded by 0 and bounded in  $[0, 1]$  for Gamma and

<i>Model</i>	Beta		Gamma <sup>8</sup>		Gaussian		AR(2)		MA(2)		Lorenz96			
<i>Parameter</i>	$\alpha$	$\beta$	$k$	$\theta$	$\mu$	$\sigma$	$\theta_1$	$\theta_2$	$\theta_1$	$\theta_2$	$b_0$	$b_1$	$\sigma_e$	$\phi$
<i>Lower bound</i>	1	1	1	1	-10	1	-1	-1	-1	0	1.4	0	1.5	0
<i>Upper bound</i>	3	3	3	3	10	10	1	0	1	1	2.2	1	2.5	1

Table 1: **Bounds of uniform priors for the considered models.**

Beta. Therefore, we directly apply SM and SSM to the Gaussian model, while TranSM and TranSSM are applied to Gamma and Beta.

**Inferred sufficient statistics and natural parametrization.** With these models, our exponential family approximation is well specified. Thus, we compare the learned  $f_w$  and  $\eta_w$  with the exact sufficient statistics and natural parameters using the Mean Correlation Coefficient (MCC, Appendix B.1.1) metric, which ranges in  $[0, 1]$  and measures how well a multivariate function is recovered. We report results obtained with SM in Table 2; the values are close to 1, indicating that the sufficient statistics and natural parameters are recovered quite well by our method. In Appendix G.1, we report values corresponding to SSM (Table 9), as well as comparisons of exact and learned embeddings in Figures 9 and 10.

<i>Model</i>	<i>MCC weak in</i>	<i>MCC weak out</i>	<i>MCC strong in</i>	<i>MCC strong out</i>
Beta (statistics)	0.964	0.958	0.723	0.723
Beta (nat. par.)	0.990	0.991	0.807	0.812
Gamma (statistics)	0.911	0.924	0.894	0.883
Gamma (nat. par.)	0.967	0.967	0.872	0.873
Gaussian (statistics)	0.944	0.937	0.808	0.824
Gaussian (nat. par.)	0.974	0.974	0.970	0.972

Table 2: **MCC for exponential family models between exact embeddings and those learned with SM.** We show weak and strong MCC values; MCC is between 0 and 1 and measures how well an embedding is recovered up to permutation and rescaling of its components (strong) or linear transformation (weak); the larger, the better. “in” denotes MCC on training data used to find the best transformation, while “out” denote MCC on test data. We used 500 samples in both training and test data sets.

**Inferred posterior distribution.** Figure 2 shows posteriors obtained with the proposed methods, for a possible observation for each model; we see that all approximate posteriors are remarkably close to the exact one; moreover, the results with SM and SSM are indistinguishable. For all methods, we also estimate the Wasserstein distance between true and approximate posterior and compute the Root Mean Squared Error (RMSE) between mean of the true and approximate posterior; this is repeated for 100 simulated observations, with

8. The scale parameter of the Gamma distribution is usually called  $\theta$ ; in contrast, in all the rest of this work we use  $\theta$  to denote all parameters. Please beware of the difference.



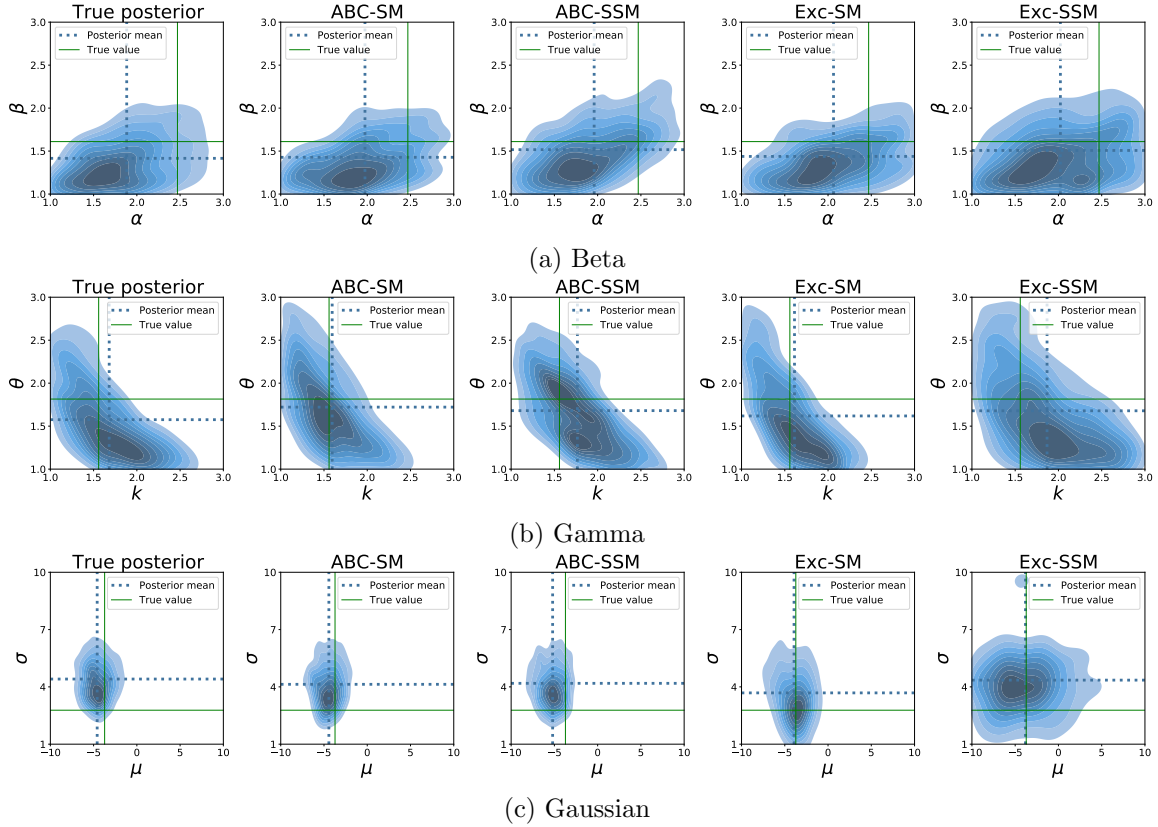


Figure 2: **True and approximate posteriors for exponential family models**, for a single observation per model. Dashed line represents posterior mean, while green solid line represents the exact parameter value.

results reported in Figure 3. ABC-FP is here the worst method; additionally, ABC-SM and ABC-SSM perform similarly to PMC-SL and PMC-RE. Finally, Exc-SM and Exc-SSM are marginally worse.

## 5.2 Time series models

The Moving Average model of order 2, or MA(2), and the AutoRegressive model of order 2, or AR(2), are special cases of the ARMA time-series model. The MA(2) model is defined by:

$$X_1 = \xi_1, \quad X_2 = \xi_2 + \theta_1 \xi_1, \quad X_j = \xi_j + \theta_1 \xi_{j-1} + \theta_2 \xi_{j-2}, \quad j = 3, \dots, t,$$

while the AR(2) is defined as:

$$X_1 = \xi_1, \quad X_2 = \xi_2 + \theta_1 X_1, \quad X_j = \xi_j + \theta_1 X_{j-1} + \theta_2 X_{j-2}, \quad j = 3, \dots, t;$$

in both,  $\xi_j$ 's are i.i.d. standard normal error terms (unobserved). We take here  $t = 100$  and we put uniform priors on the parameters of the two models, with bounds given in




Figure 3: **Performance of the different techniques for exponential family models.** Wasserstein distance from the exact posterior and RMSE between exact and approximate posterior means are reported for 100 observations using boxplots. Boxes span from 1st to 3rd quartile, whiskers span 95% probability density region and horizontal line denotes median. The numerical values are not comparable across examples, as they depend on the range of parameters. Here, SL and RE used the true sufficient statistics.

Table 1. For these models, the true likelihood can be evaluated, but they do not belong to the exponential family<sup>9</sup>.

**Inferred posterior distribution.** Figure 4 shows the posterior obtained with our proposed methods, for a possible observation for each model; again, our approximations are close to the exact posterior, with Exc-SM and Exc-SSM leading to slightly broader posteriors. Again, we assess performance with the Wasserstein distance from the true posterior and the RMSE between the means of true and approximate posterior for all methods, over 100 fresh observations. The results are reported in Figure 5; here, ABC-FP is the best method, with ABC-SM and ABC-SSM marginally worse. Exc-SM and Exc-SSM follow and perform better or comparably to PMC-RE and PMC-SL.

### 5.3 Lorenz96 meteorological model

The Lorenz96 model (Lorenz, 1996) is a toy model of chaotic atmospheric behaviour, including interacting slow and fast variables. We use here a modified version (Wilks, 2005) where the effect of the fast variables on the slow ones is replaced by a stochastic effective term depending on a set of parameters. Specifically, this model is defined by the following coupled Differential Equations (DEs):

$$\frac{dy_k}{dt} = -y_{k-1}(y_{k-2} - y_{k+1}) - y_k + 10 - g(y_k, t; \theta); \quad k = 1, \dots, K,$$

where  $y_k(t)$  is the value at time  $t$  of the  $k$ -th variable and indices are cyclic (index  $K + 1$  corresponds to 1, and so on). The effective term  $g$  depends on  $\theta = (b_0, b_1, \sigma_e, \phi)$ , and is

9. More precisely, they cannot be written as an exponential family with embedding dimension fixed with data size; in fact, MA(2) can be written as a Gaussian distribution with  $t \times t$  covariance matrix, which is an exponential family whose embedding dimension increases with data size.

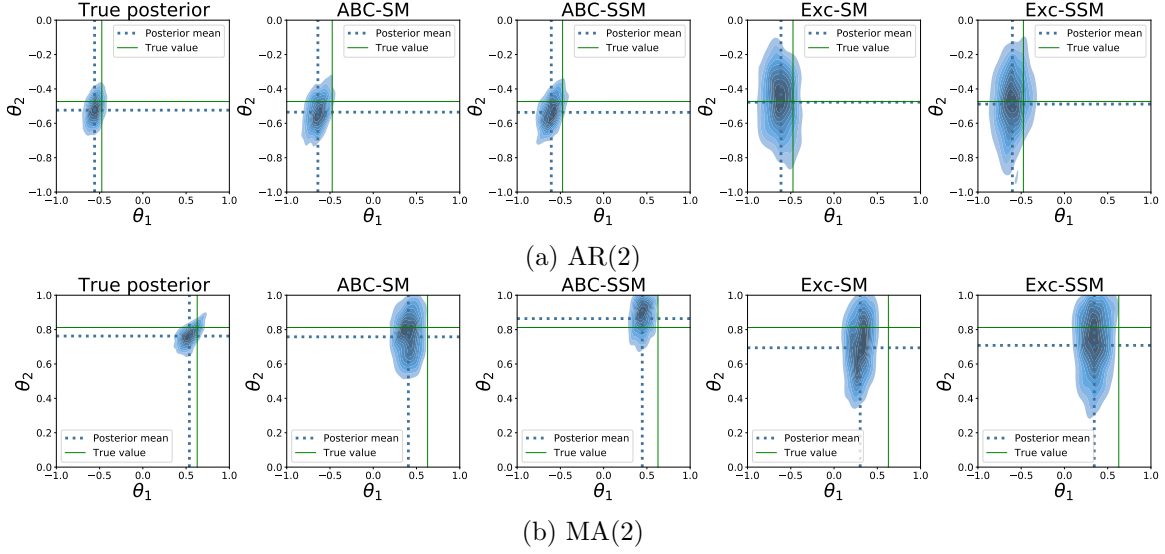


Figure 4: **True and approximate posteriors for AR(2) and MA(2) models**, for a single observation per model. Dashed line represents posterior mean, while green solid line represents the exact parameter value.

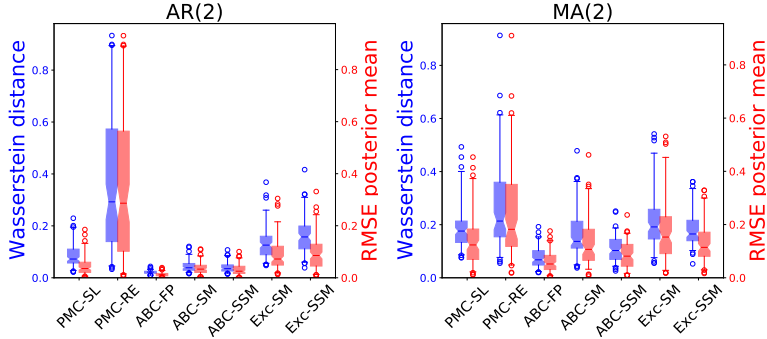


Figure 5: **Performance of the different techniques for AR(2) and MA(2) models**. Wasserstein distance from the exact posterior and RMSE between exact and approximate posterior means are reported for 100 observations using boxplots. Boxes span from 1st to 3rd quartile, whiskers span 95% probability density region and horizontal line denotes median.

defined upon discretizing the DEs with a timestep  $\Delta t$ :

$$g(y, t; \theta) = \underbrace{b_0 + b_1 y}_{\text{linear deterministic term}} + \overbrace{\phi \cdot r(t - \Delta t) + \sigma_e (1 - \phi^2)^{1/2} \eta(t)}^{\text{zero mean AR(1)}}$$

where  $\eta(t) \sim \mathcal{N}(0, 1)$ . We numerically integrate the model using 4th order Runge-Kutta method with  $\Delta t = 1/30$  in an interval  $[0, T]$ ; further, we fix the initial condition to a value  $y(0)$  which is independent on parameters. Here, the true likelihood is unaccessible, so that sampling from the exact posterior is impossible.

We consider the model in a small and large configuration: in the small one, we take  $K = 8$  and  $T = 1.5$ , which lead to 45 timesteps and a data dimension of  $d = 45 \cdot 8 = 360$ . In the large one, we take instead  $K = 40$  and  $T = 4$ , corresponding to 120 timesteps and a data dimension of  $d = 120 \cdot 40 = 4800$ .

**Inferred posterior distribution.** For both configurations, we perform inference with ABC-SSM and ABC-FP; additionally, we use Exc-SSM for the small configuration (as running the exchange algorithm in the large one is too costly). In ExchangeMCMC, we used 500 inner MCMC steps and 200 bridging steps for each outer step. Figure 6 reports posteriors for a single observation in both setups. The ABC-FP posterior is narrower than the one for ABC-SSM, but concentrated around similar parameter values. The posterior for Exc-SSM looks slightly different: it concentrates on similar parameter values as the other two for  $\theta_1$  and  $\theta_2$ , but is broader for  $\sigma_e$  and  $\phi$ .

**Posterior predictive validation.** We assess the performance of the posterior predictive distribution  $p(x|x^0) = \int p(x|\theta)\pi(\theta|x^0)d\theta$ , in which  $\pi(\theta|y)$  is the posterior obtained with one of the considered techniques. Specifically, we use the Kernel and Energy Scoring Rules (SRs) to evaluate how well the predictive  $p(x|x^0)$  predicts the observation  $x^0$ . A SR (Gneiting and Raftery, 2007) is a function of a distribution and an observation, and assesses the mismatch between them (the smaller, the better). More details on Scoring Rules are given in Appendix D, together with the definition of those used here.

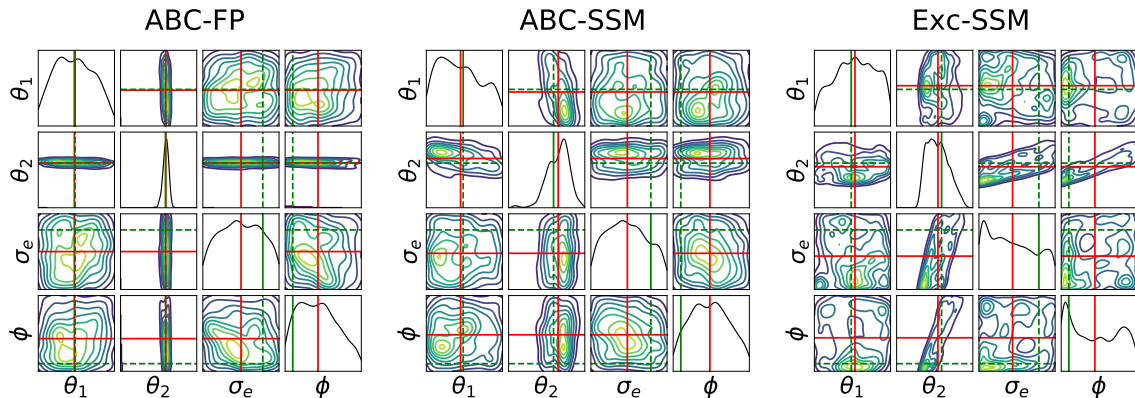
As the Lorenz96 model returns a (multivariate) time-series, we compare predictive distribution and observation at each timestep independently. We repeat this validation with 100 different observations and corresponding posteriors, for all techniques and both model configurations. We obtain cumulative scores by summing the scores over timesteps; median and some quantile values over the 100 observations are shown in Fig. 7. In both model configurations, the posterior predictive generated by ABC-SSM and Exc-SSM marginally outperform ABC-FP, according to both SRs. Score values at each timestep, together with more details on the validation technique, are reported in Appendix G.4.

## 6. Related works

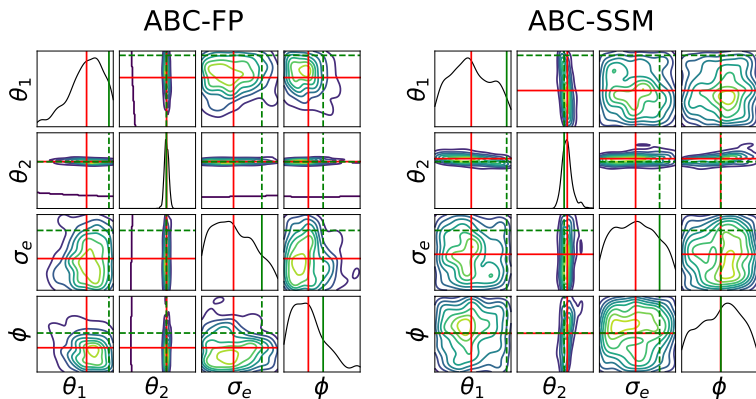
**ABC statistics with NN.** Using NNs for learning statistics for ABC has been previously suggested, first with the regression approach discussed in Sec. 2.1 (used in Jiang et al. 2017; Wiqvist et al. 2019; Åkesson et al. 2020). In Pacchiardi et al. (2020), a NN is trained so that the distance between pairs of simulated statistics best reproduces the distance between the corresponding parameter values. Chen et al. (2021) instead fit a NN by maximizing an estimate of the mutual information between statistics and parameters.

**ABC statistics from auxiliary models.** Some similarities with our work can be found in Gleim and Pigorsch (2013) and Ruli et al. (2016), which consider respectively an auxiliary model and a composite likelihood alongside the simulator model and obtain summary statistics from them; similarly, our approach can be seen as building an auxiliary exponential family model with easily accessible summary statistics.

**LFI with NNs.** The idea of sampling directly from the approximate posterior defined by  $p_w$  is related to a suite of LFI methods using NNs. A large part of these works use normaliz-



(a) Lorenz96 - Small configuration



(b) Lorenz96 - Large configuration

Figure 6: **Example of posterior inference for a single Lorenz96 observation with different approaches**, for both small (first row) and large (second row) configurations. In each panel, the diagonal plots represent the univariate marginal distributions, while the off-diagonal ones are bivariate density contour plots. Moreover, the green and red lines represent respectively exact parameter value and posterior mean. All axes span full prior range (see Table 1).

ing flows (NNs implementing invertible transformations, suitable for efficiently representing probability densities; see Papamakarios et al., 2021 for a review); specifically, Papamakarios et al. (2019); Lueckmann et al. (2019) use them to learn an approximation of the likelihood, while Papamakarios and Murray (2016); Lueckmann et al. (2017); Greenberg et al. (2019); Radev et al. (2020) learn the posterior. Most of the above approaches focus on inference for a single observation, tailoring the simulations to better approximate the posterior for the relevant parameter values at lower computational cost; instead, Radev et al. (2020) and Lueckmann et al. (2019) propose to amortize across observations, similarly to our approach. Besides normalizing flows, Klein et al. (2020) casts the LFI problem as a distributional regression one. Finally, Tabak et al. (2020) does not employ NNs but rather

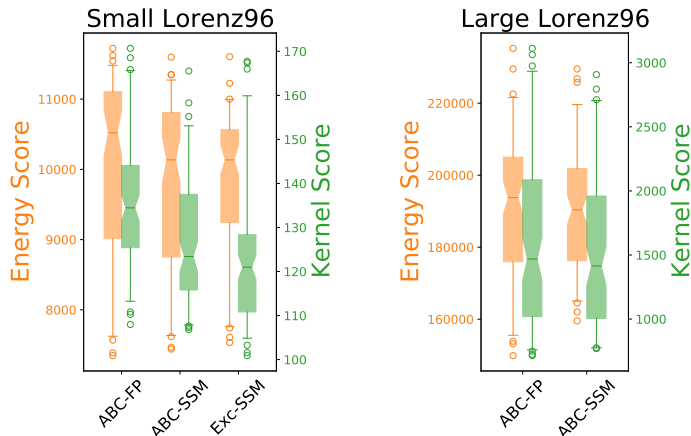


Figure 7: **Predictive performance of the different methods according to the Kernel and Energy Scoring Rules.** Each boxplot represents cumulative (i.e., summed over the time index) scoring rule value for a given method for the small (left) and large (right) Lorenz96 configuration. Boxes span from 1st to 3rd quartile, whiskers span 95% probability density region and horizontal line denotes median.

solves a Wasserstein barycenter problem to model conditional maps which allow sampling from the distribution of an observation conditional on some covariates.

**Fitting unnormalized models.** Several techniques besides Score Matching have been proposed for fitting unnormalized models: MCMC-MLE (Geyer, 1991) exploits MCMC to estimate the normalizing constant for different values of the parameter and uses that in MLE. Contrastive Divergence (CD, Hinton 2002) instead uses MCMC to obtain a stochastic approximation of the gradient of the log-likelihood; this requires a smaller number of MCMC steps with respect to MCMC-MLE, but the stochastic gradient estimate is biased. Minimum Probability Flow (MPF, Sohl-Dickstein et al. 2011) considers a dynamics from data to model distribution, and minimizes the Kullback Leibler divergence between the data distribution and the one obtained by running the dynamics for a short time; however, the efficacy of MPF depends significantly on the considered dynamics. Noise Contrastive Estimation (NCE, Gutmann and Hyvärinen 2012) converts the parameter estimation problem to ratio estimation between data distribution and a suitable noise distribution; in practice, NCE uses logistic regression to discriminate the observed data and data generated from the noise; in the loss, the normalizing constant appears explicitly and can be estimated independently. To be effective, NCE requires the noise distribution to overlap well with the data density while being easy to sample from and to evaluate, which is not easy to get. Finally, some works (Dai et al., 2019a,b) use the dual formulation of MLE to avoid estimating the normalizing constant at the price of introducing dual variables to be jointly estimated.

**Fitting unnormalized conditional models.** The approaches above running dynamics to estimate quantities (CD, MCMC-MLE, MPF) cannot easily be applied to the conditional setting. NCE has been instead used in Ton et al. (2021) with a single noise distribution for all values of the conditioning variable  $\theta$ ; it however requires an independent NN to parametrize

the normalizing constant  $Z(\theta)$ . SM can be instead easily applied to the conditional setting, as previously done in Arbel and Gretton (2018) and further demonstrated in our work. To the best of our knowledge, SSM was not applied to a conditional setting before, although the extension straightforwardly follows what is done for SM.

**Fast approximations of SM.** Besides SSM, some approximations to SM have been investigated, all of which only require first derivatives. For instance, Denoising Score Matching (Vincent, 2011) computes the Fisher divergence between the model and a Kernel Density Estimate of the data distribution, which is equal to a quantity independent on the second derivative. Alternatively, Kernel Stein Discrepancy (Liu and Wang, 2016) intrinsically depends on first derivatives only. For both techniques, however, several samples for each  $\theta$  would be required in the conditional setting. Finally, Wang et al. (2020) exploits a connection between the Fisher divergence and gradient flows in the 2-Wasserstein space to develop an approximation that only relies on first-order derivatives.

**Kernel Conditional Exponential Families (KCEFs).** In KCEFs (Arbel and Gretton, 2018), the summary statistics and natural parameters are functions in a Reproducing Kernel Hilbert Space, whose properties allow to evaluate the density although an infinite-dimensional embedding space is used (using the *kernel trick*). In Arbel and Gretton (2018), SM was used to fit KCEFs instead of our neural conditional exponential families. They build on Sriperumbudur et al. (2017), which first used SM to perform density estimation with (non-conditional) Kernel Exponential Families (KEFs). In Wenliang et al. (2019) NNs are used to parametrize the kernel in a KEF and trained with the SM loss. However, as KCEFs do not have finite-dimensional sufficient statistics, they are unsuitable for learning ABC statistics. Additionally, KCEFs have a worse complexity in terms of training dataset size with respect to NN-based methods. KCEFs have also been used with a dual MLE objective in Dai et al. (2019a).

## 7. Conclusions and extensions

We proposed a technique to approximate the likelihood using a neural conditional exponential family, trained via (Sliced) Score Matching to handle the intractable normalizing constant.

We tested this approximation in two setups: first, by using the exponential family sufficient statistics as ABC statistics, which is intuitively appealing as the exponential family is the largest class of distributions with fixed-size sufficient statistics. We empirically showed this to be comparable or outperform ABC with summaries built via the state-of-the-art regression approach (Fearnhead and Prangle, 2012; Jiang et al., 2017).

Secondly, we used MCMC for doubly-intractable distributions to sample from the posterior corresponding to the likelihood approximation, which we found to have performance comparable to the other approaches. This can be repeated for any new observation without additional model simulations, making it advantageous for expensive simulator models.

Our proposed direct sampling approach based on exponential family likelihood approximation and ExchangeMCMC could be improved as follows:

- we used ExchangeMCMC (Murray et al., 2012) to handle double intractability, but other algorithms could be more efficient, (for instance the one in Liang et al. (2016)),

which makes use of parallel computing). Alternatively, we could exploit the generalized posterior introduced in Matsubara et al. (2021), which allows standard MCMC to be used for double intractable distributions and is robust to outliers.

- To infer the posterior for a single observation, approximating the likelihood for all  $x$ 's and  $\theta$ 's as we do now is suboptimal. Similar performance may in fact be obtained with fewer simulations tailored to the observation. Sequential schemes implementing such ideas have been introduced for LFI using normalizing flows (see for instance Papamakarios et al. 2019) and could be extended to our setup.
- The motivation for the current work was learning ABC statistics, hence the exponential family formulation. However, the dot-product structure between  $f_w$  and  $\eta_w$  is not beneficial for the direct sampling approach. An *energy-based model*, which employs a single NN with input  $(x, \theta)$  in the exponent, may be more expressive and easier to train.
- An energy-based posterior approximation  $\pi_w(\theta|x)$  could be trained by minimizing the expectation over the data marginal  $p(x)$  of the (sliced) Fisher divergence with respect to the true posterior. This is complementary to the strategy employed in this work to fit  $p_w(x|\theta)$ . Interestingly,  $\pi_w(\theta|x)$  would be known up to a normalizing constant depending on  $x$  only, making use of standard MCMC possible.

## Acknowledgments

We acknowledge fruitful discussion with prof. Christian Robert, prof. Antonietta Mira and prof. Geoff Nicholls. LP is supported by the EPSRC and MRC through the OxWaSP CDT programme (EP/L016710/1), which also funds the computational resources used to perform this work. RD is funded by EPSRC (grant nos. EP/V025899/1, EP/T017112/1) and NERC (grant no. NE/T00973X/1).

All summary statistics learning strategies used in this manuscript (the regression approach and the exponential family with SM and SSM) have been implemented in the Python library ABCpy (Dutta et al., 2021b), which exploits Pytorch (Paszke et al., 2019) to train the NNs. ABCpy was also used to run PMC-SL, PMC-RE and the ABC experiments.



## Appendix A. Sufficient statistics

Consider a conditional probabilistic model  $p(x|\theta)$ ; moreover, abusing notation, we will also denote as  $p(x|t)$  the density  $p_{X|T}(X = x|T = t)$ , as well as  $p(x|t; \theta)$  the density  $p_{X|T, \Theta}(X = x|T = t; \Theta = \theta)$ , where here  $\Theta$  denotes the random variable which takes values  $\theta$ . Finally, we use  $\pi$  to denote distributions over the parameter values  $\theta$ ; specifically,  $\pi(\theta|x)$  denotes the standard posterior, and  $\pi(\theta|t)$  is an abuse of notation for the density  $\pi_{\Theta|T}(\Theta = \theta|T = t)$ .

**Definition 17** *A statistic  $t = t(x)$  is sufficient if  $p(x|t; \theta) = p(x|t)$ , where  $\theta$  is a parameter of the distribution. Alternatively, we have, in the Bayesian setting:*

$$\pi(\theta|x) = \pi(\theta|t(x)),$$

for any (non-degenerate) choice of prior distribution  $\pi(\theta)$ .

The existence of sufficient statistics implies a precise form of the distribution:

**Theorem 18 (Fisher-Neyman factorization theorem):** *A statistic is sufficient  $\iff p(x|\theta) = h(x)g(t(x)|\theta)$ , where  $h$  and  $g$  are non-negative functions.*

It is clear from the above theorem that  $f(x)$  in the exponential family is sufficient.

A stronger result regarding exponential family also exists, which goes under the name of Pitman–Koopman–Darmois theorem. This theorem says that the exponential family is the most general family of distributions for which there is a sufficient statistics whose size is fixed with the number of samples, provided that the domain of the probability distribution does not vary with the parameter  $\theta$  (Koopman, 1936).

## Appendix B. Properties of the conditional exponential family

### B.1 Linear identifiability

Let us consider the exponential family model  $p_w(x|\theta) = \exp(\eta_w(\theta)^T f_w(x))/Z_w(\theta)$ , where here  $\eta_w$  and  $f_w$  are not restricted to be Neural Networks. Khemakhem et al. (2020) studies the identifiability properties of the above family; specifically, they consider identifiability properties of the feature extractors  $\eta_w$  and  $f_w$ . Identifiability of representations is useful as it means that two different models from the above family learn similar representations when trained with different initialization on the same data set. In the framework of our likelihood-free inference task, it also means that if the true model belongs to the family we are considering, it is theoretically possible to recover the natural parameters and the true sufficient statistics.

If we consider now  $\eta_w$  and  $f_w$  to be Neural Networks, this causes problems as they are not identifiable in the standard sense. In fact, many different parameter configurations lead to the same function (as there are many symmetries in how the transformations in a Neural Network layer are defined). Therefore, Khemakhem et al. (2020) introduces, following previous works, two more general notions of identifiable representations. Subsequently, let  $\mathcal{W}$  denote the space of the possible Neural Network weights  $w$ .

**Definition 19 Weak identifiability (Section 2.2 in Khemakhem et al., 2020).** Let  $\sim_w^f$  and  $\sim_w^\eta$  be equivalence relations on  $\mathcal{W}$  defined as:

$$\begin{aligned} w \sim_w^f w' &\iff f_w(x) = Af_{w'}(x) + c \\ w \sim_w^\eta w' &\iff \eta_w(\theta) = B\eta_{w'}(\theta) + e \end{aligned}$$

where  $A$  and  $B$  are  $(d_s \times d_s)$ -matrices of rank at least  $\min(d_s, d)$  and  $\min(d_s, p)$  respectively, and  $c$  and  $e$  are vectors.

**Definition 20 Strong identifiability (Section 2.2 in Khemakhem et al., 2020).** Let  $\sim_s^f$  and  $\sim_s^\eta$  be equivalence relations on  $\mathcal{W}$  defined as:

$$\begin{aligned} w \sim_s^f w' &\iff \forall i, f_{i,w}(x) = a_i f_{\sigma(i),w'}(x) + c_i \\ w \sim_s^\eta w' &\iff \forall i, \eta_{i,w}(\theta) = b_i \eta_{\gamma(i),w'}(\theta) + e_i \end{aligned}$$

where  $\sigma$  and  $\gamma$  are permutations of  $[[1, n]]$ ,  $a_i$  and  $b_i$  are non-zero scalars and  $c_i$  and  $e_i$  are scalars.

Weak identifiability means that two parameters are equivalent if the corresponding feature extractors are the same up to linear transformation. Strong identifiability is a specific case of the weak one, in which the linear transformation is restricted to be a scaled permutation.

After introducing these concepts, Khemakhem et al. (2020) provides two theorems (Theorem 1 and 2) in which weak or strong identifiability of the representations are implied if different parameter values lead to same distributions, i.e.  $p_w(x|\theta) = p_{w'}(x|\theta) \forall x, \theta \implies w \sim_w^f w'$  or  $p_w(x|\theta) = p_{w'}(x|\theta) \forall x, \theta \implies w \sim_s^f w'$  (and similar for  $\sim_w^\eta$  and  $\sim_s^\eta$ ). These two results hold under some strict conditions on the functional form of the feature extractors  $f$  and  $\eta$  (concerning differentiability, rank of the Jacobian and other properties).

Then, they verify explicitly these conditions for a simple fully connected Neural Network architecture, in which they restrict the activation functions to be LeakyReLUs and the layer widths to be monotonically increasing or decreasing. Of course, this architecture is quite restrictive; it is in fact impossible to study theoretically the properties of more complex architectures and to show that they satisfy the necessary conditions for identifiability to hold. However, they show empirically that the conditional structure, even with more complex architectures, is helpful in increasing identifiability of the representations (as computed by the MCC, see Appendix B.1.1).

Therefore, even if the architectures which are used throughout this work do not satisfy the assumptions needed to explicitly show identifiability of the representations, we argue that the presence of such results is an hint towards the fact that identifiability (according to the above definitions) is actually achievable.

### B.1.1 MEAN CORRELATION COEFFICIENT (MCC)

In order to evaluate empirically whether the above identifiability results are satisfied, we need a way to measure similarity between two embeddings of the same set of data. By following Khemakhem et al. (2020), we describe here the Mean Correlation Coefficient

(MCC), which is a simple measure to do that. In the main text, we used this technique to assess how well our approximating family recovered the exact sufficient statistics and natural parameters, in the case where the data generating model  $p_0$  belongs to the exponential family.

In the following, we describe two versions of MCC, which are directly linked to the weak and strong identifiability definitions in Appendix B.1.

**Strong MCC.** Let us consider two sets of embeddings  $\{y_i \in \mathbb{R}^{d_s}\}_{i=1}^n$  and  $\{z_i \in \mathbb{R}^{d_s}\}_{i=1}^n$ , which can be thought of as samples from two multivariate random variables  $Y = t_1(X)$  and  $Z = t_2(X)$ , for two functions  $t_1, t_2$  and a random variable  $X$ ; in general, the order of the components of these two vectors is arbitrary, so that we cannot say, for instance, that the 1st component of  $Y$  corresponds to the 1st of  $Z$ . Then, MCC computes all the correlation coefficients between each pair of components of  $Y$  and  $Z$ . Next, it solves a linear sum assignment problem which identifies each component of  $Y$  with a component of  $Z$  aiming to maximize the sum of the absolute value of the corresponding correlation coefficient. In this way, it tries to couple the embeddings which are most linear one to the other. Finally, the MCC is computed as the mean of the absolute correlation coefficients after the right permutation of elements of the vector.

MCC is therefore a metric between 0 and 1 which measures how well each component of the original embedding (say,  $Y$ ) has been recovered independently by the other one (say,  $Z$ ). Moreover, as it relies on the correlation coefficient, it is not sensitive to rescaling or translation of each of the embeddings (in fact, the correlation coefficient of two univariate random variables is  $\pm 1$  when a deterministic linear relationship between the two exists, unless the relationship is perfectly horizontal or vertical).

Finally, in order to get a better estimate of the MCC, we can split the set of embeddings in two: one which will be used to determine the permutation and which will give an in-sample estimate of the MCC, and the other one to which the previous permutation will be applied and will give an out-of-sample estimate of MCC.

**Weak MCC.** The above definition of MCC aims to estimate how well each single element of the embedding is recovered, independently on the other; we call it, following Khemakhem et al. (2020), strong MCC. However, there may be cases where the recovered embedding is equal to the correct one up to a more generic linear transformation  $A$ . In that case, we would like to have a way to measure, up to a linear transformation, how far are two embeddings. Following Khemakhem et al. (2020), we therefore apply Canonical Correlation Analysis (CCA) Hotelling (1936) to learn  $A$ , and after that compute MCC, which we will call weak MCC. Specifically, CCA is a way to compute linear transformation  $A$  so that the correlations between corresponding components of  $A \cdot Y$  and  $Z$  is maximized. The so-defined weak MCC is therefore a number between 0 and 1 which measures how close is  $Y$  to  $Z$  after the best linear transformation is applied to  $Y$ . Similarly as before, a part of data is used to learn the best embedding; we can therefore use fresh samples to get an out-of-sample estimate along the in-sample one.

## B.2 Universal approximation capability

Khemakhem et al. (2020) provide a result (Theorem 3) in which they prove universal approximation capability of the conditional exponential family. Specifically, they show that, considering the dimension of the representations  $d_s$  as an additional parameter, it is possible to find an arbitrarily good approximation of any conditional probability density  $p_0(x|\theta)$ , provided that  $\mathcal{X}$  and  $\Theta$  are compact Hausdorff spaces. In practice, good approximation may be achieved with a value of  $d_s$  larger than  $d$  or  $p$ . Moreover, as remarked in the main text, this result is not concerned with the way the approximating family is fit; indeed, we expect that this task becomes (both statistically and computationally) more challenging when  $d_s$  increases.

## Appendix C. Some properties of Score Matching

### C.1 Proof of Theorem 6

Our Theorem 6 extends Theorem 1 in Hyvärinen (2005), that considers the case  $\mathcal{X} = \mathbb{R}^d$ , and which is recovered in the case  $a_i = -\infty$  and  $b_i = +\infty \forall i$ . Our proof follows Hyvärinen (2005), which however explicitly stated only Assumptions **A1** and **A2**. Following Yu et al. (2019), we add the additional Assumption **A3** which is required for Fubini-Tonelli theorem to apply.

In order to prove Theorem 6, Assumption **A3** can be weakened to  $\mathbb{E}_{p_0} \left| \frac{\partial^2}{\partial x_i^2} \log p_w(X) \right| < \infty, \forall w, \forall i = 1, \dots, d$ . We state the more general one in the main text as that allows Theorem 9 for SSM to be proved as well.

**Proof** Let  $s_0(x) = \nabla_x \log p_0(x)$  denote the score of the data distribution, and analogously  $s(x; w) = \nabla_x \log p_w(x)$ . Then, Eq. (4) can be rewritten as:

$$\begin{aligned} D_F(p_0 \| p_w) &= \frac{1}{2} \int_{\mathcal{X}} p_0(x) \|s_0(x) - s(x; w)\|^2 dx \\ &= \frac{1}{2} \int_{\mathcal{X}} p_0(x) [\|s_0(x)\|^2 + \|s(x; w)\|^2 - 2s_0(x)^T s(x; w)] dx \\ &= \underbrace{\frac{1}{2} \int_{\mathcal{X}} p_0(x) \|s_0(x)\|^2 dx}_C + \underbrace{\frac{1}{2} \int_{\mathcal{X}} p_0(x) \|s(x; w)\|^2 dx}_A - \underbrace{\int_{\mathcal{X}} p_0(x) s_0(x)^T s(x; w) dx}_B \end{aligned}$$

Note that we can split the integral into the three parts as the first two are assumed to be finite in **A2**; as a consequence,  $B$  is also finite thanks to  $|2ab| \leq a^2 + b^2$ .

The first element does not depend on  $w$ , so that we can safely ignore that. The second one appears as it is in the final Eq. (5). Therefore, we will focus on the last term, which we can write as:

$$B = - \sum_{i=1}^d \int_{\mathcal{X}} p_0(x) s_{0,i}(x) s_i(x; w) dx; \quad (13)$$

Let's now consider a single  $i$ , which we can rewrite as:

$$- \int_{\mathcal{X}} p_0(x) \frac{\partial \log p_0(x)}{\partial x_i} s_i(x; w) dx = - \int_{\mathcal{X}} \frac{p_0(x)}{p_0(x)} \frac{\partial p_0(x)}{\partial x_i} s_i(x; w) dx = - \int_{\mathcal{X}} \frac{\partial p_0(x)}{\partial x_i} s_i(x; w) dx.$$

Now, let's consider the integral over  $x_i$  first, and apply partial integration to it. Doing this relies on the Fubini-Tonelli's theorem, which can be safely applied due to  $B$  being finite as discussed above. We now apply partial integration to the integral over  $x_i$ :

$$\begin{aligned}
 & - \int_{a_1}^{b_1} \int_{a_2}^{b_2} \dots \int_{a_d}^{b_d} \frac{\partial p_0(x)}{\partial x_i} s_i(x; w) dx = - \int_{a_1}^{b_1} \dots \int_{a_{i-1}}^{b_{i-1}} \int_{a_{i+1}}^{b_{i+1}} \dots \int_{a_d}^{b_d} \left[ p_0(x) s_i(x; w) \Big|_{x_i=a_i}^{x_i=b_i} \right. \\
 & \quad \left. - \int_{a_i}^{b_i} p_0(x) \frac{\partial s_i(x; w)}{\partial x_i} dx_i \right] dx_1 \dots dx_{i-1} dx_{i+1} \dots dx_d \\
 & = \int_{a_1}^{b_1} \dots \int_{a_{i-1}}^{b_{i-1}} \int_{a_{i+1}}^{b_{i+1}} \dots \int_{a_d}^{b_d} \left[ \int_{a_i}^{b_i} p_0(x) \frac{\partial s_i(x; w)}{\partial x_i} dx_i \right] dx_1 \dots dx_{i-1} dx_{i+1} \dots dx_d \\
 & = \int_{\mathcal{X}} p_0(x) \frac{\partial s_i(x; w)}{\partial x_i} dx.
 \end{aligned}$$

where the second equality holds thanks to Assumption **A1**. For stating the last equality rigorously, we need to invoke Fubini-Tonelli's theorem again; this relies on the assumption:

$$\int_{\mathcal{X}} \left| p_0(x) \frac{\partial s_i(x; w)}{\partial x_i} \right| dx = \mathbb{E}_{p_0} \left| \frac{\partial s_i(X; w)}{\partial x_i} \right| < \infty,$$

which is equivalent to Assumption **A3**.

By repeating this argument for all terms in the sum in Eq. (13), we obtain that:

$$B = \int_{\mathcal{X}} p_0(x) \sum_{i=1}^d \frac{\partial s_i(x; w)}{\partial x_i} dx = \int_{\mathcal{X}} p_0(x) \sum_{i=1}^d \frac{\partial^2 \log p_w(x)}{\partial x_i^2} dx,$$

which concludes our proof. ■

## C.2 Proof of Theorem 7

We give here an extended version of Theorem 7, which we then prove following Hyvärinen (2005).

**Theorem 21** *Assume  $\exists w^* : p_0(\cdot) = p_{w^*}(\cdot)$ . Then:*

$$w = w^* \implies D_F(p_0 \| p_w) = 0.$$

*Further, if  $p_0(x) > 0 \forall x \in \mathcal{X}$ , you also have:*

$$D_F(p_0 \| p_w) = 0 \implies p_w(\cdot) = p_0(\cdot).$$

*Finally, if no other value  $w \neq w^*$  gives a pdf  $p_w$  that is equal to  $p_{w^*}$ :*

$$D_F(p_0 \| p_w) = 0 \implies w = w^*.$$

**Proof** The first statement is straightforward. For the second one, if  $D_F(p_0||p_w) = 0$  and  $p_0(x) > 0 \forall x \in \mathcal{X}$ , then  $\nabla_x \log p_w(x) = \nabla_x \log p_0(x)$ ; this in turn implies that  $\log p_w(x) = \log p_0(x) + c \forall x \in \mathcal{X}$  for a constant  $c$ , which however is 0 as both  $p_0$  and  $p_w$  are pdf's.

The third statement follows from the second as, for the additional assumption,  $w^*$  is the only choice of  $w$  which gives  $p_w = p_0$ .  $\blacksquare$

### C.3 Proof of Theorem 9

Similarly to the SM case, our Theorem 9 extends Theorem 1 in Song et al. (2020), that considers the case  $\mathcal{X} = \mathbb{R}^d$ , and which is recovered in the case  $a_i = -\infty$  and  $b_i = +\infty \forall i$ . Our proof follows Song et al. (2020), which however explicitly stated only Assumptions **A1**, **A2** and **A4**. Following Yu et al. (2019), we add the additional Assumption **A3** which is required for Fubini-Tonelli theorem to apply. The strategy of the proof is very similar to Theorem 6.

**Proof** As before, let  $s_0(x) = \nabla_x \log p_0(x)$  denote the score of the data distribution, and analogously  $s(x; w) = \nabla_x \log p_w(x)$ . Then, Eq. (6) can be rewritten as:

$$\begin{aligned} D_{FS}(p_0||p_w) &= \frac{1}{2} \int_{\mathcal{V}} \int_{\mathcal{X}} q(v) p_0(x) [v^T s_0(x) - v^T s(x; w)]^2 dx dv \\ &= \frac{1}{2} \int_{\mathcal{V}} \int_{\mathcal{X}} q(v) p_0(x) [(v^T s_0(x))^2 + (v^T s(x; w))^2 - 2(v^T s_0(x))(v^T s(x; w))] dx dv \\ &= \int_{\mathcal{V}} \int_{\mathcal{X}} q(v) p_0(x) \left[ \frac{1}{2} (v^T s(x; w))^2 - (v^T s_0(x))(v^T s(x; w)) \right] dx dv + C \end{aligned}$$

Note that we can split the integral into the three parts thanks to Assumptions **A2** and **A4**, which ensure that the expectation of each term is bounded. Additionally, we have absorbed in the constant  $C$  the term which does not depend on  $w$ .

In the last row of the above Equation, the first term in the square brackets appears in Eq. (7), which is what we want to prove. We focus therefore on the second term:

$$\begin{aligned} & - \int_{\mathcal{V}} \int_{\mathcal{X}} q(v) p_0(x) [(v^T s_0(x))(v^T s(x; w))] dx dv \\ &= - \int_{\mathcal{V}} \int_{\mathcal{X}} q(v) p_0(x) [(v^T \nabla_x \log p_0(x))(v^T \nabla_x \log p_w(x))] dx dv \\ &= - \int_{\mathcal{V}} \int_{\mathcal{X}} q(v) [(v^T \nabla_x p_0(x))(v^T \nabla_x \log p_w(x))] dx dv \tag{14} \\ &= - \sum_{i=1}^d \int_{\mathcal{V}} \int_{\mathcal{X}} q(v) \left[ v_i \frac{\partial p_0(x)}{\partial x_i} (v^T \nabla_x \log p_w(x)) \right] dx dv \end{aligned}$$

We will now consider one single element of the sum for a chosen  $i$ , and consider the integral over  $x_i$  first. Swapping the order of integration relies on Fubini-Tonelli's theorem, which can be safely applied due to the above quantity being finite as discussed above. We then

apply partial integration to the integral over  $x_i$ :

$$\begin{aligned}
 & - \int_{\mathcal{V}} \int_{a_1}^{b_1} \int_{a_2}^{b_2} \cdots \int_{a_d}^{b_d} q(v) \left[ v_i \frac{\partial p_0(x)}{\partial x_i} (v^T \nabla_x \log p_w(x)) \right] dx dv \\
 &= - \int_{\mathcal{V}} \int_{a_1}^{b_1} \cdots \int_{a_{i-1}}^{b_{i-1}} \int_{a_{i+1}}^{b_{i+1}} \cdots \int_{a_d}^{b_d} q(v) \left[ v_i p_0(x) (v^T \nabla_x \log p_w(x)) \Big|_{x_i \searrow a_i}^{x_i \nearrow b_i} \right. \\
 & \quad \left. - \int_{a_i}^{b_i} v_i p_0(x) \left( v^T \frac{\partial}{\partial x_i} \nabla_x \log p_w(x) \right) dx_i \right] dx_{-i},
 \end{aligned}$$

where  $dx_{-i} = dx_1 \dots dx_{i-1} dx_{i+1} \dots dx_d$ . The first element in the square bracket goes to 0 thanks to Assumption **A1**. We are left therefore with:

$$\begin{aligned}
 & \int_{\mathcal{V}} \int_{a_1}^{b_1} \cdots \int_{a_{i-1}}^{b_{i-1}} \int_{a_{i+1}}^{b_{i+1}} \cdots \int_{a_d}^{b_d} \left[ \int_{a_i}^{b_i} q(v) p_0(x) v_i \left( v^T \frac{\partial}{\partial x_i} \nabla_x \log p_w(x) \right) dx_i \right] dx_{-i} \\
 &= \int_{\mathcal{V}} \int_{\mathcal{X}} q(v) p_0(x) v_i \left( v^T \frac{\partial}{\partial x_i} \nabla_x \log p_w(x) \right) dx dv.
 \end{aligned}$$

The last equality again exploits Fubini-Tonelli theorem, which in this case requires:

$$\int_{\mathcal{V}} \int_{\mathcal{X}} q(v) p_0(x) \left| v_i \left( v^T \frac{\partial}{\partial x_i} \nabla_x \log p_w(x) \right) \right| dx dv = \mathbb{E}_{V \sim q} \mathbb{E}_{X \sim p_0} \left| V_i \left( V^T \frac{\partial}{\partial x_i} \nabla_x \log p_w(X) \right) \right| < \infty.$$

This is satisfied by combining Assumptions **A3** and **A4**, as in fact:

$$\begin{aligned}
 & \mathbb{E}_{V \sim q} \mathbb{E}_{X \sim p_0} \left| V_i \left( V^T \frac{\partial}{\partial x_i} \nabla_x \log p_w(X) \right) \right| = \mathbb{E}_{V \sim q} \mathbb{E}_{X \sim p_0} \left| V_i \sum_{j=1}^d V_j \frac{\partial^2}{\partial x_i \partial x_j} \log p_w(X) \right| \\
 & \leq \sum_{j=1}^d \mathbb{E}_{V \sim q} \mathbb{E}_{X \sim p_0} \left| V_i V_j \frac{\partial^2}{\partial x_i \partial x_j} \log p_w(X) \right| = \sum_{j=1}^d \mathbb{E}_{V \sim q} |V_i V_j| \cdot \mathbb{E}_{X \sim p_0} \left| \frac{\partial^2}{\partial x_i \partial x_j} \log p_w(X) \right| \\
 & \leq \sum_{j=1}^d \sqrt{\mathbb{E}_{V \sim q} V_i^2 \cdot \mathbb{E}_{V \sim q} V_j^2} \cdot \mathbb{E}_{X \sim p_0} \left| \frac{\partial^2}{\partial x_i \partial x_j} \log p_w(X) \right|,
 \end{aligned}$$

where the last inequality holds thanks to Cauchy-Schwarz inequality; Assumptions **A3** and **A4** ensure the last row is  $< \infty$ .

Back to Eq. (14), we have therefore:

$$\begin{aligned}
 & - \sum_{i=1}^d \int_{\mathcal{V}} \int_{\mathcal{X}} q(v) \left[ v_i \frac{\partial p_0(x)}{\partial x_i} (v^T \nabla_x \log p_w(x)) \right] dx dv \\
 &= \int_{\mathcal{V}} \int_{\mathcal{X}} \sum_{i=1}^d q(v) p_0(x) v_i \left( v^T \frac{\partial}{\partial x_i} \nabla_x \log p_w(x) \right) dx dv \\
 &= \int_{\mathcal{V}} \int_{\mathcal{X}} \sum_{i=1}^d \sum_{j=1}^d q(v) p_0(x) v_i v_j \frac{\partial^2}{\partial x_i \partial x_j} \log p_w(x) dx dv \\
 &= \int_{\mathcal{V}} \int_{\mathcal{X}} q(v) p_0(x) [v^T (\mathbf{H}_x \log p_w(x)) v] dx dv.
 \end{aligned}$$

■

### C.4 Proof of Theorem 10

We give here an extended version of Theorem 10. This is a version of Lemma 1 in Song et al. (2020), whose proof we adapt.

**Theorem 22** *Assume  $\exists w^* : p_0(\cdot) = p_{w^*}(\cdot)$ . Then:*

$$w = w^* \implies D_{FS}(p_0 \| p_w) = 0.$$

*Further, if  $p_0(x) > 0 \forall x \in \mathcal{X}$ , you also have:*

$$D_{FS}(p_0 \| p_w) = 0 \implies p_w(\cdot) = p_0(\cdot).$$

*Finally, if no other value  $w \neq w^*$  gives a pdf  $p_w$  that is equal to  $p_{w^*}$ :*

$$D_{FS}(p_0 \| p_w) = 0 \implies w = w^*.$$

**Proof** The first statement is straightforward.

For the second one, if  $D_F(p_0 \| p_w) = 0$  and  $p_0(x) > 0 \forall x \in \mathcal{X}$ , then:

$$\begin{aligned} & \int_{\mathcal{V}} q(v) (v^T (\nabla_x \log p_0(x) - \nabla_x \log p_w(x)))^2 dv = 0 \\ \iff & \int_{\mathcal{V}} q(v) v^T (\nabla_x \log p_0(x) - \nabla_x \log p_w(x)) (\nabla_x \log p_0(x) - \nabla_x \log p_w(x))^T v dv = 0 \\ \iff & (\nabla_x \log p_0(x) - \nabla_x \log p_w(x))^T E[VV^T] (\nabla_x \log p_0(x) - \nabla_x \log p_w(x)) = 0 \\ \stackrel{(\star)}{\iff} & \nabla_x \log p_0(x) - \nabla_x \log p_w(x) = 0 \\ \iff & \log p_w(x) = \log p_0(x) + c \quad \forall x \in \mathcal{X}, \end{aligned}$$

where in the third line above  $V$  is a random variable distributed according to  $q$ , for which therefore  $\mathbb{E}[VV^T]$  is positive definite by Assumption **A4**, which ensures equivalence  $(\star)$  holds. Additionally, as both  $p_0$  and  $p_w$  are pdf's (and therefore normalized), the constant  $c = 0$ .

The third statement follows from the second as, for the additional assumption,  $w^*$  is the only choice of  $w$  which gives  $p_w = p_0$ . ■

### C.5 Proof of Theorem 12

We give here an extended version of Theorem 12, which we then prove.

**Theorem 23** *Let  $Y = t(X) \in \mathcal{Y}$  for a bijection  $t$ , and denote by  $p_0^Y$  and  $p_w^Y$  the distributions on  $\mathcal{Y}$  induced by the distributions  $p_0$  and  $p_w$  on  $\mathcal{X}$ . Assume  $\exists w^* : p_0(\cdot) = p_{w^*}(\cdot)$ , and let  $D$  denote either  $D_F$  or  $D_{FS}$ . Then:*

$$w = w^* \implies D(p_0^Y \| p_w^Y) = 0.$$



Further, if  $p_0(x) > 0 \forall x \in \mathcal{X}$  and Assumption **A4** holds, you also have:

$$D(p_0^Y \| p_w^Y) = 0 \implies p_w(\cdot) = p_0(\cdot).$$

Finally, if no other value  $w \neq w^*$  gives a pdf  $p_w$  that is equal to  $p_{w^*}$ :

$$D(p_0^Y \| p_w^Y) = 0 \implies w = w^*.$$

**Proof** The proof relies on the equivalence between distributions for the random variables  $Y$  and  $X$ ; in fact, by fixing  $y = t(x)$  and denoting by  $|J_t(x)|$  the determinant of the Jacobian matrix of  $t$  evaluated in  $x$ , we have that  $p_0^Y(y) = \frac{p_0(x)}{|J_t(x)|}$  and  $p_w^Y(y) = \frac{p_w(x)}{|J_t(x)|}$ , so that  $p_0(\cdot) = p_w(\cdot) \iff p_0^Y(\cdot) = p_w^Y(\cdot)$ . The first and third statements follow directly from this fact by applying Theorem 21 (if  $D$  is chosen to be  $D_F$ ) or Theorem 22 (if  $D = D_{FS}$ ) to  $D(p_0^Y \| p_w^Y)$ ; for the second, notice also that  $p_0(x) > 0 \forall x \in \mathcal{X} \implies p_0^Y(y) \forall y \in \mathcal{Y} > 0$  as  $t$  is a bijection; then, Theorem 21 or Theorem 22 imply that  $p_0^Y(\cdot) = p_w^Y(\cdot) \implies p_0(\cdot) = p_w(\cdot)$ . ■

## C.6 Discussion on the positivity condition for SM

We follow here the discussion in Appendix D of Arbel and Gretton (2018).

Consider again the Fisher divergence in Eq. (4); we want to understand the conditions under which this is a divergence between probability measures, which essentially means it is 0 if and only if the probability distributions are the same.

Besides the fact that both  $p_0$  and  $p_w$  need to be continuous (otherwise the gradient would be a delta function), it turns out that a necessary condition is that  $p_0$  is positive on the whole space to which the random variable belong (let's say  $\mathcal{X}$ ), if we don't put any restrictions on  $p_w$ . Otherwise, the following scenario may happen (see Appendix D in Arbel and Gretton, 2018): consider the case in which  $p_0(x)$  is a mixture of two densities supported on disjoint sets:  $p_0(x) = \alpha_A p_A(x) + \alpha_B p_B(x)$ , with non-negative weights  $\alpha_A, \alpha_B$ , such that

$$p_A(x) > 0 \iff x \in \mathcal{X}_A, \quad p_B(x) > 0 \iff x \in \mathcal{X}_B, \quad \mathcal{X}_A, \mathcal{X}_B \in \mathcal{X}; \quad \mathcal{X}_A \cap \mathcal{X}_B = \emptyset.$$

Note that this implies  $\exists x \in \mathcal{X} : p_0(x) = 0$ . In this setting, any  $p_w$  of the form  $p_w(x) = \beta_A p_A(x) + \beta_B p_B(x)$  will give  $D_F(p_0 \| p_w) = 0$ . This can be seen by computing the Fisher divergence directly:

$$\begin{aligned}
 D_F(p_0||p_w) &= \frac{1}{2} \int_{\mathcal{X}} p_0(x) \|\nabla_x \log p_0(x) - \nabla_x \log p_w(x)\|^2 dx \\
 &= \frac{1}{2} \int_{\mathcal{X}} p_0(x) \|\nabla_x \log(\alpha_A p_A(x) + \alpha_B p_B(x)) - \nabla_x \log(\beta_A p_A(x) + \beta_B p_B(x))\|^2 dx \\
 &= \frac{1}{2} \int_{\mathcal{X}_A} p_0(x) \|\nabla_x \log(\alpha_A p_A(x)) - \nabla_x \log(\beta_A p_A(x))\|^2 dx + \\
 &\quad \frac{1}{2} \int_{\mathcal{X}_B} p_0(x) \|\nabla_x \log(\alpha_B p_B(x)) - \nabla_x \log(\beta_B p_B(x))\|^2 dx \\
 &= \frac{1}{2} \int_{\mathcal{X}_A} p_0(x) \underbrace{\|\nabla_x \log p_A(x) - \nabla_x \log p_A(x)\|^2}_{=0} dx + \\
 &\quad \frac{1}{2} \int_{\mathcal{X}_B} p_0(x) \underbrace{\|\nabla_x \log p_B(x) - \nabla_x \log p_B(x)\|^2}_{=0} dx \\
 &= 0,
 \end{aligned}$$

where the third equality relies on splitting the integration domain over the two subsets on which  $p_A$  and  $p_B$  are supported (and the other is 0) and the fourth equality relies on the presence of the logarithmic derivatives, that removes the importance of the mixture weights.

Due to the above, in the case of the conditional Fisher divergence in Eq. (9),  $p_0(x|\theta)$  needs to be supported on the whole  $\mathcal{X}$  for each  $\theta$  in order for  $D_F^E(p_0|p_w) = 0 \iff p_0(\cdot|\theta) = p_w(\cdot|\theta) \pi(\theta)$ -almost everywhere. This can be seen by considering the case of univariate parameter  $\theta$  and by choosing  $p_0(x|\theta) = p_A(x)H(\theta) + (1 - H(\theta))p_B(x)$ , where  $p_A$  and  $p_B$  are as above and  $H(\cdot)$  represents the Heaviside function. In this case, choosing  $p_w(x|\theta) = q(x)$ , where  $q(x) = \beta_A p_A(x) + \beta_B p_B(x)$  will give  $D_F^E(p_0|p_w) = 0$ , as for each fixed  $\theta$ ,  $D_F(p_0(\cdot|\theta)|p_w(\cdot|\theta))$  falls in the case considered above.

### C.7 Equivalence of Correction Factor and Transformed Score Matching

As discussed in the main text (Section 3.2), the first extension of score matching to non-negative random variables was given in Hyvärinen (2007):

$$D_F^+(p_0||p_w) = \frac{1}{2} \int_{\mathbb{R}_+^d} p_0(x) \|\nabla_x \log p_0(x) \odot x - \nabla_x \log p_w(x) \odot x\|^2 dx, \quad (15)$$

where  $\odot$  denotes element wise product between vectors and  $\mathbb{R}_+^d$  is the positive octant of  $\mathbb{R}^d$ . The correction factor  $x$  relaxes assumption **A1** to  $p_0(x) x_i^2 \frac{\partial}{\partial x_i} \log p_w(x) \rightarrow 0$ , so that it is possible to get an explicit form of the above with looser assumptions. This is an example of Corrected Score Matching (CorrSM, 3.2), in which the issue arising due to distribution having a compact support is fixed by introducing a correction factor in the formulation of the objective.

Yu et al. (2019) proposed a more general score matching for non-negative random variables by allowing freedom of choice in the factor that is used in the integrand to correct for the integration by parts step (Appendix C.1), leading to the following objective:

$$D_F^+(p_0||p_w) = \frac{1}{2} \int_{\mathbb{R}_+^d} p_0(x) \|(\nabla_x \log p_0(x)) \odot \sqrt{h(x)} - (\nabla_x \log p_w(x)) \odot \sqrt{h(x)}\|^2 dx, \quad (16)$$

where  $h(x)$  has the same dimension as  $x$ , and has positive elements almost surely.

The explicit formulation associated to Eq. (16) can be obtained under the following assumptions:

**A1b**  $p_0(x)h_j(x_j)\partial_j \log p_w(x) \rightarrow 0$  for  $x_i \searrow 0$  and  $x_i \nearrow \infty$ ,  $\forall w, i$ ,

**A2b**  $\mathbb{E}_{p_0} \|\nabla_x \log p_0(X) \odot h^{1/2}(X)\|_2^2 < \infty$ ,  $\mathbb{E}_{p_0} \|\nabla_x \log p_w(X) \odot h^{1/2}(X)\|_2^2 < \infty \forall w$ ,

**A3b**  $\mathbb{E}_{p_0} \|(\nabla_x \log p_w(X) \odot h(X))'\|_1 < \infty \forall w$ , where the prime symbol denotes element-wise differentiation.

Under the above assumption, Eq. (16) is equal to:

$$D_F^+(p_0 \| p_w) = \int_{\mathbb{R}_+^n} p_0(x) \sum_{i=1}^d \left[ \frac{1}{2} h_i(x) \left( \frac{\partial \log p_w(x)}{\partial x_i} \right)^2 + h_i(x) \left( \frac{\partial^2 \log p_w(x)}{\partial x_i^2} \right) + h_i'(x) \frac{\partial \log p_w(x)}{\partial x_i} \right] dx + C,$$

where  $C$  is a constant with respect to  $p_w$ .

In Proposition 2 in Yu et al. (2019), they give a result similar to our Theorems 7 and 12 guaranteeing that minimization of  $D_F^+(p_0 \| p_w)$  is a valid procedure for estimating a probabilistic model. When considering the finite-sample estimate of C.7, different choices of  $h(x)$  may allow to focus more on smaller/larger values of  $x$ , which may in practice have better properties than the original form for non-negative data in Hyvärinen (2007), which is recovered for  $h(x) = x^2$  (where the square is applied element-wise).

This formulation in Eq. (16), albeit originally considered for non-negative random variables only, can be extended to random variables with any bounded domain, by choosing a suitable function  $h$  and modifying **A1b** to hold for  $x_i$  going to the limits of the domain. In the next Sections, we therefore compare this approach with TranSM without specifying the domain; we will show that both the implicit and explicit formulation are the same with both TranSM and CorrSM, implying that the two are equivalent (we will show this in the specific case in which the transformation and the function  $h$  act independently on the different coordinates, but we believe this to be the case more in general; see Appendix C.9).

### C.7.1 EQUIVALENCE OF THE IMPLICIT FORM

As mentioned in the main text (Section 3.2), another approach to apply Fisher divergence to distributions with bounded domain (on one side or both) is to transform the data space to the real line and then apply standard Fisher divergence; we called this Transformation Score Matching (TranSM). Let us denote  $t$  such a transformation, which we assume to be bijective. Then, starting from  $p_0(x)$  and  $p_w(x)$  we get  $p_0(y) = \frac{p_0(x)}{|J_t(x)|}$  and  $p_w(y) = \frac{p_w(x)}{|J_t(x)|}$  for  $y = t(x)$ , where  $J_t$  is the Jacobian matrix of  $t$  and  $|\cdot|$  denotes here the determinant; here, differently from the main text, we use a lighter notation where  $p_0(y)$  and  $p_0(x)$  are two different densities associated to the different name of the argument (same for  $p_w$ ). We investigate what is the Fisher divergence between the densities of the transformed distributions. Recall that  $p_0(y)dy = p_0(x)dx$ , for  $y = t(x)$ . Moreover, we also have that:

$$\nabla_y g(y) = J_{t^{-1}}(t(x)) \nabla_x g(t(x)) = (J_t(x))^{-1} \nabla_x g(t(x)),$$

where the second equality comes from the fact that  $J_{t^{-1}}(t(x)) = (J_t(x))^{-1}$  due to the inverse function theorem. Then, we can compute the Fisher Divergence between  $p_0(y)$  and  $p_w(y)$  (corresponding to the TransSM objective):

$$\begin{aligned} D_F(p_0(y) \| p_w(y)) &= \frac{1}{2} \int p_0(y) \|\nabla_y \log p_0(y) - \nabla_y \log p_w(y)\|^2 dy \\ &= \frac{1}{2} \int p_0(x) \|(J_t(x))^{-1} [\nabla_x \log p_0(x) - \nabla_x \log |J_t(x)| - \nabla_x \log p_w(x) + \nabla_x \log |J_t(x)|]\|^2 dx \\ &= \frac{1}{2} \int p_0(x) \|(J_t(x))^{-1} [\nabla_x \log p_0(x) - \nabla_x \log p_w(x)]\|^2 dx. \end{aligned}$$

In the rather common case in which the transformation  $t$  acts on each component independently, the Jacobian matrix is diagonal; in this case, the latter is equivalent to Eq. (16) upon defining  $h(x)$  to be a vector containing the squares of the diagonal elements of the Jacobian, i.e. putting  $\sqrt{h_i(x)} = (J_t(x))_{ii}^{-1}$ .

### C.7.2 EQUIVALENCE OF THE EXPLICIT FORM

For both TransSM and CorrSM it is possible to get an explicit form of the objective (Eqs. 5 and C.7), in which the integrand does not depend on the data distribution  $p_0$ . In case in which the transformation  $t$  acts on the different components independently, the original explicit divergence for the transformed variable  $Y = t(X)$  is equivalent to the corrected explicit divergence for the original  $X$ , analogously to the implicit Fisher divergence form. In fact, by applying the definition of explicit Fisher divergence (Eq. 5) to the transformed  $Y$ , you get:

$$D_F(p_0(y) \| p_w(y)) = \int p_0(y) \underbrace{\sum_{i=1}^d \left[ \frac{1}{2} \left( \frac{\partial \log p_w(y)}{\partial y_i} \right)^2 + \left( \frac{\partial^2 \log p_w(y)}{\partial y_i^2} \right) \right]}_{\star} dy + C$$

where  $C$  is a constant with respect to  $p_w$ . Considering only the term in square brackets, denoting  $\partial_i = \frac{\partial}{\partial x_i}$  and setting  $\sqrt{h_i(x)} = (J_t(x))_{ii}^{-1}$ , we get:

$$\begin{aligned} \star &= \frac{1}{2} h_i(x) \left[ (\partial_i \log p_w(x))^2 + \frac{1}{4} (\partial_i \log h_i(x))^2 + \partial_i \log p_w(x) \cdot \partial_i \log h_i(x) \right] \\ &+ \frac{1}{2} h'_i(x) \cdot \partial_i \log p_w(x) + \frac{1}{4} h'_i(x) \partial_i \log h_i(x) + h_i(x) \partial_i^2 \log p_w(x) + \frac{1}{2} h_i(x) \partial_i^2 \log h_i(x) \\ &= \frac{1}{2} h_i(x) (\partial_i \log p_w(x))^2 + \frac{1}{8} h_i(x) (\partial_i \log h_i(x))^2 \\ &+ h'_i(x) \cdot \partial_i \log p_w(x) + \frac{1}{4} h'_i(x) \partial_i \log h_i(x) + h_i(x) \partial_i^2 \log p_w(x) + \frac{1}{2} h_i(x) \partial_i^2 \log h_i(x). \end{aligned}$$

The blue terms are the same that appear in the CorrSM explicit formulation (Eq. C.7), while all other terms are constants with respect to  $p_w$ .

We have shown therefore that CorrSM and TransSM are equivalent in both the explicit and implicit formulation if the transformation is applied independently on the elements of

$x$ . Therefore, the two approaches are completely equivalent when it comes to minimizing them.

### C.8 Specific formulation of TransSM

We discuss here the transformations we apply in this work in TransSM; specifically, we only consider the case in which the support for the multivariate  $x$  is defined by an intersection of element-wise inequalities, i.e.  $x \in \bigotimes_{i=1}^d (a_i, b_i)$ , where  $a_i, b_i$  can take on the values  $\pm\infty$  as well. In this case, then, a transformation can be applied independently on each element of  $x$ . We consider here the following transformations (which are also used in the Stan package Carpenter et al., 2017):

- When  $X \in [0, \infty)^d$ , the transformation we use is  $y_i = \log(x_i) \in \mathbb{R}^d$ . This corresponds to diagonal Jacobian with elements  $(J_t(x))_{ii}^{-1} = x_i$ , so that the above expression becomes the same as the original Fisher divergence for non-negative random variables discussed in Eq. (15).
- More generally, if  $x_i \in [a_i, +\infty)$  for  $|a_i| < \infty$ , we can transform the data as  $y_i = \log(x_i - a_i) \in \mathbb{R}$ , while if  $x_i \in (-\infty, b_i]$  for  $|b_i| < \infty$ , we simply reverse the transformation:  $y_i = \log(b_i - x_i) \in \mathbb{R}$ . These correspond to  $(J_t(x))_{ii}^{-1} = x_i - a_i$  and  $(J_t(x))_{ii}^{-1} = b_i - x_i$ .
- Finally, if  $x_i \in (a_i, b_i)$  for  $|a_i|, |b_i| < \infty$ , we can use the transformation defined as:  $y_i = t(x_i) = \text{logit}\left(\frac{x_i - a_i}{b_i - a_i}\right)$  with inverse transformation  $x_i = t^{-1}(y_i) = a + (b - a) \frac{e^{y_i}}{e^{y_i} + 1}$ . This corresponds to  $(J_t(x))_{ii}^{-1} = \frac{(x_i - a_i)(b_i - x_i)}{b_i - a_i}$ .

### C.9 Score matching for distributions with more general domain

As highlighted in the main text, across this work we are concerned with applying score matching to distributions whose support is defined by independent constraints on the different coordinates, as for instance  $\mathcal{X} = \bigotimes_{i=1}^d (a_i, b_i)$ . That is arguably the most common case in the literature. However, there have been some recent works which applied SM to more general cases. For instance, Mardia et al. (2016) devised a way to apply it to a directional distribution defined on an oriented Riemannian manifold (for instance, a sphere). It is interesting how their derivation of the explicit form from the implicit one relies on the classical divergence theorem (also known as Stokes' theorem), of which the partial integration trick used in Theorem 6 is a specific case. Liu and Kanamori (2019) introduced instead a way to apply score matching for a distribution on Euclidean space with complex truncation boundaries; their approach boils down to introducing a smart correction factor which goes to 0 at the boundary (thus allowing partial integration) but still being tractable; again, they need a more general version of Theorem 6 to obtain an objective for which the integrand does not depend on the data distribution.

In Appendix C.7, we established that CorrSM and TransSM are equivalent if the transformation is applied independently on the elements of  $x$ , which requires the domain to be defined by independent constraints on the coordinates. In the more general case of a irregular subset of Euclidean space (as in the setup of Liu and Kanamori, 2019), it is not

clear whether it is always possible to associate a correction factor to a transformation which maps the space to  $\mathbb{R}^d$ . That seems to be plausible if the domain satisfies some regularity conditions which may be related to convexity (for instance think of a triangle in  $\mathbb{R}^2$ , which can be easily stretched to cover the full space). We are not aware however of any work investigating this.

### C.10 Score matching with exponential family

We consider here the exponential family:

$$p_w(x|\theta) = \exp(\eta_w(\theta)^T f_w(x))/Z_w(\theta),$$

and want to find the value of  $w$  minimizing either  $D_F^E(p_0||p_w)$  or  $D_{FS}^E(p_0||p_w)$ , which are defined in Eq. (10).

Under the conditions discussed in Section 3.3, if  $\pi(\theta) > 0 \forall \theta$ , then  $D_F^E(p_0||p_w) = 0$  and  $D_{FS}^E(p_0||p_w) = 0 \iff p_w(x|\theta) = p_0(x|\theta) \pi(\theta)$ -almost everywhere. In this case, if  $f_w$  and  $\eta_w$  satisfy the conditions required for the theorems mentioned in Appendix B.1 to hold, then  $f_w$  and  $\eta_w$  are respectively sufficient statistics and natural parameters of  $p_0$ .

In order to find the value of the empirical estimate of the explicit for of both  $D_F^E(p_0||p_w)$  and  $D_{FS}^E(p_0||p_w)$ , we insert the definition of the exponential family with in Eq. 11, which leads to:

$$\hat{J}(w) = \frac{1}{N} \sum_{j=1}^N \left[ \sum_{i=1}^d \left( \frac{1}{2} \left( \eta_w(\theta^{(j)})^T \frac{\partial}{\partial x_i} f_w(x^{(j)}) \right)^2 + \eta_w(\theta^{(j)})^T \frac{\partial^2}{\partial x_i^2} f_w(x) \right) \right],$$

$$\hat{J}_S(w) = \frac{1}{NM} \sum_{j=1}^N \sum_{k=1}^M \left[ v^{(j,k),T} H_x(\eta_w(\theta^{(j)})^T f_w(x^{(j)})) v^{(j,k)} + \frac{1}{2} \sum_{i=1}^d \left( \eta_w(\theta^{(j)})^T \frac{\partial}{\partial x_i} f_w(x^{(j)}) \right)^2 \right].$$

Note that the objective does not change if you set  $f_w(x)$  to  $c + f_w(x)$ , for a constant vector  $c$ ; in fact, this constant gets absorbed into the normalizing constant in  $p_w(x|\theta)$ .

Similarly,  $\eta_w(\theta)^T f_w(x) = (1/c \cdot \eta_w(\theta))^T (c \cdot f_w(x))$  for some constant  $c \neq 0$ . Therefore, statistics and corresponding parameters are only defined up to a scale with respect to one another; if you use two Neural Networks to learn both of them, different network initializations may lead to different learned statistics and natural parameters, but their product should be fixed (up to translation of  $f_w(x)$ ).

However, this degeneracy may make training the approximate likelihood  $p_w$  with the score matching approach harder. In order to improve training, we usually add a Batch Normalization layer on top of the  $\eta_w$  network. Basically, Batch Normalization fixes the scale of the output of  $\eta_w$  over a training batch, therefore removing this additional degree of freedom and making training easier. We discuss in more detail this in Appendix E.2.

## Appendix D. Scoring Rules

A Scoring Rule (SR)  $S$  (Dawid and Musio, 2014; Gneiting and Raftery, 2007) is a function of a probability distribution over  $\mathcal{X}$  and of an observation in  $\mathcal{X}$ . In the framework of

probabilistic forecasting,  $S(P, y)$  represents the *penalty* which you incur when stating a forecast  $P$  for an observation  $y$ .<sup>10</sup>

If the observation  $y$  is a realization of a random variable  $Y$  with distribution  $Q$ , the expected Scoring Rule can be defined as:

$$S(P, Q) := \mathbb{E}_{Y \sim Q} S(P, Y),$$

where we overload notation in the second argument of  $S$ . The Scoring Rule  $S$  is said to be *proper* relative to a set of distributions  $\mathcal{P}(\mathcal{X})$  over  $\mathcal{X}$  if

$$S(Q, Q) \leq S(P, Q) \quad \forall P, Q \in \mathcal{P}(\mathcal{X}),$$

i.e., if the expected Scoring Rule is minimized in  $P$  when  $P = Q$ . Moreover,  $S$  is *strictly proper* relative to  $\mathcal{P}(\mathcal{X})$  if  $P = Q$  is the unique minimum:

$$S(Q, Q) < S(P, Q) \quad \forall P, Q \in \mathcal{P}(\mathcal{X}) \text{ s.t. } P \neq Q;$$

when minimizing an expected strictly proper Scoring Rule, a forecaster would provide their true belief (Gneiting and Raftery, 2007).

By following Dawid and Musio (2014), we define the divergence related to a proper Scoring Rule as  $D(P, Q) := S(P, Q) - S(Q, Q) \geq 0$ . Notice that  $P = Q \implies D(P, Q) = 0$ , but there may be  $P \neq Q$  such that  $D(P, Q) = 0$ . However, if  $S$  is strictly proper,  $D(P, Q) = 0 \iff P = Q$ , which is the commonly used condition to define a statistical divergence (as for instance the Kullback-Leibler, or KL, divergence). Therefore, each strictly proper Scoring Rule corresponds to a statistical divergence between probability distributions

In the following, we detail the two Scoring Rules which we consider in the main text (Section 5.3).

**Energy score** The energy score is given by:

$$S_E^{(\beta)}(P, y) = 2 \cdot \mathbb{E} \left[ \|X - y\|_2^\beta \right] - \mathbb{E} \left[ \|X - X'\|_2^\beta \right], \quad X \perp\!\!\!\perp X' \sim P,$$

where  $\beta \in (0, 2)$  and  $\perp\!\!\!\perp$  denotes independence between random variables. This is a strictly proper Scoring Rule for the class  $\mathcal{P}_\beta(\mathcal{X})$  of probability measures  $P$  such that  $\mathbb{E}_{X \sim P} \|X\|^\beta < \infty$  (Gneiting and Raftery, 2007). The related divergence is the square of the energy distance, which is a metric between probability distributions (Rizzo and Székely 2016)<sup>11</sup>:

$$D_E(P, Q) = 2 \cdot \mathbb{E} \left[ \|X - Y\|_2^\beta \right] - \mathbb{E} \left[ \|X - X'\|_2^\beta \right] - \mathbb{E} \left[ \|Y - Y'\|_2^\beta \right],$$

for  $X \perp\!\!\!\perp X' \sim P$  and  $Y \perp\!\!\!\perp Y' \sim Q$ .

10. Some authors (Gneiting and Raftery, 2007) use the convention of  $S(P, y)$  representing a *reward* rather than a *penalty*, which is equivalent up to change of sign.

11. The probabilistic forecasting literature (Gneiting and Raftery, 2007) use a different convention of the energy score and the subsequent kernel score, which amounts to multiplying our definitions by 1/2. We follow here the convention used in the statistical inference literature (Rizzo and Székely, 2016; Chérif-Abdellatif and Alquier, 2020; Nguyen et al., 2020)

In our case of interest (Section 5.3, Appendix G.4), we are unable to evaluate exactly  $S_E^{(\beta)}$  as we do not have a closed form for  $P$ . Thus, we obtain samples  $\{x_j\}_{j=1}^m$ ,  $x_j \sim P$ , and unbiasedly estimate the Energy Score with:

$$\hat{S}_E^{(\beta)}(\{x_j\}_{j=1}^m, y) = \frac{2}{m} \sum_{j=1}^m \|x_j - y\|_2^\beta - \frac{1}{m(m-1)} \sum_{\substack{j,k=1 \\ k \neq j}}^m \|x_j - x_k\|_2^\beta.$$

In the main text (Section 5.3), we use  $\beta = 1$ , in which case we simplify notation  $S_E^{(1)} = S_E$ .

**Kernel score** For a positive definite kernel  $k(\cdot, \cdot)$  (see Gretton et al. 2012), the kernel Scoring Rule for  $k$  is defined as (Gneiting and Raftery, 2007):

$$S_k(P, y) = \mathbb{E}[k(X, X')] - 2 \cdot \mathbb{E}[k(X, y)], \quad X \perp\!\!\!\perp X' \sim P.$$

The corresponding divergence is the squared Maximum Mean Discrepancy (MMD, Gretton et al., 2012) relative to the kernel  $k$ :

$$D_k(P, Q) = \mathbb{E}[k(X, X')] + \mathbb{E}[k(Y, Y')] - 2 \cdot \mathbb{E}[k(X, Y)],$$

for  $X \perp\!\!\!\perp X' \sim P$  and  $Y \perp\!\!\!\perp Y' \sim Q$ .

The Kernel Score is proper for the class of probability distributions for which  $\mathbb{E}[k(X, X')]$  is finite (by Theorem 4 in Gneiting and Raftery, 2007). Additionally, it is strictly proper under conditions which ensure that the MMD is a metric for probability distributions on  $\mathcal{X}$  (see for instance Gretton et al., 2012). These conditions are satisfied, among others, by the Gaussian kernel (which we use in this work):

$$k(x, y) = \exp\left(-\frac{\|x - y\|_2^2}{2\gamma^2}\right); \tag{17}$$

there,  $\gamma$  is a scalar bandwidth, which is tuned as described in Appendix G.4.1. As for the Energy Score, when the exact form of  $P$  is inaccessible and therefore  $S_k$  is impossible to be evaluated exactly, we use samples  $\{x_j\}_{j=1}^m$ ,  $x_j \sim P$ , and unbiasedly estimate the Kernel Score  $S_k(P, y)$  with:

$$\hat{S}_k(\{x_j\}_{j=1}^m, y) = \frac{1}{m(m-1)} \sum_{\substack{j,k=1 \\ k \neq j}}^m k(x_j, x_k) - \frac{2}{m} \sum_{j=1}^m k(x_j, y).$$

## Appendix E. Computational practicalities

### E.1 Computational cost of SM and SSM

For SM, as discussed in Section 3.2, exploiting automatic differentiation libraries to compute the second derivatives of the log density requires  $d$  times more backward derivative computations with respect to the first derivatives. In fact, automatic differentiation libraries are able to compute derivatives of a scalar with respect to several variables at once. One single call is therefore sufficient to obtain  $\nabla_x \log p_w(x)$ . However,  $d$  additional calls are required to



obtain the second derivatives  $\frac{\partial^2}{\partial x_i^2} \log p_w(x)$ ,  $i = 1, \dots, d$ , which are the diagonal elements of the Hessian matrix of  $\log p_w(x)$ ; each additional call computes the gradient of  $\frac{\partial}{\partial x_k} \log p_w(x)$  with respect to all components of  $x$ , for some  $k \in [1, 2, \dots, d]$ . Algorithm 2 in Song et al. (2020) gives a pseudocode implementation of this approach. Computational improvement can be obtained by implementing custom code which performs the gradient computation in the forward pass (i.e., along the computation of the neural network output for a given input  $x$ ). This avoids repeating some computations multiple times, which is done when performing the backward step repeatedly; however, the implementation is tricky and needs custom code for each different neural network type (we discuss how it can be done for a fully connected neural network in Appendix E.1.1). Additionally, the computational speed-up achievable in this way is limited with respect to what is offered by, for instance, SSM.

SSM instead requires only two backward propagation steps independently on the input size of the network. This is possible by exploiting the vector-Hessian product structure and computing the linear products with  $v$  (which is independent on the input  $x$ , thus can be swapped with gradient computation) after the gradient has been computed once, so that you only ever require the gradient of a scalar quantity. See Algorithm 1 in Song et al. (2020) for a precise description of how that can be done.

### E.1.1 FORWARD COMPUTATION OF DERIVATIVES

In the standard score matching approach, the first and second derivatives of Neural Network outputs with respect to the inputs are required, namely:

$$\frac{\partial f_w(x)}{\partial x_i} \quad \text{and} \quad \frac{\partial^2 f_w(x)}{\partial x_i^2}.$$

Neural network training is possible thanks to the use of autodifferentiation libraries, which allow to keep trace of the different operations and to automatically compute the gradients required for training. These libraries can be used to obtain the above derivatives.

However, as discussed previously, it is more efficient to compute the required derivatives during the forward pass of training (i.e. when the output of the Neural Network for a given input is computed). This requires additional coding effort specific to the chosen Neural Network architecture. For instance, Avrutskiy (2017) provide formulas to compute derivatives of any order recursively for fully connected Neural Networks. For more complex NN architectures, this approach is not viable as the coding effort becomes substantial. Additionally, with large  $d$  the improvement obtained by forward derivatives computation is much smaller than what offered, for instance, by SSM.

In the current work, the forward derivatives approach has been implemented for fully connected Neural Networks and Partially Exchangeable Networks (Appendix E.3). The computational advantage is evident for the computation of the second derivatives, as shown below. This approach allowed us to apply SM to relatively high dimensional data spaces (up to 100 dimensional for the MA(2) and AR(2) case), but not to the larger Lorenz96 model.

**Forward computation of derivatives of fully connected NNs** We revisit here the work in Avrutskiy (2017). Let us consider a fully connected Neural Networks with  $L$  layers, where the weights and biases of the  $l$ -th layer are denoted by  $W_l$  and  $b_l$ ,  $l = 1, \dots, L$ . Let

us denote by  $x$  the input of the Neural Network, and by  $z_l$  the hidden values after the  $l$ -th layer, before the activation function (denoted by  $\sigma$ ) is applied. Specifically, the Neural Network hidden values are determined by:

$$z_1 = W_1 \cdot x + b_1, \quad z_l = W_l \cdot \sigma(z_{l-1}) + b_l, \quad l = 2, \dots, L,$$

where the activation function is applied element wise. Similar recursive expressions can be given for the first derivatives:

$$\frac{\partial z_1}{\partial x_i} = (W_1)_{i,\cdot}, \quad \frac{\partial z_l}{\partial x_i} = W_l \cdot \left[ \sigma'(z_{l-1}) \odot \frac{\partial z_{l-1}}{\partial x_i} \right], \quad l = 2, \dots, L,$$

where  $(W_1)_{i,\cdot}$  denotes the  $i$ -th column of  $W_1$ , and  $\odot$  denotes element wise multiplication. Expressions for the second derivatives are instead:

$$\frac{\partial^2 z_1}{\partial x_i^2} = \mathbf{0}, \quad \frac{\partial^2 z_l}{\partial x_i^2} = W_l \cdot \left[ \sigma''(z_{l-1}) \odot \left( \frac{\partial z_{l-1}}{\partial x_i} \right)^2 + \sigma'(z_{l-1}) \odot \frac{\partial^2 z_{l-1}}{\partial x_i^2} \right], \quad l = 2, \dots, L,$$

where here  $\mathbf{0}$  denotes a 0 vector with the same size as  $z_1$ .

When instead we are interested in cross terms of the form  $\frac{\partial z_l}{\partial x_i \partial x_j}$ , we can apply the following:

$$\frac{\partial z_1}{\partial x_i \partial x_j} = \mathbf{0}, \quad \frac{\partial z_l}{\partial x_i \partial x_j} = W_l \cdot \left[ \sigma''(z_{l-1}) \odot \frac{\partial z_{l-1}}{\partial x_i} \odot \frac{\partial z_{l-1}}{\partial x_j} + \sigma'(z_{l-1}) \odot \frac{\partial^2 z_{l-1}}{\partial x_i \partial x_j} \right], \quad l = 2, \dots, L.$$

With respect to naively using auto-differentiation libraries, computing the derivatives in the forward step is much cheaper; specifically, for fully connected NNs, we found empirically the first to scale quadratically with the output size, while the second scales linearly (Figure 8).

## E.2 Batch normalization

The exponential family form used as  $p_w$  depends on  $\eta_w(\theta)^T f_w(x)$ ; if you multiply  $f_w(x)$  with an invertible matrix  $A$  and multiply  $\eta$  with  $(A^T)^{-1}$ , the product does not change. In order to remove this additional degeneracy, we use a Batch Normalization (BatchNorm) layer (Ioffe and Szegedy, 2015) to normalize the outputs of  $\eta_w(\theta)$ . Essentially, BatchNorm rescales the different features to have always the same range across different batches. More specifically, BatchNorm performs the following operation on  $y$ :

$$\tilde{y} = \frac{y - \mathbb{E}[Y]}{\sqrt{\mathbb{V}[Y] + \epsilon}} * \gamma + b,$$

where  $\epsilon$  is a small constant used for numerical stability, and  $\gamma$  and  $b$  are two (optionally learnable) sets of constants with dimension equal to  $y$  (set to 1 and 0 respectively by default). During training, the expectation  $\mathbb{E}$  and variance  $\mathbb{V}$  are estimated over the batch of training samples fed to the Neural Network. In testing mode, BatchNorm rescales the features by using a fixed estimate of  $\mathbb{E}$  and  $\mathbb{V}$ ; this estimate is obtained as a running mean over the batches; let  $s(Y)$  represent either the population expectation or variance. When the  $t$ -th

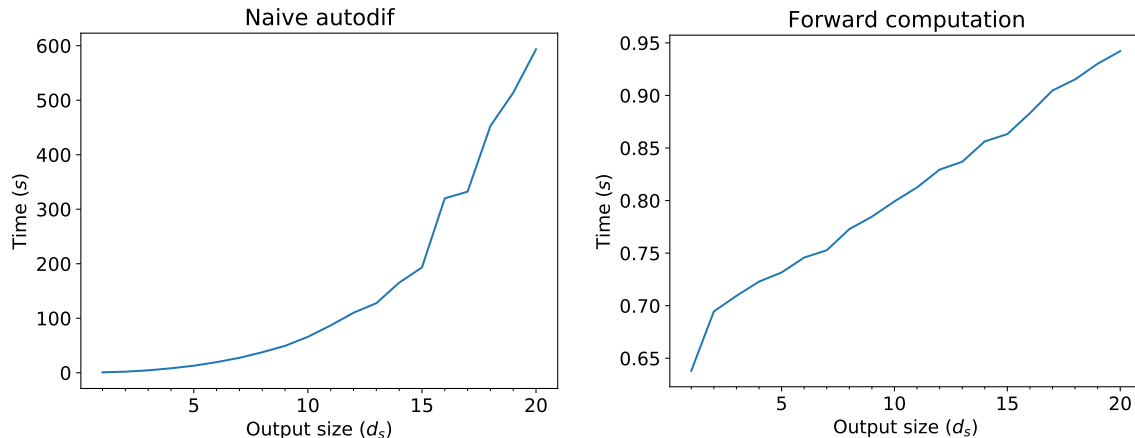


Figure 8: **Computational complexity for second order derivatives of Neural Network outputs** versus size of the output; we compare the forward computation of derivatives with naive autodifferentiation. Here, input size is fixed to 100, and one single batch of 5000 samples is fed to the network. Computations are done on a CPU machine with 8 cores.

batch is fed through the network in training mode, the running mean estimate is updated as follows:

$$\hat{s}_{new}(Y) = (1 - p) \cdot \hat{s}_{old}(Y) + p \cdot s_t(Y),$$

where  $s_t$  represent the estimate on the current batch,  $\hat{s}_{old}$  and  $\hat{s}_{new}$  respectively the old and updated running mean and  $p$  is a *momentum* constant which determines how quickly the running mean changes (the smaller it is, the slower the change of the estimate).

For instance, if the training of a Neural Network is quite unstable, the running estimate may not be a correct estimate of  $\mathbb{E}$  and  $\mathbb{V}$ , so that the test loss across training epochs may be very spiky, until network training stabilized. To solve this issue, you can either increase  $p$  (so that the running estimate “forgets” past information faster) or, alternatively, do a forward pass of the training data set (without computing gradients) before evaluating the test loss, so that the running estimate is more precisely estimated.

Across this work, we apply BatchNorm to  $y = \eta_w(\theta)$ . Moreover, we do not learn the translation parameters  $\gamma, b$ , and rather we fix them to be a vector of 1s and 0s.

### E.3 Partially Exchangeable Networks

Partially Exchangeable Networks (PENs) were introduced in Wiqvist et al. (2019) as a Neural Network architecture that satisfies the probabilistic invariance of Markovian models.

Specifically, let consider the case in which  $x = (x_1, \dots, x_d)$  comes from a Markovian model of order  $r$ , i.e.:

$$\begin{aligned} p(x|\theta) &= p(x_1|\theta)p(x_2|x_1; \theta)p(x_3|x_2, x_1; \theta) \prod_{i=4}^d p(x_i|x_1, \dots, x_{i-1}; \theta) \\ &= p(x_1|\theta) \prod_{i=2}^d p(x_i|x_{i-r}, \dots, x_{i-1}; \theta). \end{aligned}$$

This definition is an extension of the standard Markovianity assumption (of order 1), which corresponds to  $p(x|\theta) = p(x_1|\theta) \prod_{i=2}^d p(x_i|x_{i-1};\theta)$ , and it means that each element of  $x$  only depends on the last  $r$  elements. When  $r = 0$ , this corresponds to i.i.d. assumption.

When a model is  $r$ -Markovian, the probability density of an observation  $x$  is invariant to  $r$ -**block-switch transformation**, which is defined as follows:

**Definition 24  $r$ -block-switch transformation (Wiqvist et al., 2019)** Let  $x_{i:j}$  and  $x_{k:l}$  be two non-overlapping blocks with  $x_{i:(i+r)} = x_{k:(k+r)}$  and  $x_{(j-r):j} = x_{(l-r):l}$ . Then, denoting  $b = (i, j, k, l)$ , with  $j - i \geq r$  and  $l - k \geq r$ :

$$x = x_{1:i-1} \quad \boxed{x_{i:j}} \quad x_{(j+1):(k-1)} \quad \boxed{x_{k:l}} \quad x_{(l+1):M}$$

$$T_b^{(r)}(x) = x_{1:i-1} \quad \boxed{x_{k:l}} \quad x_{(j+1):(k-1)} \quad \boxed{x_{i:j}} \quad x_{(l+1):M}.$$

Otherwise, if  $x_{i:(i+r)} \neq x_{k:(k+r)}$  or  $x_{(j-r):j} \neq x_{(l-r):l}$ , then  $T_b^{(r)}(x) = x$ .

For instance, let us consider the case in which the data space is  $\mathcal{X} = \{0, 1 \dots 9\}^{16}$ , and  $r = 2$ . An example of the above transformation is the following:

$$x = 1 \quad \boxed{7 \ 2 \ 3 \ 6 \ 4 \ 5 \ 8} \quad 1 \ 7 \quad \boxed{7 \ 2 \ 9 \ 5 \ 8} \quad 1$$

$$T_{(2,8,11,15)}^{(2)}(x) = 1 \quad \boxed{7 \ 2 \ 9 \ 5 \ 8} \quad 1 \ 7 \quad \boxed{7 \ 2 \ 3 \ 6 \ 4 \ 5 \ 8} \quad 1.$$

The authors of Wiqvist et al. (2019) provide a simple Neural Network structure which is invariant to the  $r$ -block-switch transformation, motivated by the following theorem:

**Theorem 25 ((Wiqvist et al., 2019))** Let  $f : \mathcal{X}^M \rightarrow A$   $r$ -block-switch invariant. If  $\mathcal{X}$  is countable,  $\exists \phi : \mathcal{X}^{r+1} \rightarrow \mathbb{R}$  and  $\rho : \mathcal{X}^r \times \mathbb{R} \rightarrow A$  such that:

$$\forall x \in \mathcal{X}^M, f(x) = \rho \left( x_{1:r}, \sum_{i=1}^{M-r} \phi(x_{i:(i+r)}) \right). \quad (18)$$

In practice,  $\phi$  and  $\rho$  are two independent Neural Networks (which we take to be fully connected in our case), giving rise to a PEN of order  $r$ .

In Wiqvist et al. (2019), the authors show that the posterior mean of a Markovian variable of order  $r$  needs to be invariant to the  $r$ -block-switch transformation. Therefore, this motivates using a PEN for learning a summary statistics as in the approach by Fearnhead and Prangle (2012). Here, we use PEN for parametrizing the statistics in the approximating exponential family as well; this choice implies that the approximating family has the same Markovianity property as the true distribution, as we discuss in the following.

### E.3.1 RESULTS FOR THE EXPONENTIAL FAMILY

A deeper connection between  $r$ -Markovian probability models and  $r$ -block-switch transformation exists. We can in fact state the following result:

**Lemma 26** A probability model  $p(x|\theta)$  is  $r$ -Markovian  $\iff$  the function  $x \mapsto p(x|\theta)$  is  $r$ -block-switch invariant, i.e.  $p(x|\theta) = p(T_b^{(r)}(x)|\theta) \forall T_b^{(r)}$ .

**Proof** The forward direction is straightforward and can be seen by considering the decomposition of an  $r$ -Markovian model.

The converse direction can be shown by contradiction; assume in fact that  $x \mapsto p(x|\theta)$  is  $r$ -block-switch invariant but not Markovian. As it is not Markovian,  $\exists x = (x_i, x_2, \dots, x_n)$  for which  $x_{i:(i+r)} = x_{k:(k+r)}$  and  $x_{(j-r):j} = x_{(l-r):l}$  such that, defining  $b = (i, j, k, l)$ ,  $p(x|\theta) \neq p(T_b^{(r)}(x)|\theta)$ . This is however in contradiction with  $r$ -block-switch invariance, which leads to our result.  $\blacksquare$

In the case where the model we consider has a sufficient statistic, we can write  $p(x|\theta) = h(x)g(t(x)|\theta)$ . We get therefore the following corollary, which can be seen by applying the above result:

**Corollary 27** *Consider a distribution  $p(x|\theta) = h(x)g(t(x)|\theta)$ ; if the function  $h(x)$  and  $t(x)$  are  $r$ -block-switch invariant, then  $p(x|\theta)$  is  $r$ -Markovian.*

Without any further assumptions, this result is not enough to say that the sufficient statistics  $t(x)$  for a Markovian model is  $r$ -block-switch invariant; in fact, the choice  $t(x) = x$  always constitutes a sufficient statistic; moreover, in the decomposition  $p(x|\theta) = h(x)g(t(x)|\theta)$ , it may be that the function  $t(x)$  is not  $r$ -block-switch invariant but  $g(t(x)|\theta)$  is, or otherwise that the product  $h(x)g(t(x)|\theta)$  is  $r$ -block-switch invariant even if the individual terms are not. We can however get the following result:

**Lemma 28** *Consider a  $r$ -Markovian distribution  $p(x|\theta) = h(x)g(t(x)|\theta)$ ; if  $x \mapsto t(x)$  is not an injection mapping, then  $x \mapsto h(x)$  is  $r$ -block-switch invariant.*

**Proof** If  $x \mapsto t(x)$  is not an injection,  $\exists x, x'$  such that  $t(x) = t(x')$ . If the density is not degenerate, moreover,  $\exists \theta : p(x|\theta), p(x'|\theta) > 0$ . Therefore, we consider the following ratio:

$$\frac{p(x|\theta)}{p(x'|\theta)} = \frac{h(x)}{h(x')} \cdot \frac{g(t(x)|\theta)}{g(t(x')|\theta)} = \frac{h(x)}{h(x')}.$$

Now, the left hand side is  $r$ -block-switch invariant with respect to both  $x$  and  $x'$  independently, implying that  $h(x)$  is as well.  $\blacksquare$

We remark that it does not seem possible in general to say anything about  $t(x)$ , as in fact it may be that the function  $t(x)$  is not  $r$ -block-switch invariant but  $g(t(x)|\theta)$  is. In the specific case of an exponential family, however, a more specific result can be obtained:

**Lemma 29** *Consider an exponential family distribution  $p(x|\theta) = h(x) \exp(\eta(\theta)^T f(x))/Z(\theta)$  which is  $r$ -Markovian distribution; if  $x \mapsto f(x)$  is not an injection mapping, then  $x \mapsto f(x)$  is  $r$ -block-switch invariant.*

**Proof** Without loss of generality, we consider the case in which all elements of  $\eta(\theta)$  are not constant with respect to  $\theta$ ; if this is not the case, in fact, you can redefine the exponential family by incorporating the elements of  $f(x)$  corresponding to the constant ones of  $\eta$  in the  $h(x)$  factor.

Now,  $h(x)$  is  $r$ -block-switch invariant thanks to Lemma 28. Consider now the following decomposition:

$$\log p(x|\theta) = \log h(x) - \log Z(\theta) + \sum_i f_i(x)\eta_i(\theta);$$

as that needs to be  $r$ -block-switch invariant for any  $\theta$ , this can happen only if each of the  $f_i(x)$  elements are  $r$ -block-switch invariant.  $\blacksquare$

Overall, these results imply that an exponential family in which  $f$  is parametrized with a PEN of order  $r$  is  $r$ -Markovian. Moreover, provided that  $f$  is not an injection mapping, all  $r$ -Markovian exponential families have  $f$  which satisfy the  $r$ -block-switch invariant property, which is imposed by using a PEN network of order  $r$ .

### E.3.2 FORWARD COMPUTATION OF DERIVATIVES FOR PENS

We give here the derivation for the forward computation of derivatives with PENS. If we pick here  $\phi$  and  $\rho$  to be fully connected Neural Networks, we can moreover apply the forward computation of derivatives for them and we are able to compute the derivatives for PENS at a much lower cost with respect to using automatic differentiation libraries.

In Eq. (18), let us denote for brevity  $z = \sum_{i=1}^{M-r} \phi(x_{i:(i+r)})$ . We are interested now in computing the derivative:

$$\frac{\partial f}{\partial x_j} = \frac{\partial}{\partial x_j} \rho(x_{1:r}, z),$$

where note that  $z$  depends in general on  $x_i$ . Therefore, in computing the above, we need to compute the derivative with respect to both arguments; let us denote by  $\frac{\partial'}{\partial x_j}$  the derivative with respect to the first argument. Then, we have:

$$\frac{\partial f}{\partial x_j} = \frac{\partial'}{\partial x_j} \rho(x_{1:r}, z) \cdot \mathbb{1}[j \leq r] + \frac{\partial}{\partial z} \rho(x_{1:r}, z) \cdot \frac{\partial z}{\partial x_j},$$

where the second term is (note that  $z$  and  $\rho$  are multivariate, so that  $\frac{\partial'}{\partial x_j} \rho(x_{1:r}, z)$  is a Jacobian matrix):

$$\frac{\partial z}{\partial x_j} = \frac{\partial}{\partial x_j} \sum_{i=1}^{M-r} \phi(x_{i:(i+r)}) = \sum_{\substack{i=j-r \\ i \geq 1}}^j \frac{\partial}{\partial x_j} \phi(x_{i:(i+r)}),$$

where all other terms of the sum disappear as they do not contain  $x_j$ .

Now, we are interested in obtaining the second derivative terms:

$$\begin{aligned} \frac{\partial^2 f}{\partial x_j^2} &= \frac{\partial}{\partial x_j} \frac{\partial f}{\partial x_j} = \frac{\partial}{\partial x_j} \left[ \frac{\partial'}{\partial x_j} \rho(x_{1:r}, z) \cdot \mathbb{1}[j \leq r] + \sum_k \frac{\partial}{\partial z_k} \rho(x_{1:r}, z) \frac{\partial z_k}{\partial x_j} \right] \\ &= \frac{\partial'^2}{\partial x_j^2} \rho(x_{1:r}, z) \cdot \mathbb{1}[j \leq r] + 2 \sum_k \frac{\partial'}{\partial x_j} \frac{\partial}{\partial z_k} \rho(x_{1:r}, z) \cdot \frac{\partial z_k}{\partial x_j} \cdot \mathbb{1}[j \leq r] \\ &\quad + \sum_{k,k'} \frac{\partial^2}{\partial z_k \partial z_{k'}} \rho(x_{1:r}, z) \frac{\partial z_k}{\partial x_j} \frac{\partial z_{k'}}{\partial x_j} + \sum_k \frac{\partial}{\partial z_k} \rho(x_{1:r}, z) \frac{\partial^2 z_k}{\partial x_j^2}; \end{aligned}$$

in the above expression,  $\sum_k$  runs over the elements of  $z$  and  $\frac{\partial^2}{\partial x_j^2}$  denotes second derivative with respect to the first element. Note that all terms appearing in the above formulas contain first and second derivatives of the Neural Networks  $\phi$  and  $\rho$  with respect to one single input, except for the terms highlighted in red. In order to compute that, obtaining the full hessian matrix of  $\rho$  is required. We remark that the latter can be very large in case the input dimension is large, therefore leading to memory overflow issues.

#### E.4 Exchange MCMC

For convenience, we describe here the ExchangeMCMC algorithm by Murray et al. (2012). We consider the task of sampling from a posterior distribution  $\pi(\theta|x)$ . We can evaluate an unnormalized version of the likelihood  $\tilde{p}(x|\theta)$ , and we denote the normalized version as  $p(x|\theta) = \tilde{p}(x|\theta)/Z(\theta)$ ,  $Z(\theta)$  being an intractable normalizing constant. We want to build an MCMC chain by using a proposal distribution  $q(\cdot|\theta; x)$  (which optionally depends on the considered  $x$  as well). Usually, the standard Metropolis acceptance threshold for a proposal  $\theta'$  is defined as:

$$\alpha = \frac{\pi(\theta'|x)q(\theta|\theta'; x)}{\pi(\theta|x)q(\theta'|\theta; x)} = \frac{\tilde{p}(x|\theta')q(\theta|\theta'; x)\pi(\theta')}{\tilde{p}(x|\theta)q(\theta'|\theta; x)\pi(\theta)} \cdot \frac{Z(\theta)}{Z(\theta')},$$

where the last factor cannot be evaluated, as we do not have access to the normalizing constant.

The ExchangeMCMC algorithm proposed by Murray et al. (2012) bypasses this issue by drawing an auxiliary observation  $x' \sim p(\cdot|\theta')$  and defining the acceptance probability as:

$$\alpha = \frac{p(x|\theta')q(\theta|\theta'; x)\pi(\theta')}{p(x|\theta)q(\theta'|\theta; x)\pi(\theta)} \cdot \frac{p(x'|\theta)}{p(x'|\theta')} = \frac{\tilde{p}(x|\theta')\tilde{p}(x'|\theta)q(\theta|\theta'; x)\pi(\theta')}{\tilde{p}(x|\theta)\tilde{p}(x'|\theta')q(\theta'|\theta; x)\pi(\theta)} \cdot \frac{Z(\theta)}{Z(\theta')} \cdot \frac{Z(\theta')}{Z(\theta)}. \quad (19)$$

Here, all the normalizing constants cancel out, so that the acceptance threshold can be evaluated explicitly, at the expense of drawing a simulation from the likelihood for each MCMC step. Murray et al. (2012) showed that an MCMC chain using the above acceptance rate targets the correct posterior  $\pi(\theta|x)$ . The resulting algorithm is given in Algorithm 1:

---

**Algorithm 1** Original exchangeMCMC algorithm (Murray et al., 2012).

---

**Require:** Initial  $\theta$ , number of iterations  $T$ , proposal distribution  $q$ , observation  $x$ .

- 1: **for**  $i = 1$  **to**  $T$  **do**
  - 2:     Propose  $\theta' \sim q(\theta'|\theta; x)$
  - 3:     Generate auxiliary observation  $x' \sim p(\cdot|\theta')$
  - 4:     Compute acceptance threshold  $\alpha$  as in Eq. (19)
  - 5:     With probability  $\alpha$ , set  $\theta \leftarrow \theta'$
  - 6: **end for**
- 

**Bridging.** When considering more closely the acceptance rate in Eq. (19), it can be seen that it depends on two ratios:  $\frac{p(x|\theta')}{p(x|\theta)}$  represents how well the proposed parameter value explains the observation with respect to the previous parameter value, while instead  $\frac{p(x'|\theta)}{p(x'|\theta')}$  measures how well the auxiliary variable (generated using  $\theta'$ ) can be explained with

parameter  $\theta$ . Therefore, even if the former is large and  $\theta$  would be a suitable parameter value,  $\alpha$  can still be small if the auxiliary random variable is not explained well by the previous parameter value. This can lead to slow mixing of the chain; to improve on this, Murray et al. (2012) proposed to sample a set of auxiliary variables  $(x'_0, x'_1, \dots, x'_K)$  from intermediate distributions in the following way<sup>12</sup>: consider a set of densities

$$\tilde{p}_k(x|\theta, \theta') = \tilde{p}(x|\theta')^{\beta_k} \tilde{p}(x|\theta)^{1-\beta_k}, \quad \beta_k = \frac{K - k + 1}{K + 1};$$

$x'_0$  is generated from  $p(\cdot|\theta')$  as before, and then each  $x'_k$  is generated from  $R(\cdot|x'_{k-1}; \theta, \theta')$ , which denotes a Metropolis-Hastings transition kernel starting from  $x'_{k-1}$  with stationary density  $\tilde{p}_k(\cdot|\theta, \theta')$ . Then, the acceptance rate is modified as follows:

$$\alpha = \frac{\tilde{p}(x|\theta')q(\theta|\theta'; x)\pi(\theta')}{\tilde{p}(x|\theta)q(\theta|\theta; x)\pi(\theta)} \cdot \prod_{k=0}^K \frac{\tilde{p}_{k+1}(x'_k|\theta, \theta')}{\tilde{p}_k(x'_k|\theta, \theta')}. \quad (20)$$

The overall algorithm is given in Algorithm 2. Note that  $K = 0$  recovers the original ExchangeMCMC. This procedure generally improves the acceptance rate as it basically introduces a sequence of intermediate updates to the auxiliary data which by smoothening out the difference between the two distributions.

---

**Algorithm 2** ExchangeMCMC algorithm with bridging (Murray et al., 2012).

---

**Require:** Initial  $\theta$ , number of iterations  $T$ , proposal distribution  $q$ , number of bridging steps  $K$ , observation  $x$ .

- 1: **for**  $i = 1$  **to**  $T$  **do**
  - 2:     Propose  $\theta' \sim q(\theta'|\theta; x)$
  - 3:     Generate auxiliary observation  $x'_0 \sim p(\cdot|\theta')$
  - 4:     **for**  $k = 1$  **to**  $K$  **do** ▷ Bridging steps
  - 5:         Generate  $x'_k \sim R(\cdot|x'_{k-1}; \theta, \theta')$
  - 6:     **end for**
  - 7:     Compute acceptance threshold  $\alpha$  as in Eq. (20)
  - 8:     With probability  $\alpha$ , set  $\theta \leftarrow \theta'$
  - 9: **end for**
- 

**ExchangeMCMC without perfect simulations.** If, as in the setup considered across this work, we are not able to sample from  $p(\cdot|\theta')$  as it is required in the ExchangeMCMC algorithm (line 3 in Alg. 1), Murray et al. (2012) suggested to run  $T_{in}$  steps of an MCMC chain on  $x$  targeting  $p(\cdot|\theta')$  at each step of ExchangeMCMC; if  $T_{in}$  is large enough, the last sample can be considered as (approximately) drawn from  $p(\cdot|\theta')$  itself and used in place of the unavailable perfect simulation. In practice, however, this only leads to an approximate ExchangeMCMC algorithm, as at each iteration of the inner chain a finite  $T_{in}$  is used, so that the inner chain would not perfectly converge to its target; for this reason, even an infinitely long outer chain would not target the right posterior for any finite

---

12. Differently from the rest of the work, here subscripts do not denote vector components, but rather different auxiliary variables.



$T_{in}$ . Nonetheless, this approach was shown empirically to work satisfactorily in Caimo and Friel (2011); Everitt (2012); Liang (2010). Some theoretical guarantees (albeit under strong conditions), are given in Appendix B by Everitt (2012), which bounds the total variation distance between target of approximate ExchangeMCMC with finite  $T_{in}$  and the target of the exact one, and shows that they become equal when  $T_{in} \rightarrow \infty$ .

In Liang (2010), they argue that starting the inner chain from the observation value improves convergence; we adapt this approach in our implementation (Algorithm 3).

---

**Algorithm 3** ExchangeMCMC algorithm (Murray et al., 2012) with inner MCMC.

---

**Require:** Initial  $\theta$ , number of iterations  $T$  and  $T_{in}$ , proposal distributions  $q$  and  $q_x$ , observation  $x$ .

```

1: for  $i = 1$  to  $T$  do                                     ▷ Outer chain
2:   Propose  $\theta' \sim q(\theta'|\theta; x)$ 
3:   Set  $x' = x$                                              ▷ Start inner chain from the observation
4:   for  $j = 1$  to  $T_{in}$  do                                   ▷ Inner chain
5:     Propose  $x'' \sim q_x(x''|x')$ 
6:     With probability  $\frac{p_w(x''|\theta')q_x(x'|x'')}{p_w(x'|\theta')q_x(x''|x')}$ , set  $x' = x''$ 
7:   end for
8:   Compute acceptance threshold  $\alpha$  as in Eq. (19)     ▷ This uses last point of inner
MCMC
9:   With probability  $\alpha$ , set  $\theta \leftarrow \theta'$ 
10: end for

```

---

Note that it is still possible to run bridging steps after the inner MCMC to sample from  $p(\cdot|\theta')$ ; the algorithm combining bridging and inner MCMC, which is used across this work, is given in Algorithm 4).

**Related algorithms.** Algorithms for sampling from doubly-intractable targets which are suitable for parallel computing exist, for instance Caimo and Friel (2011) propose a parallel-chain MCMC algorithm, while Everitt et al. (2017) build instead an SMC-type algorithm which is also capable of recycling information from past simulations. However, in this work we stick to using ExchangeMCMC, which turned out to be relatively cheap to use and easy to implement.

Finally, we remark that Liang et al. (2016) proposed an algorithm which is inspired from ExchangeMCMC and, in the case of impossible perfect sampling, still targets the right invariant distribution. It works by considering a set of parallel chains targeting  $p(\cdot|\theta^{(i)})$  for a fixed set of  $\{\theta^{(i)}\}$  and iteratively updating those and the main chain over  $\theta$ . The algorithm relies on some assumptions which are probably satisfied in practice (as discussed in Park and Haran, 2018). It also requires some hand-tuning and needs to keep in memory a large amount of data, which may hinder its applicability. For this reason, we do not investigate using that here.

#### E.4.1 IMPLEMENTATION DETAILS

**Acceptance rate tuning.** In our implementation of ExchangeMCMC, we discarded some burn-in steps to make sure the chain forgets its initial state. During burn-in, moreover, after

---

**Algorithm 4** ExchangeMCMC algorithm (Murray et al., 2012) with inner MCMC and bridging.

---

**Require:** Initial  $\theta$ , number of iterations  $T$  and  $T_{in}$ , number of bridging steps  $K$ , proposal distributions  $q$  and  $q_x$ , observation  $x$ .

```

1: for  $i = 1$  to  $T$  do ▷ Outer chain
2:   Propose  $\theta' \sim q(\theta'|\theta; x)$ 
3:   Set  $x' = x$  ▷ Start inner chain from the observation
4:   for  $j = 1$  to  $T_{in}$  do ▷ Inner chain
5:     Propose  $x'' \sim q_x(x''|x')$ 
6:     With probability  $\frac{p_w(x''|\theta')q_x(x'|x'')}{p_w(x'|\theta')q_x(x''|x')}$ , set  $x' = x''$ 
7:   end for
8:   Set  $x'_0 = x'$ 
9:   for  $k = 1$  to  $K$  do ▷ Bridging steps
10:    Generate  $x'_k \sim R(\cdot|x'_{k-1}; \theta, \theta')$ 
11:  end for
12:  Compute acceptance threshold  $\alpha$  as in Eq. (20)
13:  With probability  $\alpha$ , set  $\theta \leftarrow \theta'$ 
14: end for

```

---

each window of 100 outer steps, we tuned the proposal sizes considering the acceptance rate in the window, and increasing (respectively decreasing) if that is too large (small) with respect to a chosen interval (see below). Notice that we applied this strategy independently to the proposal size for the outer chain, the inner chain and the bridging step, when the latter is used. When tuning the proposal size for the inner chain or bridging steps, we considered the overall acceptance rate over each window of 100 outer steps. We remark how this was only done during burn-in, so that the convergence properties of the chain were not affected.

For inner and outer steps as well as bridging, we found that a target acceptance rate in the interval  $[0.2, 0.5]$  lead to good performance; this is consistent with the recommended range for Metropolis-Hastings MCMC (Roberts et al., 1997).

**MCMC on bounded space.** When the inner (or outer) MCMC chain is run on a bounded domain, we apply the transformations discussed in Appendix C.8 to map it to an unbounded domain, and therefore run the MCMC on that space. Notice that the Jacobian factor arising from the transformations has therefore to be taken into account when computing the acceptance rate.

## Appendix F. Details on Neural Networks training

For all experiments, we used the `Pytorch` library (Paszke et al., 2019) to train NNs. In Tables 3, 4, 5, 6 and 7 we report the Neural Network architectures used in the different experiments.  $FC(n, m)$  denotes a fully connected layer with  $n$  inputs and  $m$  outputs. For the time-series and Lorenz96 experiments,  $\phi_w$  and  $\rho_w$  represent the two Neural Networks used to build the PEN network  $f_w$  as described in Eq. (18), and similarly  $\phi_\beta$  and  $\rho_\beta$  represent the ones used in building  $s_\beta$ . Finally,  $BN(p)$  represents a BatchNorm layer with momentum

$p$  (as described in Appendix E.2), with  $\gamma, b$ , fixed respectively to be vectors of 1s and 0s. We remark that the momentum value does not impact on the training of the network, but it modifies the evaluation of the test loss, which we use for early stopping, as discussed below.

<i>Network</i>	$f_w$	$\eta_w$	$s_\beta$
<i>Structure</i>	FC(10,30)	FC(2,15)	FC(10,30)
	FC(30,50)	FC(15,30)	FC(30,50)
	FC(50,50)	FC(30,30)	FC(50,50)
	FC(50,20)	FC(30,15)	FC(50,20)
	FC(20,3)	FC(15,2)	FC(20,2)
		BN(0.9)	

Table 3: **Architectures used for the exponential family models** (Gaussian, Gamma and Beta).

<i>Network</i>	$f_w$		$\eta_w$	$s_\beta$	
	$\phi_w$	$\rho_w$		$\phi_\beta$	$\rho_\beta$
<i>Structure</i>	FC(3,50)	FC(22,50)	FC(2,15)	FC(3,50)	FC(22,50)
	FC(50,50)	FC(50,50)	FC(15,30)	FC(50,50)	FC(50,50)
	FC(50,50)	FC(50,3)	FC(30,30)	FC(50,30)	FC(50,2)
	FC(30,20)		FC(30,15)	FC(30,20)	
			FC(15,2)		
		BN(0.9)			

Table 4: **Architectures used for the AR(2) model.**

<i>Network</i>	$f_w$		$\eta_w$	$s_\beta$	
	$\phi_w$	$\rho_w$		$\phi_\beta$	$\rho_\beta$
<i>Structure</i>	FC(11,50)	FC(30,50)	FC(2,15)	FC(11,50)	FC(30,50)
	FC(50,50)	FC(50,50)	FC(15,30)	FC(50,50)	FC(50,50)
	FC(50,30)	FC(50,3)	FC(30,30)	FC(50,30)	FC(50,2)
	FC(30,20)		FC(30,15)	FC(30,20)	
			FC(15,2)		
		BN(0.9)			

Table 5: **Architectures used for the MA(2) model.**

In all experiments, stochastic gradient descent with a batch size of 1000 samples is used with Adam optimizer (Kingma and Ba, 2015), whose parameters are left to the default values as implemented in `Pytorch`. Finally, we evaluate the test loss on a test set with same size as the training set at intervals of  $T_{\text{check}}$  epochs and early-stop training if the test loss increased with respect to the last evaluation. In order to have a better running estimate

<i>Network</i>	$f_w$		$\eta_w$	$s_\beta$	
	$\phi_w$	$\rho_w$		$\phi_\beta$	$\rho_\beta$
<i>Structure</i>	FC(16,50)	FC(28,40)	FC(4,30)	FC(16,50)	FC(28,40)
	FC(50,100)	FC(40,90)	FC(30,50)	FC(50,100)	FC(40,90)
	FC(100,50)	FC(90,35)	FC(50,50)	FC(100,50)	FC(90,35)
	FC(50,20)	FC(35,5)	FC(50,30)	FC(50,20)	FC(35,4)
			FC(30,4)		
			BN(0.9)		

Table 6: **Architectures used for the Lorenz96 model in the small setup.**

<i>Network</i>	$f_w$		$\eta_w$	$s_\beta$	
	$\phi_w$	$\rho_w$		$\phi_\beta$	$\rho_\beta$
<i>Structure</i>	FC(80,120)	FC(60,80)	FC(4,30)	FC(80,120)	FC(60,80)
	FC(120,160)	FC(80,100)	FC(30,50)	FC(120,160)	FC(80,100)
	FC(160,120)	FC(100,80)	FC(50,50)	FC(160,120)	FC(100,80)
	FC(120,20)	FC(80,5)	FC(50,30)	FC(120,20)	FC(80,4)
			FC(30,4)		
			BN(0.9)		

Table 7: **Architectures used for the Lorenz96 model in the large setup.**

of the quantities of interest for the BatchNorm layer, before each test epoch we perform a forward pass of the whole training data set, without computing gradients, as discussed in Appendix E.2. We do not perform early stopping before epoch  $T_{\text{start}}$ . The nets are trained for a maximum of  $T$  training epochs with  $N_{\text{train}}$  training samples. In some experiments, we used an exponential learning rate scheduler, which decreases progressively the learning rate by multiplying it by a factor  $\zeta < 1$  at each epoch. We fixed  $\zeta = 0.99$ . The values of the parameters are reported in Table 8 for all the different models and setups, together with the values of the learning rates (lr) used and whether the scheduler was used or not (Sch).

We remark that, in order to make the comparison fair, the learning rates for the FP experiments were chosen by cross validation with several learning rates choices. Similarly, we hand-picked the best learning rate values for training the Neural Networks with SM and SSM, even if in that case it was not possible, due to computational constraints, to perform a full search on a large set of values of learning rates for all experiments. In fact, the computational cost of training the exponential family approximation with SM or SSM is larger than the cost of learning the summary statistics with the FP approach. Moreover, in this scenario we have two independent learning rates values to tune, as we train two networks simultaneously. For this reason, we did not spend too much time trying different lr values for models for which we already got satisfactory results.

<i>Model</i>	<i>Setup</i>	<i>Learning rates</i>	$N_{train}$	$T$	$T_{start}$	$T_{check}$	<i>Sch</i>
Gaussian	SM	$\text{lr}(f_w) = 0.0003, \text{lr}(\eta_w) = 0.003$	$10^4$	500	150	10	Yes
	SSM	$\text{lr}(f_w) = 0.001, \text{lr}(\eta_w) = 0.001$	$10^4$	500	200	10	Yes
	FP	$\text{lr}(s_\beta) = 0.01$	$10^4$	1000	300	25	No
Gamma	SM	$\text{lr}(f_w) = 0.001, \text{lr}(\eta_w) = 0.001$	$10^4$	500	200	10	Yes
	SSM	$\text{lr}(f_w) = 0.001, \text{lr}(\eta_w) = 0.001$	$10^4$	500	200	10	Yes
	FP	$\text{lr}(s_\beta) = 0.001$	$10^4$	1000	300	25	Yes
Beta	SM	$\text{lr}(f_w) = 0.001, \text{lr}(\eta_w) = 0.001$	$10^4$	500	200	10	Yes
	SSM	$\text{lr}(f_w) = 0.001, \text{lr}(\eta_w) = 0.001$	$10^4$	500	200	10	Yes
	FP	$\text{lr}(s_\beta) = 0.01$	$10^4$	1000	250	50	Yes
AR(2)	SM	$\text{lr}(f_w) = 0.001, \text{lr}(\eta_w) = 0.001$	$10^4$	500	100	25	Yes
	SSM	$\text{lr}(f_w) = 0.001, \text{lr}(\eta_w) = 0.001$	$10^4$	500	100	25	Yes
	FP	$\text{lr}(s_\beta) = 0.001$	$10^4$	1000	500	25	Yes
MA(2)	SM	$\text{lr}(f_w) = 0.001, \text{lr}(\eta_w) = 0.001$	$10^4$	500	100	25	Yes
	SSM	$\text{lr}(f_w) = 0.001, \text{lr}(\eta_w) = 0.001$	$10^4$	500	100	25	Yes
	FP	$\text{lr}(s_\beta) = 0.001$	$10^4$	1000	500	25	Yes
Lorenz96 Small	SSM	$\text{lr}(f_w) = 0.001, \text{lr}(\eta_w) = 0.001$	$10^4$	1000	500	50	Yes
	FP	$\text{lr}(s_\beta) = 0.001$	$10^4$	1000	200	25	Yes
Lorenz96 Large	SSM	$\text{lr}(f_w) = 0.001, \text{lr}(\eta_w) = 0.001$	$10^4$	1000	500	50	Yes
	FP	$\text{lr}(s_\beta) = 0.001$	$10^4$	1000	200	25	Yes

Table 8: **Hyperparameter values for NN training:** “FP” denotes the least squares regression by Fearnhead and Prangle (2012); Jiang et al. (2017) used for the ABC-FP experiment, while “SM” and “SSM” denote respectively exponential family trained with Score Matching and Sliced Score Matching.

## Appendix G. Additional experimental results

### G.1 Exponential family models

#### G.1.1 MEAN CORRELATION COEFFICIENTS FOR NEURAL NETWORKS TRAINED WITH SSM

We report in Table 9 the weak and strong Mean Correlation Coefficient (MCC, Appendix B.1.1) for Neural Networks trained with SSM, for the exponential family models. MCC is a metric in  $[0, 1]$ , with 1 denoting perfect recovery up to a linear transformation (weak) or permutation (strong). As it can be seen in Table 9, our method leads to values quite close to 1, particularly for the weak MCC, implying that our method is able to recover the embeddings up to a linear transformation, as expected.

<i>Model</i>	<i>MCC weak in</i>	<i>MCC weak out</i>	<i>MCC strong in</i>	<i>MCC strong out</i>
Beta (statistics)	0.990	0.986	0.982	0.979
Beta (nat. par.)	0.987	0.989	0.983	0.985
Gamma (statistics)	0.939	0.928	0.723	0.709
Gamma (nat. par.)	0.977	0.977	0.792	0.794
Gaussian (statistics)	0.874	0.844	0.623	0.638
Gaussian (nat. par.)	0.861	0.862	0.581	0.543

Table 9: **MCC for exponential family models between exact embeddings and those learned with SSM.** We show weak and strong MCC values; MCC is between 0 and 1 and measures how well an embedding is recovered up to permutation and rescaling of its components (strong) or linear transformation (weak); the larger, the better. “in” denotes MCC on training data used to find the best transformation, while “out” denote MCC on test data. We used 500 samples in both training and test data sets.!

#### G.1.2 LEARNED AND EXACT EMBEDDINGS FOR THE EXPONENTIAL FAMILY MODELS

We compare here the exact and learned sufficient statistics and natural parameters of the exponential family models; precisely, we draw samples  $(x^{(j)}, \theta^{(j)})$  and then plot the learned statistics  $f_w(x^{(j)})$  versus the exact one, and similarly for the natural parameters. Figure 9 reports the results for Neural Networks trained with SM, while Figure 10 reports the results for Neural Networks trained with SSM.

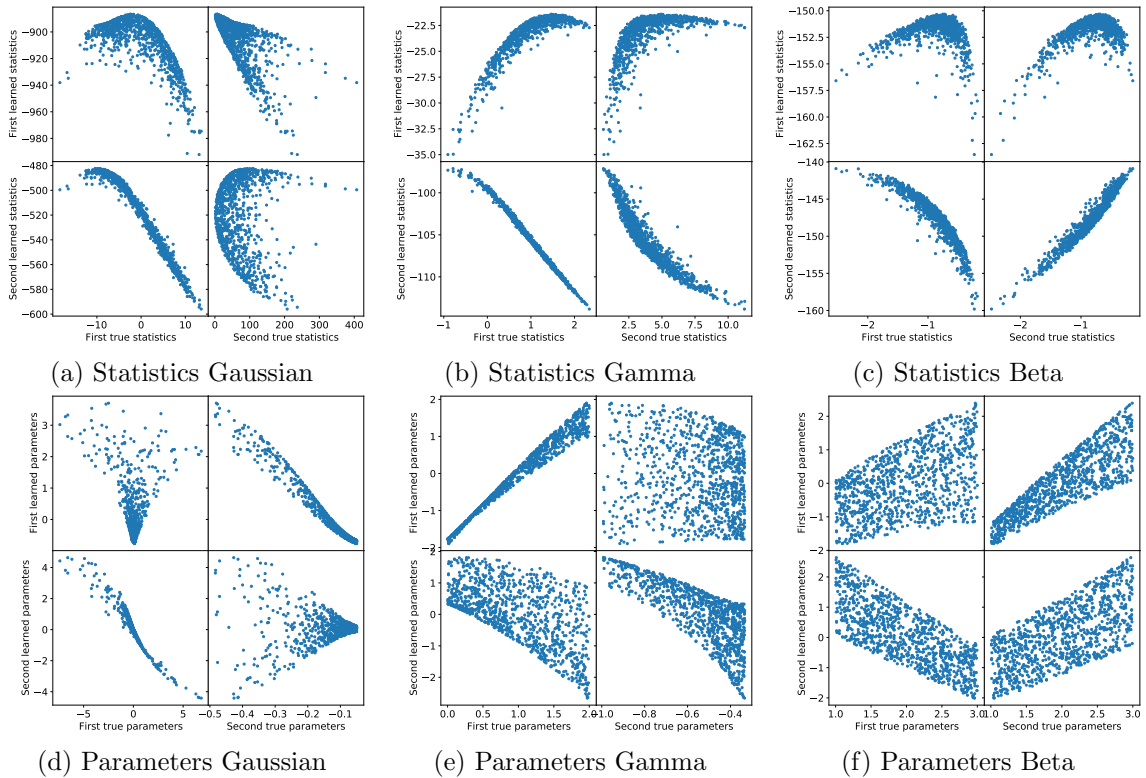


Figure 9: **Learned and exact embeddings for exponential family models obtained with SM.** Each point represents a different  $x$  (for the statistics) or  $\theta$  (for the natural parameters); 1000 of each were used here.

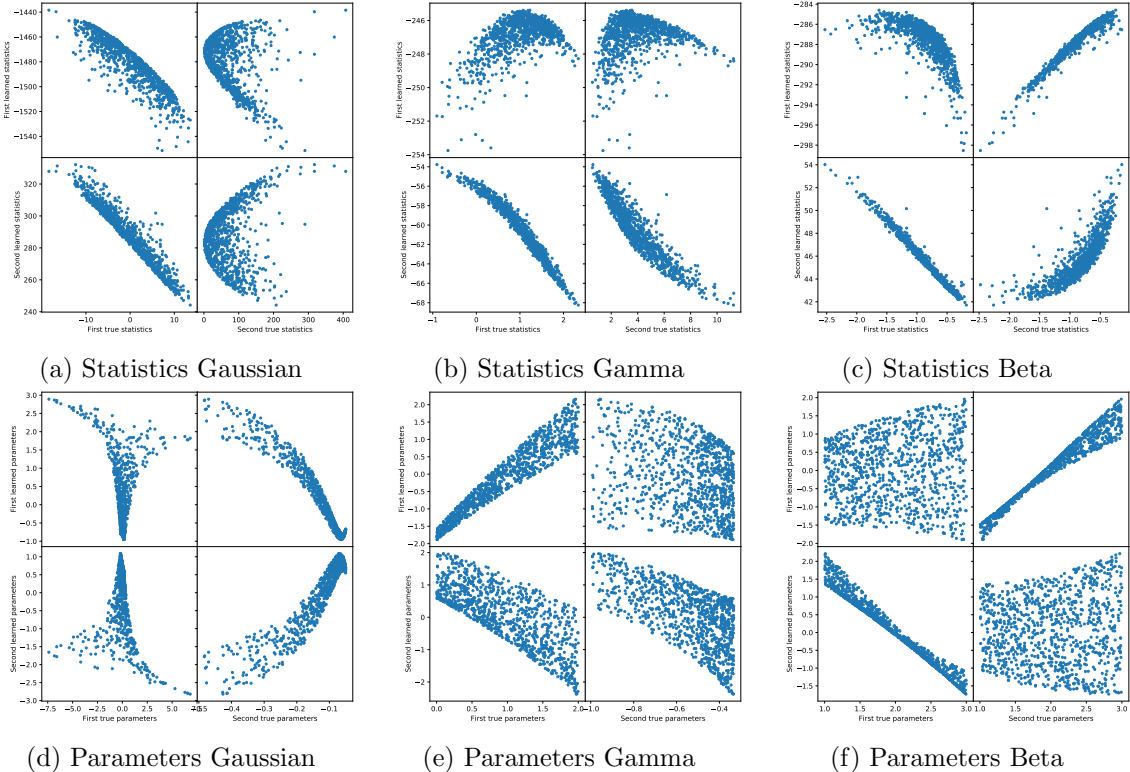


Figure 10: **Learned and exact embeddings for exponential family models obtained with SSM.** Each point represents a different  $x$  (for the statistics) or  $\theta$  (for the natural parameters); 1000 of each were used here.



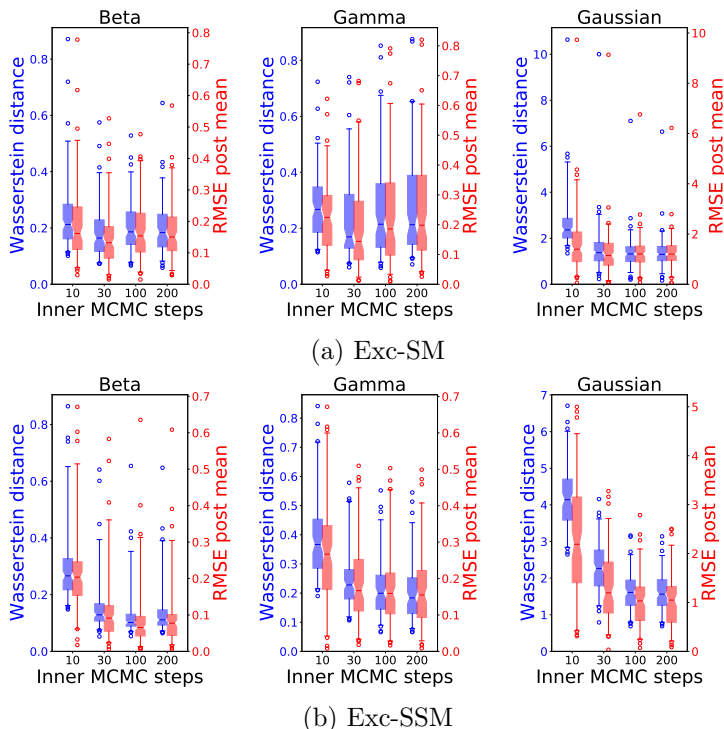


Figure 11: **Performance of Exc-SM with different number of inner MCMC steps, for exponential family models.** Wasserstein distance from the exact posterior and RMSE between exact and approximate posterior means are reported for 100 observations using boxplots. Boxes span from 1st to 3rd quartile, whiskers span 95% probability density region and horizontal line denotes median. The numerical values are not comparable across examples, as they depend on the range of parameters.

### G.1.3 PERFORMANCE OF EXC-SM AND EXC-SSM

We study here the performance of Exc-SM and Exc-SSM with different numbers of inner steps in the ExchangeMCMC algorithm (Algorithm 4) for the Exponential family models. Specifically, we run the inference with 10, 30, 100 and 200 inner MCMC steps, and we evaluate the performance in these 4 cases (Figure 11); considering the different models, we observe that the performance with 30 steps is almost equivalent to the one with 100 and 200, albeit being faster (see the computational time in Table 10). In the main text, we therefore present results using 30 inner MCMC steps.

<i>Inner MCMC steps</i>	10	30	100	200
<i>Time (minutes)</i>	$\approx 2$	$\approx 4$	$\approx 16$	$\approx 28$

Table 10: **Approximate computational time of Exc-SM with different number of inner MCMC steps for the exponential family models.** These values were obtained by running on a single core.

<i>Inner MCMC steps</i>	10	30	100	200
<i>Time (minutes)</i>	$\approx 2$	$\approx 5$	$\approx 17$	$\approx 29$

Table 11: **Approximate computational time of Exc-SM with different number of inner MCMC steps, for the AR(2) and MA(2) models.** These values were obtained by running on a single core.

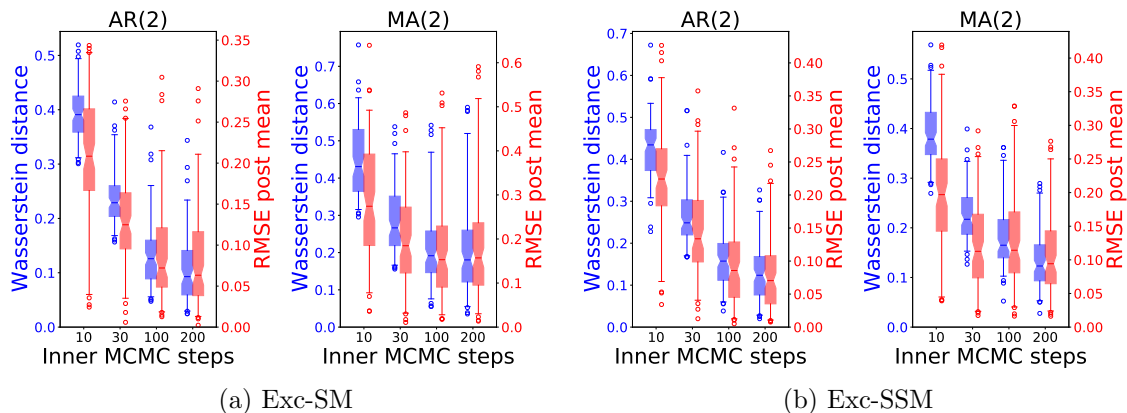


Figure 12: **Performance of Exc-SM with different number of inner MCMC steps, for AR(2) and MA(2) models.** Wasserstein distance from the exact posterior and RMSE between exact and approximate posterior means are reported for 100 observations using boxplots. Boxes span from 1st to 3rd quartile, whiskers span 95% probability density region and horizontal line denotes median.

## G.2 AR(2) and MA(2) models

We study here the performance of Exc-SM and Exc-SSM with different numbers of inner steps in the ExchangeMCMC algorithm (Algorithm 4) for the AR(2) and MA(2) models. Specifically, we run the inference with 10, 30, 100 and 200 inner MCMC steps, and we evaluate the performance in these 4 cases (Figure 12); considering the different models, we observe that the performance with 30 steps is almost equivalent to the one with 100 and 200, albeit being faster (see the computational time in Table 11). In the main text, we therefore present results using 30 inner MCMC steps.

## G.3 Simulation cost to reach equivalent performance as Exc-SM

Above, we have showed that Exc-SM and Exc-SSM are competitive with the other approaches, even if they require no additional model simulations. Here, we quantify how many model simulations are needed for the other techniques to reach the same performance of Exc-SM in the exponential family and time-series models. The same analysis could be done with respect to Exc-SSM but, as its performance is generally close to the one achieved by Exc-SM, we avoid repeating it.

We compute therefore the performance of ABC-SM, ABC-SSM, ABC-FP, PMC-SL, PMC-RE at each iteration for all 100 observations and find when their median performance

(as quantifies by the Wasserstein distance with respect to the true posterior) becomes comparable or better than the median one achieved by Exc-SM. The number of required simulations are reported in Table 12. Notice that some techniques are not able to outperform Exc-SM for some models; we highlight that by a dash in the Table. SL and RE reach similar performance to Exc-SM with one single iteration, when they are able to do so; in fact, we found empirically that the performance of them does not significantly improve with iterations. Still, we remark that one single iteration of SL and RE requires a very large number of model simulations.

In Table 12, we also give the number of simulations required for the preliminary NN training in the methods which use one; further, we compute the overall cost of inference in terms of model simulations (taking into account both NN training and inference) for different number of observations; we remark that, as discussed previously, our method requires no additional model simulations to perform inference after the NNs have been trained. From Table 12, it can be seen that, for all models, ABC-SM, ABC-FP, PMC-RE and PMC-SL require a number of simulations larger than the one needed to train the NNs in Exc-SM to reach the performance achieved by Exc-SM, which makes the latter an interesting option for models in which simulations are very expensive.

In Figures 13 and 14 we represent the performance attained by ABC-FP, ABC-SM, ABC-SSM, PMC-SL and PMC-RE at each iteration of the iterative algorithm. On the horizontal axis of all plots, we report the number of model simulations corresponding to the iteration of the algorithm.

#### G.4 Validation with Scoring Rules for Lorenz96 model

Recall that the Lorenz96 model (Section 5.3) is a multivariate time-series model. Therefore, we use the Scoring Rules (on which more details are given in Appendix D) to evaluate the predictive performance at each timestep separately.

First, let us recall the definition of the posterior predictive density:

$$p(x|x^0) = \int p(x|\theta)\pi(\theta|x^0),$$

whose corresponding posterior predictive distribution we denote as  $P(\cdot|x^0)$ . In the above expression,  $\pi(\theta|x^0)$  may represent the posterior obtained with any of the considered methods (Exc-SM, Exc-SSM, ABC-FP, ABC-SM, ABC-SSM).

Let  $P^{(t)}(\cdot|x^0)$  denote the posterior predictive distribution at time  $t$  conditioned on an observation  $x^0$ , and let  $x^{0,(t)}$  denote the  $t$ -th timestep of the observation. Then, we are interested in:

$$S_E(P^{(t)}(\cdot|x^0), x^{0,(t)}) \quad \text{and} \quad S_k(P^{(t)}(\cdot|x^0), x^{0,(t)});$$

we estimate these using samples from  $P^{(t)}(\cdot|x^0)$  with the unbiased estimators discussed in Appendix D. In the Kernel Score, the bandwidth of the Gaussian kernel is set from simulations as described in Appendix G.4.1.

In this way, we obtain a score for an observation  $x^0$  at each timestep of the model. This procedure is repeated for 100 different observations; the scores at each timestep are reported in Figure 15, while the summed scores over timesteps were reported in Figure 7 in the main text. In both Figures, we report the median value and various quantiles over the considered

Beta						
	Exc-SM	ABC-SM	ABC-SSM	ABC-FP	PMC-SL	PMC-RE
NN training	$2 \cdot 10^4$	$2 \cdot 10^4$	$2 \cdot 10^4$	$2 \cdot 10^4$	-	-
Inference	0	$2.5 \cdot 10^4$	$1.8 \cdot 10^4$	-	$1 \cdot 10^5$	$1 \cdot 10^6$
Total 1 obs	$2 \cdot 10^4$	$4.5 \cdot 10^4$	$3.8 \cdot 10^4$	-	$1 \cdot 10^5$	$1 \cdot 10^6$
Total 3 obs	$2 \cdot 10^4$	$9.5 \cdot 10^4$	$7.4 \cdot 10^4$	-	$3 \cdot 10^5$	$3 \cdot 10^6$
Total 100 obs	$2 \cdot 10^4$	$2.52 \cdot 10^6$	$1.82 \cdot 10^6$	-	$1 \cdot 10^7$	$1 \cdot 10^8$
Gamma						
	Exc-SM	ABC-SM	ABC-SSM	ABC-FP	PMC-SL	PMC-RE
NN training	$2 \cdot 10^4$	$2 \cdot 10^4$	$2 \cdot 10^4$	$2 \cdot 10^4$	-	-
Inference	0	$4.7 \cdot 10^4$	$2.9 \cdot 10^4$	-	$1 \cdot 10^5$	$1 \cdot 10^6$
Total 1 obs	$2 \cdot 10^4$	$6.7 \cdot 10^4$	$4.9 \cdot 10^4$	-	$1 \cdot 10^5$	$1 \cdot 10^6$
Total 3 obs	$2 \cdot 10^4$	$1.61 \cdot 10^5$	$1.07 \cdot 10^5$	-	$3 \cdot 10^5$	$3 \cdot 10^6$
Total 100 obs	$2 \cdot 10^4$	$4.72 \cdot 10^6$	$2.92 \cdot 10^6$	-	$1 \cdot 10^7$	$1 \cdot 10^8$
Gaussian						
	Exc-SM	ABC-SM	ABC-SSM	ABC-FP	PMC-SL	PMC-RE
NN training	$2 \cdot 10^4$	$2 \cdot 10^4$	$2 \cdot 10^4$	$2 \cdot 10^4$	-	-
Inference	0	$2.7 \cdot 10^4$	$2.6 \cdot 10^4$	-	-	$1 \cdot 10^6$
Total 1 obs	$2 \cdot 10^4$	$4.7 \cdot 10^4$	$4.6 \cdot 10^4$	-	-	$1 \cdot 10^6$
Total 3 obs	$2 \cdot 10^4$	$1.01 \cdot 10^5$	$9.8 \cdot 10^4$	-	-	$3 \cdot 10^6$
Total 100 obs	$2 \cdot 10^4$	$2.72 \cdot 10^6$	$2.62 \cdot 10^6$	-	-	$1 \cdot 10^8$
AR2						
	Exc-SM	ABC-SM	ABC-SSM	ABC-FP	PMC-SL	PMC-RE
NN training	$2 \cdot 10^4$	$2 \cdot 10^4$	$2 \cdot 10^4$	$2 \cdot 10^4$	-	-
Inference	0	$3.2 \cdot 10^4$	$3.1 \cdot 10^4$	$2.3 \cdot 10^4$	$1 \cdot 10^5$	-
Total 1 obs	$2 \cdot 10^4$	$5.2 \cdot 10^4$	$5.1 \cdot 10^4$	$4.3 \cdot 10^4$	$1 \cdot 10^5$	-
Total 3 obs	$2 \cdot 10^4$	$1.16 \cdot 10^5$	$1.13 \cdot 10^5$	$8.9 \cdot 10^4$	$3 \cdot 10^5$	-
Total 100 obs	$2 \cdot 10^4$	$3.22 \cdot 10^6$	$3.12 \cdot 10^6$	$2.32 \cdot 10^6$	$1 \cdot 10^7$	-
MA2						
	Exc-SM	ABC-SM	ABC-SSM	ABC-FP	PMC-SL	PMC-RE
NN training	$2 \cdot 10^4$	$2 \cdot 10^4$	$2 \cdot 10^4$	$2 \cdot 10^4$	-	-
Inference	0	$3.1 \cdot 10^4$	$2.1 \cdot 10^4$	$2.0 \cdot 10^4$	$1 \cdot 10^5$	-
Total 1 obs	$2 \cdot 10^4$	$5.1 \cdot 10^4$	$4.1 \cdot 10^4$	$4.0 \cdot 10^4$	$1 \cdot 10^5$	-
Total 3 obs	$2 \cdot 10^4$	$1.13 \cdot 10^5$	$8.3 \cdot 10^4$	$8.0 \cdot 10^4$	$3 \cdot 10^5$	-
Total 100 obs	$2 \cdot 10^4$	$3.12 \cdot 10^6$	$2.12 \cdot 10^6$	$2.02 \cdot 10^6$	$1 \cdot 10^7$	-

Table 12: **Model simulations needed for the different techniques**, for both NN training and inference; for ABC-FP, ABC-SM, ABC-SSM, PMC-SL and PMC-RE, we report simulations needed to obtain performance at least as good as Exc-SM; in case the approach does not reach the same performance as Exc-SM, we denote that by a dash. Notice that PMC-SL and PMC-RE do not require NN training before performing inference. We also show the total number of simulations needed to apply the different approaches on 1, 3 and 100 observations, by taking into account NN training and inference steps.

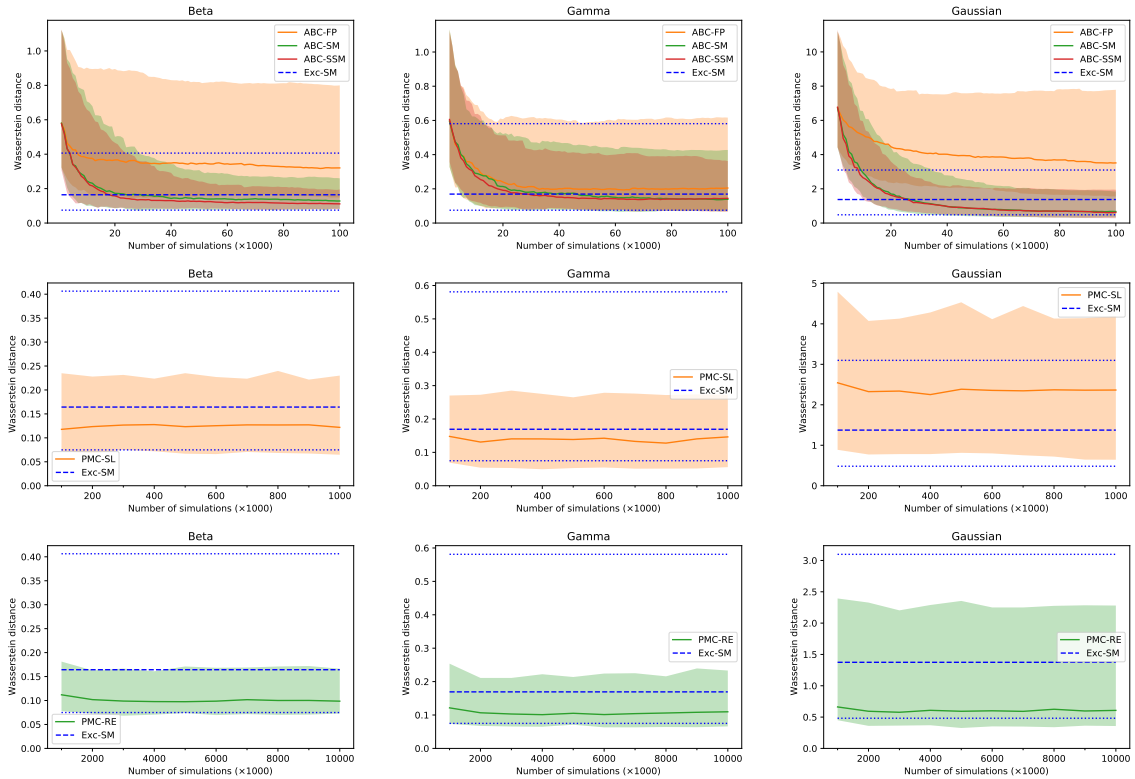


Figure 13: **Wasserstein distance between approximate and exact posterior at different iterations of the sequential algorithms for the exponential family models, for 100 different observations.** The solid line denotes median, while colored regions denote 95% probability density region; an horizontal line denoting the value obtained with Exc-SM is also represented, 95% probability density region denoted by dotted horizontal lines. The horizontal axis reports the number of model simulations corresponding to the iteration of the different algorithms.

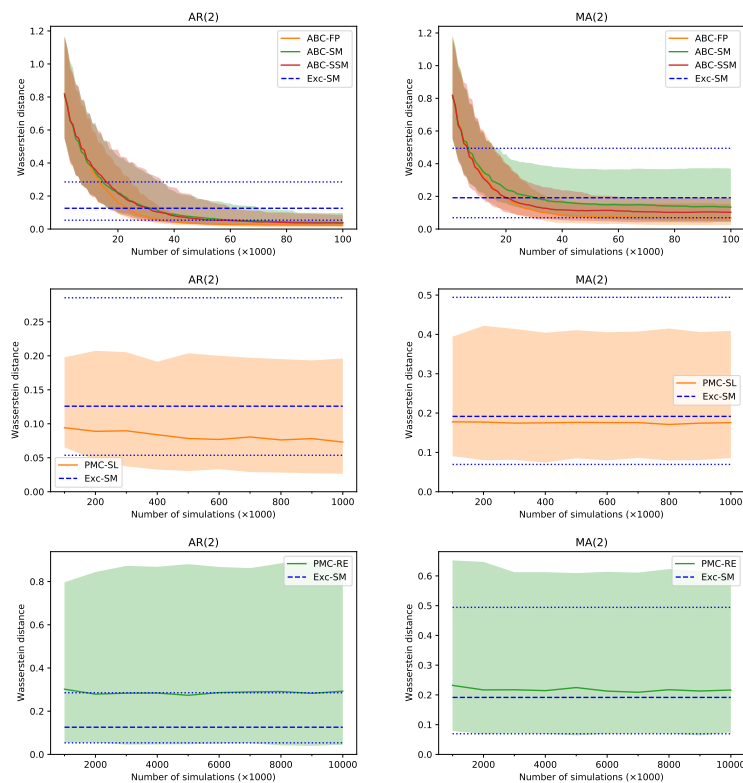


Figure 14: **Wasserstein distance between approximate and exact posterior at different iterations of the sequential algorithms for the AR(2) and MA(2) models, for 100 different observations.** The solid line denotes median, while colored regions denote 95% probability density region; an horizontal line denoting the value obtained with Exc-SM is also represented, 95% probability density region denoted by dotted horizontal lines. The horizontal axis reports the number of model simulations corresponding to the iteration of the different algorithms.

100 observations. It can be seen that ABC-FP is slightly outperformed by our proposed methods for both the large and small Lorenz96 configuration. Additionally, notice how, in the small configuration, the Energy score has an increasing trend over  $t$ , while the Kernel one instead decreases.

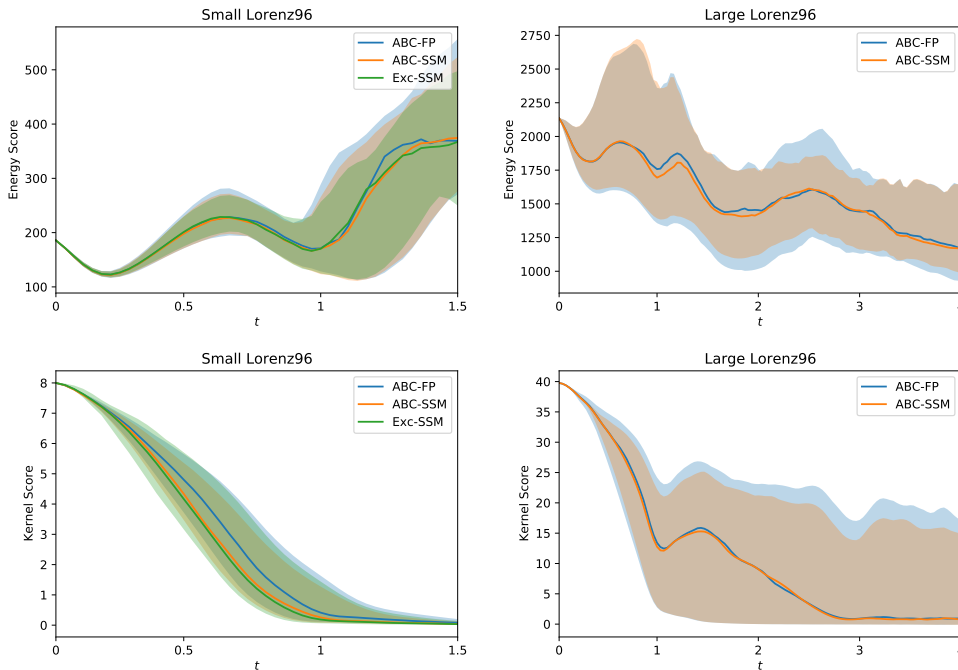


Figure 15: **Posterior predictive performance of the different methods at each timestep according to the Kernel and Energy Scores**; the smaller, the better. Samples from the posterior predictive were obtained for 100 observations, and both Scoring Rules estimated. The solid lines denote medians over the 100 observations, while colored regions denote 95% probability density region.

#### G.4.1 SETTING $\gamma$ IN THE KERNEL SCORING RULE

In the kernel score  $S_k$ , the Gaussian kernel in Eq. (17) is used across this work. There, the kernel bandwidth  $\gamma$  is a free parameter. In order to ensure comparability between the Scoring Rule values for different observations and inference methods, the value of  $\gamma$  needs to be fixed independently on both.

Inspired by Park et al. (2016), we exploit an empirical procedure to set  $\gamma$  for the specific case of multivariate time-series models (of which our Lorenz96 model is an instance, see Section 5.3). Specifically, we use the following procedure:

1. Draw a set of parameter values  $\theta_j \sim \pi(\theta)$  and simulations  $x_j \sim p(\cdot|\theta_j)$ , for  $j = 1, \dots, m$ .
2. Estimate the median of  $\{\|x_j^{(t)} - x_k^{(t)}\|_2\}_{jk}^m$  and call it  $\hat{\gamma}^{(t)}$ , for all values of  $t$ .
3. Set the estimate for  $\gamma$  as the median of  $\hat{\gamma}^{(t)}$  over all considered timesteps  $t$ .

Empirically, we use  $m = 1000$ . Note that the above strategy uses medians rather than means as those are more robust to outliers in the estimates. With this method, we obtain  $\gamma \approx 1.54$  for the small version of the Lorenz96 model and  $\gamma = 6.38$  for the large one.

## References

- Simon Aeschbacher, Mark A Beaumont, and Andreas Futschik. A novel approach for choosing summary statistics in approximate Bayesian computation. *Genetics*, 192(3):1027–1047, 2012.
- Mattias Åkesson, Prashant Singh, Fredrik Wrede, and Andreas Hellander. Convolutional neural networks as summary statistics for approximate Bayesian computation. *arXiv preprint arXiv:2001.11760*, 2020.
- Carlo Albert, R. Künsch Hans, and Andreas Scheidegger. A simulated annealing approach to approximate Bayesian computations. *Statistics and Computing*, 25:1217–1232, 2015.
- Pierre Alquier, Nial Friel, Richard Everitt, and Aidan Boland. Noisy Monte Carlo: Convergence of Markov chains with approximate transition kernels. *Statistics and Computing*, 26(1-2):29–47, 2016.
- Ziwen An, Leah F South, David J Nott, and Christopher C Drovandi. Accelerating Bayesian synthetic likelihood with the graphical lasso. *Journal of Computational and Graphical Statistics*, 28(2):471–475, 2019.
- Ziwen An, David J Nott, and Christopher Drovandi. Robust Bayesian synthetic likelihood via a semi-parametric approach. *Statistics and Computing*, 30(3):543–557, 2020.
- Michael Arbel and Arthur Gretton. Kernel conditional exponential family. In *International Conference on Artificial Intelligence and Statistics*, pages 1337–1346. PMLR, 2018.
- V.I. Avrutskiy. Backpropagation generalized for output derivatives. *arXiv preprint arXiv:1712.04185*, 2017.
- Mark A Beaumont. Approximate Bayesian computation in evolution and ecology. *Annual review of ecology, evolution, and systematics*, 41:379–406, 2010.
- Michael GB Blum, Maria Antonieta Nunes, Dennis Prangle, and Scott A Sisson. A comparative review of dimension reduction methods in approximate Bayesian computation. *Statistical Science*, 28(2):189–208, 2013.
- Alberto Caimo and Nial Friel. Bayesian inference for exponential random graph models. *Social Networks*, 33(1):41–55, 2011.
- Olivier Cappé, Arnaud Guillin, Jean-Michel Marin, and Christian P Robert. Population Monte Carlo. *Journal of Computational and Graphical Statistics*, 13(4):907–929, 2004.
- Bob Carpenter, Andrew Gelman, Matthew D Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell. Stan: A probabilistic programming language. *Journal of statistical software*, 76(1), 2017.



- Yanzhi Chen, Dinghuai Zhang, Michael U Gutmann, Aaron Courville, and Zhanxing Zhu. Neural approximate sufficient statistics for implicit models. In *Ninth International Conference on Learning Representations 2021*, 2021.
- Badr-Eddine Chérif-Abdellatif and Pierre Alquier. MMD-Bayes: Robust Bayesian estimation via maximum mean discrepancy. In *Symposium on Advances in Approximate Bayesian Inference*, pages 1–21. PMLR, 2020.
- Bo Dai, Hanjun Dai, Arthur Gretton, Le Song, Dale Schuurmans, and Niao He. Kernel exponential family estimation via doubly dual embedding. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 2321–2330. PMLR, 2019a.
- Bo Dai, Zhen Liu, Hanjun Dai, Niao He, Arthur Gretton, Le Song, and Dale Schuurmans. Exponential family estimation via adversarial dynamics embedding. In *Advances in Neural Information Processing Systems*, pages 10979–10990, 2019b.
- Alexander Philip Dawid and Monica Musio. Theory and applications of proper scoring rules. *Metron*, 72(2):169–183, 2014.
- Pierre Del Moral, Arnaud Doucet, and Ajay Jasra. An adaptive sequential Monte Carlo method for approximate Bayesian computation. *Statistics and Computing*, 22(5):1009–1020, 2012.
- Ritabrata Dutta, Bastien Chopard, Jonas Lätt, Frank Dubois, Karim Zouaoui Boudjeltia, and Antonietta Mira. Parameter estimation of platelets deposition: Approximate Bayesian computation with high performance computing. *Frontiers in physiology*, 9, 2018.
- Ritabrata Dutta, Susana N Gomes, Dante Kalise, and Lorenzo Pacchiardi. Using mobility data in the design of optimal lockdown strategies for the covid-19 pandemic. *PLoS Computational Biology*, 17(8):e1009236, 2021a.
- Ritabrata Dutta, Marcel Schoengens, Lorenzo Pacchiardi, Avinash Ummadisingu, Nicole Widmer, Pierre Künzli, Jukka-Pekka Onnela, and Antonietta Mira. ABCpy: A high-performance computing perspective to approximate bayesian computation. *Journal of Statistical Software*, 100(7):1–38, 2021b. doi: 10.18637/jss.v100.i07. URL <https://www.jstatsoft.org/index.php/jss/article/view/v100i07>.
- Richard G Everitt. Bayesian parameter estimation for latent Markov random fields and social networks. *Journal of Computational and graphical Statistics*, 21(4):940–960, 2012.
- Richard G Everitt, Dennis Prangle, Philip Maybank, and Mark Bell. Marginal sequential Monte Carlo for doubly intractable models. *arXiv preprint arXiv:1710.04382*, 2017.
- Matteo Fasiolo, Simon N Wood, Florian Hartig, and Mark V Bravington. An extended empirical saddlepoint approximation for intractable likelihoods. *Electronic Journal of Statistics*, 12(1):1544–1578, 2018.

- Paul Fearnhead and Dennis Prangle. Constructing summary statistics for approximate Bayesian computation: semi-automatic approximate Bayesian computation [with Discussion]. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 74(3):419–474, 2012. ISSN 1369-7412.
- Charles J Geyer. Markov chain Monte Carlo maximum likelihood. In *Computing science and statistics: Proceedings of 23rd Symposium on the Interface Interface Foundation, Fairfax Station, 1991*, pages 156–163, 1991.
- Alexander Gleim and Christian Pigorsch. Approximate Bayesian computation with indirect summary statistics. *Draft paper: <http://ect-pigorsch.mee.uni-bonn.de/data/research/papers>*, 2013.
- Tilmann Gneiting and Adrian E Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American statistical Association*, 102(477):359–378, 2007.
- David Greenberg, Marcel Nonnenmacher, and Jakob Macke. Automatic posterior transformation for likelihood-free inference. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 2404–2414. PMLR, 09–15 Jun 2019. URL <http://proceedings.mlr.press/v97/greenberg19a.html>.
- Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *The Journal of Machine Learning Research*, 13(1):723–773, 2012.
- Michael U Gutmann and Aapo Hyvärinen. Noise-contrastive estimation of unnormalized statistical models, with applications to natural image statistics. *Journal of Machine Learning Research*, 13(Feb):307–361, 2012.
- Michael U Gutmann, Jukka Corander, et al. Bayesian optimization for likelihood-free inference of simulator-based statistical models. *Journal of Machine Learning Research*, 2016.
- J Hakkarainen, A Ilin, A Solonen, M Laine, H Haario, J Tamminen, E Oja, and H Järvinen. On closure parameter estimation in chaotic systems. *Nonlinear processes in Geophysics*, 19(1):127–143, 2012.
- Joeri Hermans, Volodimir Begy, and Gilles Louppe. Likelihood-free MCMC with amortized approximate ratio estimators. In *International Conference on Machine Learning*, pages 4239–4248. PMLR, 2020.
- Geoffrey E Hinton. Training products of experts by minimizing contrastive divergence. *Neural computation*, 14(8):1771–1800, 2002.
- Harold Hotelling. Relations between two sets of variates. *Biometrika*, 28(3/4):321–377, 1936.
- Aapo Hyvärinen. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(Apr):695–709, 2005.

- Aapo Hyvärinen. Some extensions of score matching. *Computational statistics & data analysis*, 51(5):2499–2512, 2007.
- Sergey Ioffe and Christian Szegedy. Batch Normalization: Accelerating deep network training by reducing internal covariate shift. In Francis Bach and David Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 448–456, Lille, France, 07–09 Jul 2015. PMLR. URL <http://proceedings.mlr.press/v37/ioffe15.html>.
- Marko Järvenpää, Michael U Gutmann, Arius Pleska, Aki Vehtari, and Pekka Marttinen. Efficient acquisition rules for model-based approximate Bayesian computation. *Bayesian Analysis*, 14(2):595–622, 2019.
- Marko Jarvenpaa, Aki Vehtari, and Pekka Marttinen. Batch simulations and uncertainty quantification in gaussian process surrogate approximate Bayesian computation. In *Conference on Uncertainty in Artificial Intelligence*, pages 779–788. PMLR, 2020.
- Marko Järvenpää, Michael U Gutmann, Aki Vehtari, and Pekka Marttinen. Parallel Gaussian process surrogate Bayesian inference with noisy likelihood evaluations. *Bayesian Analysis*, 16(1):147–178, 2021.
- Bai Jiang, Tung-yu Wu, Charles Zheng, and Wing H Wong. Learning summary statistic for approximate Bayesian computation via deep neural network. *Statistica Sinica*, pages 1595–1618, 2017.
- Paul Joyce and Paul Marjoram. Approximately sufficient statistics and Bayesian computation. *Statistical applications in genetics and molecular biology*, 7(1), 2008.
- Ilyes Khemakhem, Ricardo Pio Monti, Diederik P. Kingma, and Aapo Hyvärinen. ICE-BeeM: Identifiable conditional energy-based deep models. *CoRR*, abs/2002.11537, 2020. URL <https://arxiv.org/abs/2002.11537>.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL <http://arxiv.org/abs/1412.6980>.
- Nadja Klein, David J Nott, and Michael Stanley Smith. Marginally calibrated deep distributional regression. *Journal of Computational and Graphical Statistics*, pages 1–17, 2020.
- Bernard Osgood Koopman. On distributions admitting a sufficient statistic. *Transactions of the American Mathematical society*, 39(3):399–409, 1936.
- Olivier Ledoit and Michael Wolf. A well-conditioned estimator for large-dimensional covariance matrices. *Journal of multivariate analysis*, 88(2):365–411, 2004.
- Faming Liang. A double Metropolis–Hastings sampler for spatial models with intractable normalizing constants. *Journal of Statistical Computation and Simulation*, 80(9):1007–1022, 2010.

- Faming Liang, Ick Hoon Jin, Qifan Song, and Jun S Liu. An adaptive exchange algorithm for sampling from distributions with intractable normalizing constants. *Journal of the American Statistical Association*, 111(513):377–393, 2016.
- 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. ISSN 1076836X. doi: 10.1093/sysbio/syw077. URL <https://doi.org/10.1093/sysbio/syw077>.
- Qiang Liu and Dilin Wang. Stein variational gradient descent: A general purpose Bayesian inference algorithm. In *Advances in neural information processing systems*, pages 2378–2386, 2016.
- Song Liu and Takafumi Kanamori. Estimating density models with complex truncation boundaries. *arXiv preprint arXiv:1910.03834*, 2019.
- Edward N Lorenz. Predictability: A problem partly solved. In *Proc. Seminar on predictability*, volume 1, 1996.
- 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. In *Advances in Neural Information Processing Systems*, pages 1289–1299, 2017.
- Jan-Matthis Lueckmann, Giacomo Bassetto, Theofanis Karaletsos, and Jakob H Macke. Likelihood-free inference with emulator networks. In *Symposium on Advances in Approximate Bayesian Inference*, pages 32–53. PMLR, 2019.
- Kanti V Mardia, John T Kent, and Arnab K Laha. Score matching estimators for directional distributions. *arXiv preprint arXiv:1604.08470*, 2016.
- Jean-Michel Marin, Pierre Pudlo, Christian P Robert, and Robin J Ryder. Approximate Bayesian computational methods. *Statistics and Computing*, 22(6):1167–1180, 2012.
- Paul Marjoram, John Molitor, Vincent Plagnol, and Simon Tavaré. Markov chain Monte Carlo without likelihoods. *Proceedings of the National Academy of Sciences*, 100(26):15324–15328, 2003.
- Pekka Marttinen, Nicholas J Croucher, Michael U Gutmann, Jukka Corander, and William P Hanage. Recombination produces coherent bacterial species clusters in both core and accessory genomes. *Microbial Genomics*, 1(5), 2015.
- Takuo Matsubara, Jeremias Knoblauch, François-Xavier Briol, Chris Oates, et al. Robust generalised Bayesian inference for intractable likelihoods. *arXiv preprint arXiv:2104.07359*, 2021.
- Trevelyan J McKinley, Ian Vernon, Ioannis Andrianakis, Nicky McCreesh, Jeremy E Oakley, Rebecca N Nsubuga, Michael Goldstein, Richard G White, et al. Approximate Bayesian computation and simulation-based inference for complex stochastic epidemic models. *Statistical science*, 33(1):4–18, 2018.

- Edward Meeds and Max Welling. GPS-ABC: Gaussian process surrogate approximate Bayesian computation. In *Proceedings of the Thirtieth Conference on Uncertainty in Artificial Intelligence*, pages 593–602, 2014.
- Amanda Minter and Renata Retkute. Approximate Bayesian computation for infectious disease modelling. *Epidemics*, 29:100368, 2019.
- Matthew T Moores, Christopher C Drovandi, Kerrie Mengersen, and Christian P Robert. Pre-processing for approximate Bayesian computation in image analysis. *Statistics and Computing*, 25(1):23–33, 2015.
- Iain Murray, Zoubin Ghahramani, and David MacKay. MCMC for doubly-intractable distributions. *arXiv preprint arXiv:1206.6848*, 2012.
- Hien Duy Nguyen, Julyan Arbel, Hongliang Lü, and Florence Forbes. Approximate Bayesian computation via the energy statistic. *IEEE Access*, 8:131683–131698, 2020.
- Matthew A Nunes and David J Balding. On optimal selection of summary statistics for approximate Bayesian computation. *Statistical applications in genetics and molecular biology*, 9(1), 2010.
- Lorenzo Pacchiardi, Pierre Künzli, Marcel Schöngens, Bastien Chopard, and Ritabrata Dutta. Distance-learning for approximate Bayesian computation to model a volcanic eruption. *Sankhya B*, Jan 2020. ISSN 0976-8394. doi: 10.1007/s13571-019-00208-8. URL <https://doi.org/10.1007/s13571-019-00208-8>.
- George Papamakarios and Iain Murray. Fast  $\varepsilon$ -free inference of simulation models with Bayesian conditional density estimation. In *Advances in Neural Information Processing Systems*, pages 1028–1036, 2016.
- George Papamakarios, David Sterratt, and Iain Murray. Sequential neural likelihood: Fast likelihood-free inference with autoregressive flows. In Kamalika Chaudhuri and Masashi Sugiyama, editors, *Proceedings of Machine Learning Research*, volume 89 of *Proceedings of Machine Learning Research*, pages 837–848. PMLR, 16–18 Apr 2019. URL <http://proceedings.mlr.press/v89/papamakarios19a.html>.
- George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. *Journal of Machine Learning Research*, 22(57):1–64, 2021. URL <http://jmlr.org/papers/v22/19-1028.html>.
- Jaewoo Park and Murali Haran. Bayesian inference in the presence of intractable normalizing functions. *Journal of the American Statistical Association*, 113(523):1372–1390, 2018.
- Mijung Park, Wittawat Jitkrittum, and Dino Sejdinovic. K2-ABC: Approximate Bayesian computation with kernel embeddings. In *Artificial Intelligence and Statistics*, 2016.

- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.
- 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.
- Stefan T Radev, Ulf K Mertens, Andreas Voss, Lynton Ardizzone, and Ullrich Köthe. BayesFlow: Learning complex stochastic models with invertible neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 2020.
- Maria L Rizzo and Gábor J Székely. Energy distance. *Wiley interdisciplinary reviews: Computational statistics*, 8(1):27–38, 2016.
- Gareth O Roberts, Andrew Gelman, Walter R Gilks, et al. Weak convergence and optimal scaling of random walk Metropolis algorithms. *The annals of applied probability*, 7(1):110–120, 1997.
- Erlis Ruli, Nicola Sartori, and Laura Ventura. Approximate Bayesian computation with composite score functions. *Statistics and Computing*, 26(3):679–692, 2016.
- Jascha Sohl-Dickstein, Peter B Battaglino, and Michael R DeWeese. New method for parameter estimation in probabilistic models: minimum probability flow. *Physical review letters*, 107(22):220601, 2011.
- Yang Song, Sahaj Garg, Jiaxin Shi, and Stefano Ermon. Sliced score matching: A scalable approach to density and score estimation. In Ryan P. Adams and Vibhav Gogate, editors, *Proceedings of The 35th Uncertainty in Artificial Intelligence Conference*, volume 115 of *Proceedings of Machine Learning Research*, pages 574–584, Tel Aviv, Israel, 22–25 Jul 2020. PMLR. URL <http://proceedings.mlr.press/v115/song20a.html>.
- Bharath Sriperumbudur, Kenji Fukumizu, Arthur Gretton, Aapo Hyvärinen, and Revant Kumar. Density estimation in infinite dimensional exponential families. *The Journal of Machine Learning Research*, 18(1):1830–1888, 2017.
- Esteban G Tabak, Giulio Trigila, and Wenjun Zhao. Conditional density estimation and simulation through optimal transport. *Machine Learning*, pages 1–24, 2020.
- Simon Tavaré, David J Balding, Robert C Griffiths, and Peter Donnelly. Inferring coalescence times from DNA sequence data. *Genetics*, 145(2):505–518, 1997.
- Owen Thomas, Ritabrata Dutta, Jukka Corander, Samuel Kaski, Michael U Gutmann, et al. Likelihood-free inference by ratio estimation. *Bayesian Analysis*, 2020.

- Jean-Francois Ton, CHAN Lucian, Yee Whye Teh, and Dino Sejdinovic. Noise contrastive meta-learning for conditional density estimation using kernel mean embeddings. In *International Conference on Artificial Intelligence and Statistics*, pages 1099–1107. PMLR, 2021.
- Tina Toni, David Welch, Natalja Strelkowa, Andreas Ipsen, and Michael PH Stumpf. Approximate Bayesian computation scheme for parameter inference and model selection in dynamical systems. *Journal of the Royal Society Interface*, 6(31):187–202, 2009.
- Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural computation*, 23(7):1661–1674, 2011.
- Ziyu Wang, Shuyu Cheng, Li Yueru, Jun Zhu, and Bo Zhang. A Wasserstein minimum velocity approach to learning unnormalized models. In Silvia Chiappa and Roberto Calandra, editors, *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pages 3728–3738. PMLR, 26–28 Aug 2020. URL <https://proceedings.mlr.press/v108/wang20j.html>.
- Li Wenliang, Dougal Sutherland, Heiko Strathmann, and Arthur Gretton. Learning deep kernels for exponential family densities. In *International Conference on Machine Learning*, pages 6737–6746. PMLR, 2019.
- Richard Wilkinson. Accelerating ABC methods using Gaussian processes. In *Artificial Intelligence and Statistics*, pages 1015–1023. PMLR, 2014.
- Daniel S Wilks. Effects of stochastic parametrizations in the Lorenz’96 system. *Quarterly Journal of the Royal Meteorological Society*, 131(606):389–407, 2005.
- Samuel Wqvist, Pierre-Alexandre Mattei, Umberto Picchini, and Jes Frellsen. Partially exchangeable networks and architectures for learning summary statistics in approximate Bayesian computation. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 6798–6807. PMLR, 09–15 Jun 2019. URL <http://proceedings.mlr.press/v97/wqvist19a.html>.
- Simon N Wood. Statistical inference for noisy nonlinear ecological dynamic systems. *Nature*, 466(7310):1102, 2010.
- Shiqing Yu, Mathias Drton, and Ali Shojaie. Generalized score matching for non-negative data. *Journal of Machine Learning Research*, 20(76):1–70, 2019. URL <http://jmlr.org/papers/v20/18-278.html>.