

LEARNING SUMMARY STATISTIC FOR APPROXIMATE BAYESIAN COMPUTATION VIA DEEP NEURAL NETWORK

Bai Jiang, Tung-Yu Wu, Charles Zheng and Wing H. Wong

Stanford University

Abstract: Approximate Bayesian Computation (ABC) methods are used to approximate posterior distributions in models with unknown or computationally intractable likelihoods. Both the accuracy and computational efficiency of ABC depend on the choice of summary statistic, but outside of special cases where the optimal summary statistics are known, it is unclear which guiding principles can be used to construct effective summary statistics. In this paper we explore the possibility of automating the process of constructing summary statistics by training deep neural networks to predict the parameters from artificially generated data: the resulting summary statistics are approximately posterior means of the parameters. With minimal model-specific tuning, our method constructs summary statistics for the Ising model and the moving-average model, which match or exceed theoretically-motivated summary statistics in terms of the accuracies of the resulting posteriors.

Key words and phrases: Approximate Bayesian computation, deep learning, summary statistic.

1. Introduction

1.1. Approximate Bayesian computation

Bayesian inference is traditionally centered around the ability to compute or sample from the posterior distribution of the parameters, having conditioned on the observed data. Suppose data X is generated from a model \mathcal{M} with parameter θ , the prior of which is denoted by $\pi(\theta)$. If the closed form of the likelihood function $l(\theta) = p(X|\theta)$ is available, the posterior distribution of θ given observed data x_{obs} can be computed via Bayes' rule

$$\pi(\theta|x_{obs}) = \frac{\pi(\theta)p(x_{obs}|\theta)}{p(x_{obs})}.$$

Alternatively, if the likelihood function can only be computed conditionally or up to a normalizing constant, one can still draw samples from the posterior by using

stochastic simulation techniques such as Markov Chain Monte Carlo (MCMC) and rejection sampling (Asmussen and Glynn (2007)).

In many applications, the likelihood function $l(\theta) = p(X|\theta)$ cannot be explicitly obtained, or is intractable to compute; this precludes the possibility of direct computation or MCMC sampling. In these cases, approximate inference can still be performed as long as 1) it is possible to draw θ from the prior $\pi(\theta)$, and 2) it is possible to simulate X from the model \mathcal{M} given θ , using the methods of Approximate Bayesian Computation (ABC) (See e.g. Beaumont, Zhang and Balding (2002); Toni et al. (2009); Lopes and Beaumont (2010); Beaumont (2010); Csilléry et al. (2010); Marin et al. (2012); Sunnåker et al. (2013)).

While many variations of the core approach exist, the fundamental idea underlying ABC is quite simple: that one can use rejection sampling to obtain draws from the posterior distribution $\pi(\theta|x_{obs})$ without computing any likelihoods. We draw parameter-data pairs (θ', X') from the prior $\pi(\theta)$ and the model \mathcal{M} , and accept only the θ' such that $X' = x_{obs}$, which occurs with conditional probability $P(X = x_{obs}|\theta')$ for any θ' . Algorithm 1 describes the ABC method for discrete data (Tavaré et al. (1997)), which yields an i.i.d. sample $\{\theta^{(i)}\}_{1 \leq i \leq n}$ of the exact posterior distribution $\pi(\theta|X = x_{obs})$.

Algorithm 1 ABC rejection sampling 1

```

for  $i = 1, \dots, n$  do
  repeat
    Propose  $\theta' \sim \pi(\theta)$ 
    Draw  $X' \sim \mathcal{M}$  given  $\theta'$ 
  until  $X' = x_{obs}$  (acceptance criterion)
  Accept  $\theta'$  and let  $\theta^{(i)} = \theta'$ 
end for

```

The success of Algorithm 1 depends on acceptance rate of proposed parameter θ' . For continuous x_{obs} and X' , the event $X' = x_{obs}$ happens with probability 0, and hence Algorithm 1 is unable to produce any draws. As a remedy, one can relax the acceptance criterion $X' = x_{obs}$ to be $\|X' - x_{obs}\| < \epsilon$, where $\|\cdot\|$ is a norm and ϵ is the tolerance threshold. The choice of ϵ is crucial for balancing efficiency and approximation error, since with smaller ϵ the approximation error decreases while the acceptance probability also decreases.

1.2. Summary statistic

When data vectors x_{obs} , X are high-dimensional, the inefficiency of rejection sampling in high dimensions results in either extreme inaccuracy, or accuracy

at the expense of an extremely time-consuming procedure. To circumvent the problem, one can introduce low-dimensional summary statistic S and further relax the acceptance criterion to be $\|S(X') - S(x_{obs})\| < \epsilon$. The use of summary statistics results in Algorithm 2, which was first proposed as the extension of Algorithm 1 in population genetics application (Fu and Li (1997); Weiss and von Haeseler (1998); Pritchard et al. (1999)).

Algorithm 2 ABC rejection sampling 2

```

for  $i = 1, \dots, n$  do
  repeat
    Propose  $\theta' \sim \pi$ 
    Draw  $X' \sim \mathcal{M}$  with  $\theta'$ 
  until  $\|S(X') - S(x_{obs})\| < \epsilon$  (relaxed acceptance criterion)
  Accept  $\theta'$  and let  $\theta^{(i)} = \theta'$ 
end for

```

Instead of the exact posterior distribution, the resulting sample $\{\theta^{(i)}\}_{1 \leq i \leq n}$ obtained by Algorithm 2 follows an approximate posterior distribution

$$\pi(\theta \| \|S(X') - S(x_{obs})\| < \epsilon) \approx \pi(\theta | S(X) = S(x_{obs})) \quad (1.1)$$

$$\approx \pi(\theta | X = x_{obs}). \quad (1.2)$$

The choice of the summary statistic is crucial for the approximation quality of ABC posterior distribution. An effective summary statistic should offer a good trade-off between two approximation errors (Blum et al. (2013)). The approximation error (1.1) is introduced when one replaces “equal” with “similar” in the first relaxation of the acceptance criterion. Under appropriate regularity conditions, it vanishes as $\epsilon \rightarrow 0$. The approximation error (1.2) is introduced when one compares summary statistics $S(X)$ and $S(x_{obs})$ rather than the original data X and x_{obs} . In essence, this is just the information loss of mapping high-dimensional X to low-dimensional $S(X)$. A summary statistic S of higher dimension is in general more informative, hence reduces the approximation error (1.2). At the same time, increasing the dimension of the summary statistic slows down the rate that the approximation error (1.1) vanishes in the limit of $\epsilon \rightarrow 0$. Ideally, we seek a statistic which is simultaneously low-dimensional and informative.

A sufficient statistic is an attractive option, since sufficiency, by definition, implies that the approximation error (1.2) is zero (Kolmogorov (1942); Lehmann and Casella (1998)). However, the sufficient statistic has generally the same dimensionality as the sample size, except in special cases such as exponential families. And even when a low-dimensional sufficient statistic exists, it may be

intractable to compute.

The main task of this article is to construct low-dimensional and informative summary statistics for ABC methods. Since our goal is to compare methods of constructing summary statistics (rather than present a complete methodology for ABC), the relatively simple Algorithm 2 suffices. In future work, we plan to use our approach for constructing summary statistics alongside more sophisticated variants of ABC methods, such as those which combine ABC with Markov chain Monte Carlo or sequential techniques (Marjoram et al. (2003); Sisson, Fan and Tanaka (2007)). Hereafter all ABC procedures mentioned use Algorithm 2.

1.3. Related work and our DNN approach

Existing methods for constructing summary statistics can be roughly classified into two classes, both of which require a set of candidate summary statistics $\mathcal{S}_c = \{S_{c,k}\}_{1 \leq k \leq K}$ as input. The first class consists of approaches for *best subset selection*. Subsets of \mathcal{S}_c are evaluated according to various information-based criteria, e.g. measure of sufficiency (Joyce and Marjoram (2008)), entropy (Nunes and Balding (2010)), Akaike and Bayesian information criteria (Blum et al. (2013)), and the “best” subset is chosen to be the summary statistic. The second class is *linear regression* approach, which constructs summary statistics by linear regression of response θ on candidate summary statistics \mathcal{S}_c (Wegmann, Leuenberger and Excoffier (2009); Fearnhead and Prangle (2012)). Regularization techniques have also been considered to reduce overfitting in the regression models (Blum et al. (2013)). Many of these methods rely on expert knowledge to provide candidate summary statistics.

In this paper, we propose to automatically learn summary statistics for high-dimensional X by using deep neural networks (DNN). Here DNN is expected to effectively learn a good approximation to the posterior mean $\mathbb{E}_\pi[\theta|X]$ when constructing a minimum squared error estimator $\hat{\theta}(X)$ on a large data set $\{(\theta^{(i)}, X^{(i)})\}_{1 \leq i \leq N} \sim \pi \times \mathcal{M}$. The minimization problem is given by

$$\min_{\beta} \frac{1}{N} \sum_{i=1}^N \left\| f_{\beta}(X^{(i)}) - \theta^{(i)} \right\|_2^2,$$

where f_{β} denotes a DNN with parameter β . The resulting estimator $\hat{\theta}(X) = f_{\hat{\beta}}(X)$ approximates $\mathbb{E}_\pi[\theta|X]$ and further serves as the summary statistic for ABC.

Our motivation for using (an approximation to) $\mathbb{E}_\pi[\theta|X]$ as a summary statistic for ABC is inspired by the semi-automatic method in (Fearnhead and Prangle

(2012)). Their idea is that $\mathbb{E}_\pi[\theta|X]$ as summary statistic leads to an ABC posterior, which has the same mean as the exact posterior in the limit of $\epsilon \rightarrow 0$. Therefore they proposed to linearly regress θ on candidate summary statistics $\mathcal{S}_c = \{S_{c,k}\}_{1 \leq k \leq K}$

$$\min_{\beta} \frac{1}{N} \sum_{i=1}^N \left\| \beta_0 + \sum_{k=1}^K \beta_k S_{c,k}(X^{(i)}) - \theta^{(i)} \right\|_2^2,$$

and to use the resulting minimum squared error estimator $\hat{\theta}(X)$ as the summary statistic for ABC. In their semi-automatic method, \mathcal{S}_c could be expert-designed statistics or polynomial bases (e.g. power terms of each component X_j).

Our DNN approach aims to achieve a more accurate approximation $\hat{\theta}(X) \approx \mathbb{E}_\pi[\theta|X]$ and a higher degree of automation in constructing summary statistics than the semi-automatic method. First, DNN with multiple hidden layers offers stronger representational power, compared to the semi-automatic method using linear regression. A DNN is expected to better approximate $\mathbb{E}_\pi[\theta|X]$ if the posterior mean is a highly non-linear function of X . Second, DNNs simply use the original data vector X as the input, and automatically learn the appropriate nonlinear transformations as summaries from the raw data, in contrast to the semi-automatic method and many other existing methods requiring a set of expert-designed candidate summary statistics or a basis expansion. Therefore our approach achieves a higher degree of automation in constructing summary statistics.

Blum and François (2010) have considered fitting a feed-forward neural network (FFNN) with single hidden layer by regressing $\theta^{(i)}$ on $X^{(i)}$. Their method significantly differs from ours, as theirs was originally motivated by reducing the error between the ABC posterior and the true posterior, rather than constructing summary statistics. Specifically, their method assumes that the appropriate summary statistic S has already been given, and adjusts each draw (θ, X) from the ABC procedure using summary statistic S in the way

$$\theta^* = m(S(x_{obs})) + [\theta - m(S(X))] \times \frac{\sigma(S(x_{obs}))}{\sigma(S(X))}.$$

Both $m(\cdot)$ and $\sigma(\cdot)$ are non-linear functions represented by FFNNs. Another key difference is the network size: the FFNNs in Blum and François (2010) contained four hidden neurons in order to reduce dimensionality of summary statistics, while our DNN approach contains hundreds of hidden neurons in order to gain representational power.

1.4. Organization

The rest of the article is organized as follows. In Section 2, we show how to approximate the posterior mean $\mathbb{E}_\pi[\theta|X]$ by training DNNs. In Sections 3 and 4, we report simulation studies on the Ising model and the moving average model of order 2, respectively. We describe in the supplementary materials the implementation details of DNNs and how consistency can be obtained by using the posterior mean of a basis of functions of the parameters.

2. Methods

Throughout the paper, we denote by $X \in \mathbb{R}^p$ the data, and by $\theta \in \mathbb{R}^q$ the parameter. We assume it is possible to obtain a large number of independent draws X from the model \mathcal{M} given θ despite the unavailability of $p(X|\theta)$. Denote by x_{obs} the observed data, π the prior of θ , S the summary statistic, $\|\cdot\|$ the norm to measure $S(X) - S(x_{obs})$, and ϵ the tolerance threshold. Let $\pi_{ABC}^\epsilon(\theta) = \pi(\theta \| \|S(X) - S(x_{obs})\| < \epsilon)$ denote the approximate posterior distribution obtained by Algorithm 2.

The main task is to construct a low-dimensional and informative summary statistic S for high-dimensional X , which will enable accurate approximation of π_{ABC}^ϵ . We are interested mainly in the regime where ABC is most effective: settings in which the dimension of X is moderately high (e.g. $p = 100$) and the dimension of θ is low (e.g. $q = 1, 2, 3$). Given a prior π for θ , our approach is as follows.

- (1) Generate a data set $\{(\theta^{(i)}, X^{(i)})\}_{1 \leq i \leq N}$ by repeatedly drawing $\theta^{(i)}$ from π and drawing $X^{(i)}$ from \mathcal{M} with $\theta^{(i)}$.
- (2) Train a DNN with $\{X^{(i)}\}_{1 \leq i \leq N}$ as input and $\{\theta^{(i)}\}_{1 \leq i \leq N}$ as target.
- (3) Run ABC Algorithm 2 with prior π and the DNN estimator $\hat{\theta}(X)$ as summary statistic.

Our motivation for training such a DNN is that the resulting statistic (estimator) should approximate the posterior mean $S(X) = \hat{\theta}(X) \approx \mathbb{E}_\pi[\theta|X]$.

2.1. Posterior mean as summary statistic

The main advantage of using the posterior mean $\mathbb{E}_\pi[\theta|X]$ as a summary statistic is that the ABC posterior $\pi_{ABC}^\epsilon(\theta) = \pi(\theta \| \|S(X) - S(x_{obs})\| < \epsilon)$ will then have the same mean as the exact posterior in the limit of $\epsilon \rightarrow 0$. That is to say, $\mathbb{E}_\pi[\theta|X]$ does not lose any first-order information when summarizing X .

This theoretical result has been discussed in Theorem 3 in Fearnhead and Prangle (2012), but their proof is not rigorous. We provide in Theorem 1 a more rigorous and general proof.

Theorem 1. *If $\mathbb{E}_\pi[|\theta|] < \infty$, then $S(x) = \mathbb{E}_\pi[\theta|X = x]$ is well defined. The ABC procedure with observed data x_{obs} , summary statistics S , norm $\|\cdot\|$, and tolerance threshold ϵ produces a posterior distribution*

$$\pi_{ABC}^\epsilon(\theta) = \pi(\theta|\|S(X) - S(x_{obs})\| < \epsilon),$$

with

$$\begin{aligned} \|\mathbb{E}_{\pi_{ABC}^\epsilon}[\theta] - S(x_{obs})\| &< \epsilon, \\ \lim_{\epsilon \rightarrow 0} \mathbb{E}_{\pi_{ABC}^\epsilon}[\theta] &= \mathbb{E}_\pi[\theta|X = x_{obs}]. \end{aligned}$$

Proof. First, we show $S(X) = \mathbb{E}_\pi[\theta|X]$ is a version of conditional expectation of θ given $S(X)$. Denote by $\sigma(X), \sigma(S(X))$ the σ -algebras of X and $S(X)$, respectively. $S(X)$ is clearly measurable with respect to $\sigma(X)$, thus $\sigma(S(X)) \subseteq \sigma(X)$. Then

$$\begin{aligned} S(X) &= \mathbb{E}_\pi[S(X)|S(X)] && [S(X) \text{ is known in } \sigma(S(X))] \\ &= \mathbb{E}_\pi[\mathbb{E}_\pi[\theta|X]|S(X)] && [\text{Definition of } S] \\ &= \mathbb{E}_\pi[\theta|S(X)] && [\text{Tower property, } \sigma(S(X)) \subseteq \sigma(X)]. \end{aligned}$$

As $A = \{\|S(X) - S(x_{obs})\| < \epsilon\} \in \sigma(S(X))$, we have by the definition of conditional expectation

$$\mathbb{E}_\pi[\theta\mathbb{I}_A] = \mathbb{E}_\pi[S(X)\mathbb{I}_A].$$

It follows that

$$\mathbb{E}_{\pi_{ABC}^\epsilon}[\theta] = \mathbb{E}_\pi[\theta|A] = \mathbb{E}_\pi[S(X)|A],$$

implying by Jensen's inequality that

$$\begin{aligned} \|\mathbb{E}_{\pi_{ABC}^\epsilon}[\theta] - S(x_{obs})\| &= \|\mathbb{E}_\pi[S(X)|A] - S(x_{obs})\| \\ &\leq \mathbb{E}_\pi[\|S(X) - S(x_{obs})\||A] \\ &< \epsilon. \end{aligned}$$

Letting $\epsilon \rightarrow 0$ yields $\mathbb{E}_{\pi_{ABC}^\epsilon}[\theta] \rightarrow S(x_{obs}) = \mathbb{E}_\pi[\theta|X = x_{obs}]$.

ABC procedures often give the sample mean of the ABC posterior as the point estimate for θ . Theorem 1 shows ABC procedure using $\mathbb{E}_\pi[\theta|X]$ as the summary statistic maximizes the point-estimation accuracy in the sense that the exact mean of ABC posterior $\mathbb{E}_{\pi_{ABC}^\epsilon}[\theta]$ is an ϵ -approximation to the Bayes estimator $\mathbb{E}_\pi[\theta|X = x_{obs}]$ under squared error loss.

Users of Bayesian inference generally desire more than just point estimates: ideally, one approximates the posterior $\pi(\theta|x_{obs})$ globally. We observe that such a global approximation result is possible when extending Theorem 1: if one considers a basis of functions on the parameters, $b(\theta) = (b_1(\theta), \dots, b_K(\theta))$, and uses the K -dimensional statistic $\mathbb{E}_\pi[b(\theta)|X]$ as the summary statistic(s), the ABC posterior weakly converges to the exact posterior as $\epsilon \rightarrow 0$ and $K \rightarrow \infty$ at the appropriate rate. We state this result in the supplementary material.

There is a nice connection between the posterior mean and the sufficient statistics, especially minimal sufficient statistics in the exponential family. If there exists a sufficient statistic S^* for θ , then from the concept of the sufficiency in the Bayesian context (Kolmogorov (1942)) it follows that for almost every x , $\pi(\theta|X = x) = \pi(\theta|S^*(X) = S^*(x))$, and further $S(x) = \mathbb{E}_\pi[\theta|X = x] = \mathbb{E}_\pi[\theta|S^*(X) = S^*(x)]$ is a function of $S^*(x)$. In the special case of an exponential family with minimal sufficient statistic S^* and parameter θ , the posterior mean $S(X) = \mathbb{E}_\pi[\theta|X]$ is a one-to-one function of $S^*(X)$, and thus is a minimal sufficient statistic.

2.2. Structure of deep neural network

At a high level, a deep neural network merely represents a non-linear function for transforming input vector X into output $\hat{\theta}(X)$. The structure of a neural network can be described as a series of L nonlinear transformations applied to X . Each of these L transformations is described as a *layer*: where the original input is X , the output of the first transformation is the 1st layer, the output of the second transformation is the 2nd layer, and so on, with the output as the $(L+1)$ th layer. The layers 1 to L are called *hidden layers* because they represent intermediate computations, and we let $H^{(l)}$ denote the l -th hidden layer. Then the explicit form of the network is

$$\begin{aligned} H^{(1)} &= \tanh(W^{(0)}H^{(0)} + b^{(0)}), \\ H^{(2)} &= \tanh(W^{(1)}H^{(1)} + b^{(1)}), \\ &\dots \\ H^{(L)} &= \tanh(W^{(L-1)}H^{(L-1)} + b^{(L-1)}), \\ \hat{\theta} &= W^{(L)}H^{(L)} + b^{(L)}, \end{aligned}$$

where $H^{(0)} = X$ is the input, $\hat{\theta}$ is the output, $W^{(l)}$ and $b^{(l)}$ are the parameters controlling how the inputs of layer l are transformed into the outputs of layer l . Let $n^{(l)}$ denote the size of the l -th layer: then $W^{(l)}$ is an $n^{(l+1)} \times n^{(l)}$ matrix,




Figure 1. An example of DNN with three hidden layers.




Figure 2. Neuron j in the hidden layer $l + 1$.

called the *weight matrix*, and $b^{(l)}$ is an $n^{(l+1)}$ -dimensional vector, called the *bias vector*. The $n^{(l)}$ components of each layer $H^{(l)}$ are also described evocatively as “neurons” or “hidden units”. Figure 1 illustrates an example of 3-layer DNN with input $X \in \mathbb{R}^4$ and 5/5/3 neurons in the 1st/2nd/3rd hidden layer, respective.

The role of layer $l + 1$ is to apply a nonlinear transformation to the outputs of layer l , $H^{(l)}$, and then output the transformed outputs as $H^{(l+1)}$. First, a linear transformation is applied to the previous layer $H^{(l)}$, yielding $W^{(l)}H^{(l)} + b^{(l)}$. The nonlinearity (in this case \tanh) is applied to each element of $W^{(l)}H^{(l)} + b^{(l)}$ to yield the output of the current layer, $H^{(l+1)}$. The nonlinearity is traditionally called the “activation” function, drawing an analogy to the properties of biological neurons. We choose the function \tanh as an activation function due to smoothness and computational convenience. Other popular choices for activation function are $\text{sigmoid}(t) = 1/(1 + \exp(-t))$ and $\text{ReLU}(t) = \max\{t, 0\}$. To better explain the activity of each individual neuron, we illustrate how neuron j in the hidden layer $l + 1$ works in Figure 2.

The output layer takes the top hidden layer $H^{(L)}$ as input and predicts $\hat{\theta} = W^{(L)}H^{(L)} + b^{(L)}$. In many existing applications of deep learning (e.g. computer vision and natural language processing), the goal is to predict a categorical target. In those cases, it is common to use a softmax transformation in the output layer. However, since our goal is prediction rather than classification, it suffices to use a linear transformation.

2.3. Approximating posterior mean by DNN

We use the DNN to construct a summary statistic: a function which maps x to an approximation of $\mathbb{E}_\pi[\theta|X]$. First, we generate a training set $\mathcal{D}_\pi = \{(\theta^{(i)}, X^{(i)}), 1 \leq i \leq N\}$ by drawing samples from the joint distribution $\pi(\theta, x)$. Next, we train the DNN to minimize the squared error loss between training target $\theta^{(i)}$ and estimation $\hat{\theta}(X^{(i)})$. Thus we minimize (2.1) with respect to the DNN parameters $\beta = (W^{(0)}, b^{(0)}, \dots, W^{(L)}, b^{(L)})$,

$$J(\beta) = \frac{1}{N} \sum_{i=1}^N \|f_\beta(X^{(i)}) - \theta^{(i)}\|_2^2. \quad (2.1)$$

We compute the derivatives using backpropagation (LeCun et al. (1998)) and optimize the objective function by stochastic gradient descent method. See the supplementary material for details.

Our approach is based on the fact that any function which minimizes the squared error risk for predicting θ from X may be viewed as an approximation of the posterior mean $\mathbb{E}_\pi[\theta|X]$. Hence, any supervised learning approach could be used to construct a prediction rule for predicting θ from x , and thereby provide an approximation of $\mathbb{E}_\pi[\theta|X]$. Since in many applications of ABC, we can expect $\mathbb{E}_\pi[\theta|X]$ to be a highly nonlinear and smooth function, it is important to choose a supervised learning approach which has the power to approximate such nonlinear smooth functions.

DNNs appear to be a good choice given their rich representational power for approximating nonlinear functions. More and more practical and theoretical results of deep learning in several areas of machine learning, especially computer vision and natural language processing (Hinton and Salakhutdinov (2006); Hinton, Osindero and Teh (2006); Bengio, Courville and Vincent (2013); Schmidhuber (2015)), show that deep architectures composed of simple learning modules in multiple layers can model high-level abstraction in high-dimensional data. It is speculated that by increasing the depth and width of the network, the DNN gains the power to approximate any continuous function; however, rigorous proof of the

approximation properties of DNNs remains an important open problem (Faragó and Lugosi (1993); Sutskever and Hinton (2008); Le Roux and Bengio (2010)). Nonetheless we expect that DNNs can effectively learn a good approximation to the posterior mean $\mathbb{E}_\pi[\theta|X]$ given a sufficiently large training set.

2.4. Avoiding overfitting

DNN consists of simple learning modules in multiple layers and thus has very rich representational power to learn very complicated relationships between the input X and the output θ . However, DNN is prone to overfitting given limited training data. In order to avoid overfitting, we consider three methods: generating a large training set, early stopping, and regularization on parameter.

Sufficiently Large Training Data. This is the fundamental way to avoid overfitting and improve the generalization, that, however, is impossible in many applications of machine learning. Fortunately, in applications of Approximate Bayesian Computation, an arbitrarily large training set can be generated by repeatedly sampling $(\theta^{(i)}, X^{(i)})$ from the prior π and the model \mathcal{M} , and dataset sampling can be parallelized. In our experiments, DNNs contains 3 hidden layers, each of which has 100 neurons, and has around $3 \times 100 \times 100 = 3 \times 10^4$ parameters, while the training set contains 10^6 data samples.


Early Stopping (Caruana, Lawrence and Giles (2000)). This divides the available data into three subsets: the training set, the validation set and the testing set. The training set is used to compute the gradient and update the parameter. At the same time, we monitor both the training error and the validation error. The validation error usually decreases as does the training error in the early phase of the training process. However, when the network begins to overfit, the validation error begins to increase and we stop the training process. The testing error is reported only for evaluation.

Regularization. This adds an extra term to the loss function that will penalize complexity in neural networks (Nowlan and Hinton (1992)). Here we consider \mathcal{L}_2 regularization (Ng (2004)) and minimize the objective function

$$J(\beta; \lambda) = \frac{1}{N} \sum_{i=1}^N \|f_\beta(X^{(i)}) - \theta^{(i)}\|_2^2 + \lambda \sum_{l=1}^L \|W^{(l)}\|_F^2, \quad (2.2)$$

where $\|W\|_F$ is the Frobenius norm of W , the square root of the sum of the absolute squares of its elements.

More sophisticated methods like dropout (Srivastava et al. (2014)) and tuning network size can probably better combat overfitting and learn better summary

Figure 3. Ising model on 4×4 lattice.

statistic. We only use the simple methods and do minimal model-specific tuning in the simulation studies. Our goal is to show a relatively simple DNN can learn a good summary statistic for ABC.

3. Example: Ising Model

3.1. ABC and summary statistics

The Ising model consists of discrete variables ($+1$ or -1) arranged in a lattice (Figure 3). Each binary variable, called a spin, is allowed to interact with its neighbors. The inverse-temperature parameter $\theta > 0$ characterizes the extent of interaction. Given θ , the probability mass function of the Ising model on $m \times m$ lattice is

$$p(X|\theta) = \frac{\exp\left(\theta \sum_{j \sim k} X_j X_k\right)}{Z(\theta)},$$

where $X_j \in \{-1, +1\}$, $j \sim k$ means X_j and X_k are neighbors, and the normalizing constant is

$$Z(\theta) = \sum_{x' \in \{-1, +1\}^{m \times m}} \exp\left(\theta \sum_{j \sim k} x'_j x'_k\right).$$

Since the normalizing constant requires an exponential-time computation, the probability mass function $p(x|\theta)$ is intractable except in small cases.

Despite the unavailability of probability mass function, data X can be still simulated given θ using Monte Carlo methods such as Metropolis algorithm (Assmussen and Glynn (2007)). It allows use of ABC for parameter inference. The sufficient statistic $S^*(X) = \sum_{j \sim k} X_j X_k$ is the ideal summary statistic, because

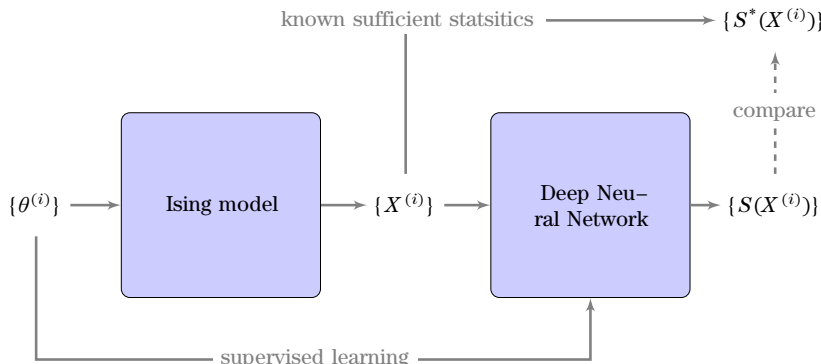


Figure 4. Experimental design on Ising model.

S^* is univariate, speeds up the convergence of approximation error (1.1) in the limit of $\epsilon \rightarrow 0$, and losses no information in the approximation (1.2).

Since S^* results in the ABC posterior with the highest quality, we take it as the gold standard and compare the DNN-based summary statistic to it. The DNN-based summary statistic, if approximating $\mathbb{E}_\pi[\theta|X]$ well, should be an approximately increasing function of $S^*(X)$. As $\mathbb{E}_\pi[\theta|X]$ is an increasing function of $S^*(X)$, it is a sufficient statistic as well. To see this, view the posterior as an exponential family with $S^*(X)$ as “parameter” and θ as “sufficient statistic”,

$$\pi(\theta|X) \propto \pi(\theta)p_\theta(X) = \underbrace{\pi(\theta)e^{-\log Z(\theta)}}_{\text{carrier measure}} \exp(S^*(X) \cdot \theta),$$

and then use the mean reparametrization result of exponential family. As $\mathbb{E}_\pi[\theta|X]$ and $S^*(X)$ are highly non-linear functions in the high-dimensional space $\{-1, +1\}^{m \times m}$, they are challenging to approximate.

3.2. Experimental design

Figure 4 outlines the whole experimental scheme. We generate a training set by the Metropolis algorithm, train a DNN to learn a summary statistic S , and then compare $S(X)$ to the “gold standard” $S^*(X)$.

Metropolis algorithm generated training, validation, testing sets of size 10^6 , 10^5 , 10^5 , respectively, from the Ising model on the 10×10 lattice with a prior $\pi(\theta) \sim \text{Exp}(\theta_c)$. The value $\theta_c = 0.4406$ is the phase transition point of Ising model on infinite lattice: when $\theta < \theta_c$, the spins tend to be disordered; when $\theta > \theta_c$ is large enough, the spins tend to have the same sign due to the strong neighbor-to-neighbor interactions (Onsager (1944)). The Ising model on a finite lattice undergoes a smooth phase transition around θ_c as θ increases, which

Table 1. The root-mean-square error (RMSE) and training time of the semi-automatic, FFNN, and DNN methods to predict θ given X . λ is the penalty coefficient in the regularized objective function (2.2). Stochastic gradient descent fits each FFNN or DNN by 200 full passes (epochs) through the training set.

| Method | Training RMSE | Testing RMSE | Time (s) |
|-------------------------|---------------|---------------|----------|
| Semi-automatic | 0.4401 | 0.4406 | 4.36 |
| FFNN, $\lambda = 0$ | 0.2541 | 0.2541 | 480.08 |
| DNN, $\lambda = 0$ | 0.2319 | 0.2318 | 1,348.17 |
| FFNN, $\lambda = 0.001$ | 0.2583 | 0.2584 | 447.07 |
| DNN, $\lambda = 0.001$ | 0.2514 | 0.2512 | 1,378.33 |

is slightly different than the sharp phase transition on infinite lattice (Landau (1976)).

A 3-layer DNN with 100 neurons on each hidden layer was trained to predict θ from X . For the purpose of comparison, the semi-automatic method with components of raw vector X as candidate summary statistics was used. We also tested an FFNN with a single hidden layer of 100 neurons and considered the regularization technique (2.2) with $\lambda = 0.001$. The FFNN used is totally different from that used by Blum and François (2010). See details in Section 1.3.

Summary statistics learned by different methods led to different ABC posteriors. They were compared to those ABC posteriors resulting from the ideal summary statistic S^* .

3.3. Results

As shown in Table 1, DNN learns a better prediction rule than the semi-automatic method and FFNN, although it takes more training time. The regularization technique does not improve the performance, probably because overfitting is not a significant issue given that the training data ($N = 10^6$) outnumbers the $\approx 3 \times 10^4$ parameters.

Figure 5a displays a scatterplot which compares the DNN-based summary statistic S and the sufficient statistic S^* . Points in the scatterplot represent to $(S^*(x), S(x))$ for an instance x in the testing set. A large number of the instances are concentrated at $S^* = 192, 200$, which appear as points in the top-right corner of the scatterplot. These instances are relatively uninteresting, so we display a heatmap of $(S(x), S^*(x))$ excluding them in Figure 5b. It shows that the DNN-based summary statistic $S(X)$ approximates an increasing function of $S^*(X)$.

The semi-automatic method constructs a summary statistic that fails to approximate $\mathbb{E}_\pi[\theta|X]$ (an increasing function of $S^*(X)$) but centers around the




Figure 5. DNN-based summary statistic S v.s. sufficient statistic S^* on the test dataset. (a) Scatterplot of 10^5 test instances. Each point represents to $(S^*(x), S(x))$ for a single test instance x . (b) Heatmap excluding instances with $S^*(x) = 192, 200$.




Figure 6. Summary statistic S constructed by the semi-automatic method v.s. sufficient statistic S^* on the test dataset. (a) Scatterplot of 10^5 test instances. Each point represents to $(S^*(x), S(x))$ for a single test instance x . (b) Heatmap excluding instances with $S^*(x) = 192, 200$.

prior mean $\theta_c = 0.4406$ (Figure 6). This is not surprising since the semi-automatic construction, a linear combination of X_j , is unable to capture the non-linearity of $\mathbb{E}_\pi[\theta|X]$.

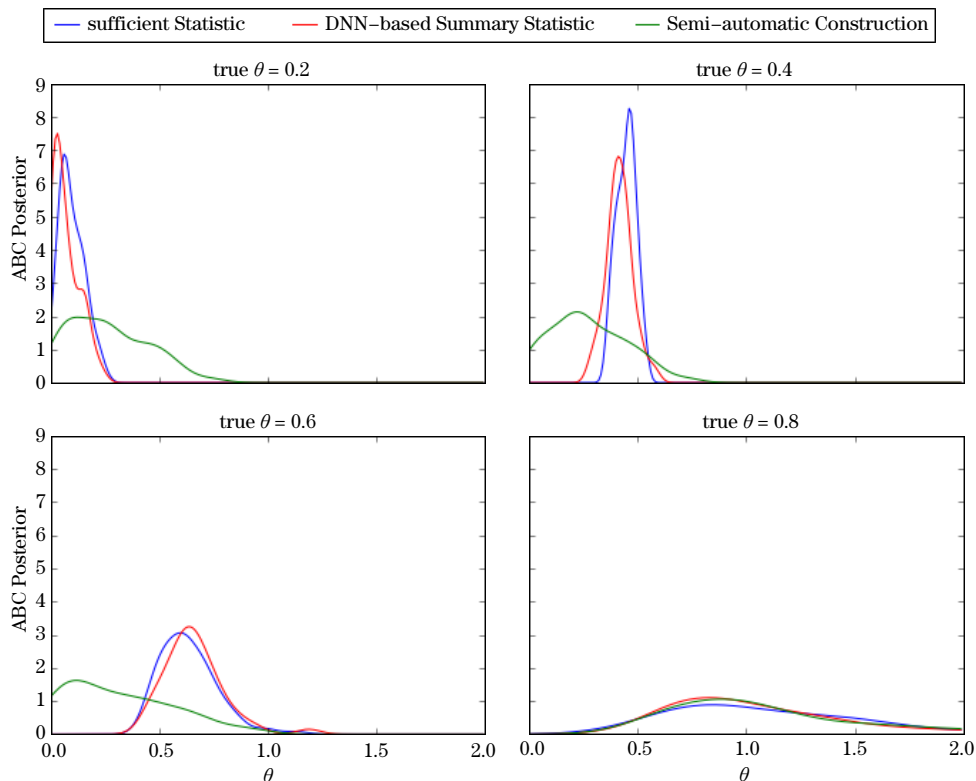


Figure 7. ABC posterior distributions for x_{obs} generated with true $\theta = 0.2, 0.4, 0.6, 0.8$.

ABC posterior distributions were obtained with the sufficient statistic S^* and the summary statistics S constructed by DNN and the semi-automatic method. For the sufficient statistic S^* , we set the tolerance level $\epsilon = 0$ so that the ABC posterior sample follows the exact posterior $\pi(\theta|X = x_{obs})$. For each summary statistic S , we set the tolerance threshold ϵ small enough so that 0.1% of 10^6 proposed θ 's were accepted. We repeated the comparison for four different observed data x_{obs} , generated from $\theta = 0.2, 0.4, 0.6, 0.8$, respectively; in each case, we compared the posterior obtained from S^* with the posteriors obtained from S , in Figure 7.

We highlight the case with true $\theta = 0.8$ (lower-right subplot in Figure 7). Since with high probability the spins X_i have the same sign when θ is large, it becomes difficult to distinguish different values of θ above the critical point θ_c based on the data x_{obs} . Hence we should expect the posterior to be small below θ_c and have a similar shape to the prior distribution above θ_c . All three ABC posteriors demonstrate this property.

4. Example: Moving Average of Order 2

4.1. ABC and summary statistics

The moving-average model is widely used in time series analysis. With X_1, \dots, X_p the observations, the moving-average model of order q , denoted by $MA(q)$, is given by

$$X_j = Z_j + \theta_1 Z_{j-1} + \theta_2 Z_{j-2} + \dots + \theta_q Z_{j-q}, \quad j = 1, \dots, p,$$

where Z_j are unobserved white noise error terms. We took $Z_j \stackrel{i.i.d.}{\sim} N(0, 1)$ in order to enable exact calculation of the posterior distribution $\pi(\theta|x_{obs})$, and then evaluation of the ABC posterior distribution. If the Z_j 's are non-Gaussian, the exact posterior $\pi(\theta|x_{obs})$ is computationally intractable, but ABC is still applicable.

Approximate Bayesian Computation has been applied to study the posterior distribution of the MA(2) model using the auto-covariance as the summary statistic (Marin et al. (2012)). The auto-covariance is a natural choice for the summary statistic in the MA(2) model because it converges to a one-to-one function of underlying parameter $\theta = (\theta_1, \theta_2)$ in probability as $p \rightarrow \infty$ by the Weak Law of Large Numbers,

$$AC_1 = \frac{1}{p-1} \sum_{j=1}^{p-1} X_j X_{j+1} \rightarrow \mathbb{E}(X_1 X_2) = \theta_1 + \theta_1 \theta_2,$$

$$AC_2 = \frac{1}{p-2} \sum_{j=1}^{p-2} X_j X_{j+2} \rightarrow \mathbb{E}(X_1 X_3) = \theta_2.$$

4.2. Experimental design

The MA(2) model is identifiable over the triangular region

$$\theta_1 \in [-2, 2], \quad \theta_2 \in [-1, 1], \quad \theta_2 \pm \theta_1 \geq -1,$$

so we took a uniform prior π over this region, and generated the training, validation, testing sets of size 10^6 , 10^5 , 10^5 , respectively. Each instance was a time series of length $p = 100$.

A 3-layer DNN with 100 neurons on each hidden layer was trained to predict θ from X . For purposes of comparison, we constructed the semi-automatic summary statistic by fitting linear regression of θ on candidate summary statistics - polynomial bases X_j, X_j^2, X_j^3, X_j^4 . We also test an FFNN with a single hidden layer of 100 neurons and considered the regularization technique (2.2)

Table 2. The root-mean-square error (RMSE) and training time of the semi-automatic, FFNN and DNN methods to predict (θ_1, θ_2) given X . λ is the penalty coefficient in the regularized objective function (2.2). Stochastic gradient descent fits each FFNN or DNN by 200 full passes (epochs) through the training set.

| Method | Training RMSE | | Testing RMSE | | Time (s) |
|-------------------------|---------------|---------------|---------------|---------------|----------|
| | θ_1 | θ_2 | θ_1 | θ_2 | |
| Semi-automatic | 0.8150 | 0.3867 | 0.8174 | 0.3857 | 45.63 |
| FFNN, $\lambda = 0$ | 0.1857 | 0.2091 | 0.1884 | 0.2115 | 543.42 |
| DNN, $\lambda = 0$ | 0.1272 | 0.1355 | 0.1293 | 0.1378 | 1,402.02 |
| FFNN, $\lambda = 0.001$ | 0.2642 | 0.2522 | 0.2679 | 0.2546 | 432.27 |
| DNN, $\lambda = 0.001$ | 0.1958 | 0.1939 | 0.1980 | 0.1956 | 1,282.66 |

with $\lambda = 0.001$. The FFNN used here is different than that used by Blum and François (2010). See details in Section 1.3.

Next, we generated some true parameters θ from the prior, drew the observed data x_{obs} , and numerically computed the exact posterior $\pi(\theta|x_{obs})$. Then we computed ABC posteriors using the auto-covariance statistic (AC_1, AC_2) , the DNN-based summary statistics (S_1, S_2) , and the semi-automatic summary statistic. The resulting ABC posteriors are compared to the exact posterior and evaluated in terms of the accuracies of the posterior mean of θ , the posterior marginal variances of θ_1, θ_2 , and the posterior correlation between (θ_1, θ_2) .

4.3. Results

Again DNN learns a better prediction rule than the semi-automatic method and FFNN, but takes more training time (Table 2, Figures 8 and 9). The regularization technique does not improve the performance.

We ran ABC procedures for an observed datum x_{obs} generated by true parameter $\theta = (0.6, 0.2)$, with three different choices of summary statistic: the DNN-based summary statistic, the auto-covariance, and also the semi-automatic summary statistic. The tolerance threshold ϵ was set to accept 0.1% of 10^5 proposed θ' in ABC procedures. Figure 10 compares the ABC posterior draws to the exact posterior which is numerically computed.

The DNN-based summary statistic gives a more accurate ABC posterior than either the ABC posterior obtained by the auto-covariance statistic or the semi-automatic construction. One of the important features of the DNN-based summary statistic is that its ABC posterior correctly captures the correlation between θ_1 and θ_2 , while the auto-covariance statistic and the semi-automatic statistic appear to be insensitive to this information (Table 3).




Figure 8. DNN predicting θ_1, θ_2 on the test dataset of 10^5 instances.




Figure 9. DNN predicting θ_1, θ_2 on the test dataset of 10^5 instances.

Table 3. Mean and covariance of exact/ABC posterior distributions for observed data x_{obs} generated with $\theta = (0.6, 0.2)$ in Figure 10.

| Posterior | mean(θ_1) | mean(θ_2) | std(θ_1) | std(θ_2) | cor(θ_1, θ_2) |
|-----------------|--------------------|--------------------|-------------------|-------------------|-----------------------------|
| Exact | 0.6418 | 0.2399 | 0.1046 | 0.1100 | 0.6995 |
| ABC (DNN) | 0.6230 | 0.2300 | 0.1210 | 0.1410 | 0.4776 |
| ABC (auto-cov) | 0.7033 | 0.1402 | 0.1218 | 0.2111 | 0.2606 |
| ABC (semi-auto) | 0.0442 | 0.1159 | 0.5160 | 0.4616 | -0.0645 |

We repeated the comparison for 100 different x_{obs} . As Table 4 shows, the ABC procedure with the DNN-based statistic better approximates the posterior




Figure 10. ABC posterior draws (top: DNN-based summary statistics, middle: auto-covariance, bottom: semi-automatic construction) for observed data x_{obs} generated with $\theta = (0.6, 0.2)$, compared to the exact posterior distribution contours.

moments than those using the auto-covariance statistic and the semi-automatic construction.

5. Discussion

We address how to automatically construct low-dimensional and informative summary statistics for ABC methods, with minimal need of expert knowledge. We base our approach on the desirable properties of the posterior mean as a

Table 4. Mean squared error (MSE) between mean and covariance of exact/ABC posterior distributions for 100 different x_{obs} .

| Posterior | MSE for | | | | |
|-----------------|--------------------|--------------------|-------------------|-------------------|-----------------------------|
| | mean(θ_1) | mean(θ_2) | std(θ_1) | std(θ_2) | cor(θ_1, θ_2) |
| ABC (DNN) | 0.0096 | 0.0089 | 0.0025 | 0.0026 | 0.0517 |
| ABC (auto-cov) | 0.0111 | 0.0184 | 0.0041 | 0.0065 | 0.1886 |
| ABC (semi-auto) | 0.5405 | 0.1440 | 0.4794 | 0.0891 | 0.3116 |

summary statistic for ABC, though it is generally intractable. We take advantage of the representational power of DNNs to construct an approximation of the posterior mean as a summary statistic.

We only heuristically justify our choice of DNNs to construct the approximation but obtain promising empirical results. The Ising model has a univariate sufficient statistic that is the ideal summary statistic and results in the best achievable ABC posterior. It is a challenging task to construct a summary statistic akin to it due to its high non-linearity and high-dimensionality, but we see in our experiments that the DNN-based summary statistic approximates an increasing function of the sufficient statistic. In the moving-average model of order 2, the DNN-based summary statistic outperforms the semi-automatic construction. The DNN-based summary statistic, which is automatically constructed, outperforms the auto-covariances; the auto-covariances in the MA(2) model can be transformed to yield a consistent estimate of the parameters, and have been widely used in the literature.

A DNN is prone to overfitting given limited training data, but this is not an issue when constructing summary statistics for ABC. In the setting of Approximate Bayesian Computation, arbitrarily many training samples can be generated by repeatedly sampling $(\theta^{(i)}, X^{(i)})$ from the prior π and the model \mathcal{M} . In our experiments, the size of the training data (10^6) is much larger than the number of parameters in the neural networks (10^4), and there is little discrepancy between the prediction error losses on the training data and the testing data. The regularization technique does not significantly improve the performance.

We compared the DNN with three hidden layers with the FFNN with a single hidden layer. Our experimental comparison indicates that FFNNs are less effective than DNNs for the task of summary statistics construction.

Supplementary Materials

The supplementary materials contain an extension of Theorem 1 and show

the convergence of the posterior expectation of $b(\theta)$ under the posterior obtained by ABC using $S_b(X) = \mathbb{E}_\pi[b(\theta)|X]$ as the summary statistic. This extension establishes a global approximation to the posterior distribution. Implementation details of backpropagation and stochastic gradient descent algorithms when training deep neural network are provided. The derivatives of squared error loss function with respect to network parameters are computed. They are used by stochastic gradient descent algorithms to train deep neural networks.

Acknowledgment

The authors gratefully acknowledge the National Science Foundation grants DMS1407557 and DMS1330132.

References

- Asmussen, S. and Glynn, P. W. (2007). *Stochastic simulation: Algorithms and analysis* (Vol. 57). Springer Science & Business Media.
- Beaumont, M. A., Zhang, W. and Balding, D. J. (2002). Approximate Bayesian computation in population genetics. *Genetics* **162**, 2025-2035.
- Beaumont, M. A. (2010). Approximate Bayesian computation in evolution and ecology. *Annual Review of Ecology, Evolution, and Systematics* **41**, 379-406.
- Bengio, Y., Courville, A. and Vincent, P. (2013). Representation learning: A review and new perspectives. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **35**, 1798-1828.
- Blum, M. G., Nunes, M. A., Prangle, D. and Sisson, S. A. (2013). A comparative review of dimension reduction methods in approximate Bayesian computation. *Statistical Science* **28**, 189-208.
- Blum, M. G. and François, O. (2010). Non-linear regression models for Approximate Bayesian Computation. *Statistics and Computing* **20**, 63-73.
- Caruana, R., Lawrence, S. and Giles, C.L. (2000). Overfitting in neural nets: backpropagation, conjugate gradient, and early stopping. In *Advances in Neural Information Processing Systems 13: Proceedings of the 2000 Conference* **13**, 402-408.
- Csilléry, K., Blum, M. G., Gaggiotti, O. E. and François, O. (2010). Approximate Bayesian computation (ABC) in practice. *Trends in Ecology & Evolution* **25**, 410-418.
- Faragó, A. and Lugosi, G. (1993). Strong universal consistency of neural network classifiers. *Information Theory, IEEE Transactions on* **39**, 1146-1151.
- Fearnhead, P. and Prangle, D. (2012). Constructing summary statistics for approximate Bayesian computation: semi-automatic approximate Bayesian computation. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **74**, 419-474.
- Fu, Y. X. and Li, W. H. (1997). Estimating the age of the common ancestor of a sample of DNA sequences. *Molecular Biology and Evolution* **14**, 195-199.
- Hinton, G. E. and Salakhutdinov, R. R. (2006). Reducing the dimensionality of data with neural networks. *Science* **313**, 504-507.

- Hinton, G. E., Osindero, S. and Teh, Y. W. (2006). A fast learning algorithm for deep belief nets. *Neural Computation* **18**, 1527-1554.
- Joyce, P. and Marjoram, P. (2008). Approximately sufficient statistics and Bayesian computation. *Statistical Applications in Genetics and Molecular Biology* **7**.
- Kolmogorov, A. N. (1942) Definition of center of dispersion and measure of accuracy from a finite number of observations. *Izv. Akad. Nauk S.S.S.R. Ser. Mat.* **6**, 3-32 (in Russian).
- Landau, D. P. (1976). Finite-size behavior of the Ising square lattice. *Physical Review B* **13**, 2997.
- LeCun, Y., Bottou, L., Bengio, Y. and Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE* **86**, 2278-2324.
- Lehmann, E. L. and Casella, G. (1998). *Theory of Point Estimation* (Vol. 31). Springer Science & Business Media.
- Le Roux, N. and Bengio, Y. (2010). Deep belief networks are compact universal approximators. *Neural Computation* **22**, 2192-2207.
- Lopes, J. S. and Beaumont, M. A. (2010). ABC: a useful Bayesian tool for the analysis of population data. *Infection, Genetics and Evolution* **10**, 825-832.
- Marin, J. M., Pudlo, P., Robert, C. P. and Ryder, R. J. (2012). Approximate Bayesian computational methods. *Statistics and Computing* **22**, 1167-1180.
- Marjoram, P., Molitor, J., Plagnol, V. and Tavaré, S. (2003). Markov chain Monte Carlo without likelihoods. *Proceedings of the National Academy of Sciences* **100**, 15324-15328.
- Ng, A. Y. (2004). Feature selection, L1 vs. L2 regularization, and rotational invariance. In *Proceedings of the Twenty-First International Conference on Machine Learning* (p. 78). ACM.
- Nowlan, S. J. and Hinton, G. E. (1992). Simplifying neural networks by soft weight-sharing. *Neural Computation* **4**, 473-493.
- Nunes, M. A. and Balding, D. J. (2010). On optimal selection of summary statistics for approximate Bayesian computation. *Statistical Applications in Genetics and Molecular Biology* **9**.
- Onsager, L. (1944). Crystal statistics. I. A two-dimensional model with an order-disorder transition. *Physical Review* **65**, 117.
- Pritchard, J. K., Seielstad, M. T., Perez-Lezaun, A. and Feldman, M. W. (1999). Population growth of human Y chromosomes: a study of Y chromosome microsatellites. *Molecular Biology and Evolution* **16**, 1791-1798.
- Schmidhuber, J. (2015). Deep learning in neural networks: An overview. *Neural Networks* **61**, 85-117.
- Sisson, S. A., Fan, Y. and Tanaka, M. M. (2007). Sequential monte carlo without likelihoods. *Proceedings of the National Academy of Sciences* **104**, 1760-1765.
- Srivastava, N., Hinton, G. E., Krizhevsky, A., Sutskever, I. and Salakhutdinov, R. (2014). Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research* **15**, 1929-1958.
- Sunnåker, M., Busetto, A. G., Numminen, E., Corander, J., Foll, M. and Dessimoz, C. (2013). Approximate Bayesian computation. *PLoS Comput. Biol.* **9**, e1002803.
- Sutskever, I. and Hinton, G. E. (2008). Deep, narrow sigmoid belief networks are universal approximators. *Neural Computation* **20**, 2629-2636.

- Tavaré, S., Balding, D. J., Griffiths, R. C. and Donnelly, P. (1997). Inferring coalescence times from DNA sequence data. *Genetics* **145**, 505-518.
- Toni, T., Welch, D., Strelkowa, N., Ipsen, A. and Stumpf, M. P. (2009). Approximate Bayesian computation scheme for parameter inference and model selection in dynamical systems. *Journal of the Royal Society Interface* **6**, 187-202.
- Wegmann, D., Leuenberger, C. and Excoffier, L. (2009). Efficient approximate Bayesian computation coupled with Markov chain Monte Carlo without likelihood. *Genetics* **182**, 1207-1218.
- Weiss, G. and von Haeseler, A. (1998). Inference of population history using a likelihood approach. *Genetics* **149**, 1539-1546.

Department of Statistics, Stanford University, Stanford, California 94305 USA

E-mail: baijiang@stanford.edu

Institute of Computational and Mathematical Engineering, Stanford, Stanford, California 94305 USA

E-mail: tungyuwu@stanford.edu

Department of Statistics, Stanford University, Stanford, California 94305 USA

E-mail: snarles@stanford.edu

Department of Statistics, Stanford University, Stanford, California 94305 USA

E-mail: whwong@stanford.edu

(Received October 2015; accepted September 2016)