Statistica Sinica Preprint No: SS-2020-0348	
Title	Deep Learning for Marginal Bayesian Posterior Inference with Recurrent Neural Networks
Manuscript ID	SS-2020-0348
URL	http://www.stat.sinica.edu.tw/statistica/
DOI	10.5705/ss.202020.0348
Complete List of Authors	Thayer Fisher, Alex Luedtke, Marco Carone and Noah Simon
Corresponding Author	Thayer Fisher
E-mail	thayerf@uw.edu
Notice: Accepted version subject to English editing.	

Statistica Sinica

Marginal Bayesian Posterior Inference using Recurrent Neural Networks with Application to Sequential Models

Thayer Fisher, Alex Luedtke, Marco Carone, and Noah Simon

University of Washington, Department of Biostatistics

Abstract

In Bayesian data analysis, it is often important to evaluate quantiles of the posterior distribution of a parameter of interest (e.g., to form posterior intervals). In multi-dimensional problems, when non-conjugate priors are used, this is often difficult generally requiring either an analytic or sampling-based approximation, such as Markov chain Monte-Carlo (MCMC), Approximate Bayesian computation (ABC) or variational inference. We discuss a general approach that reframes this as a multi-task learning problem and uses recurrent deep neural networks (RNNs) to approximately evaluate posterior quantiles. As RNNs carry information along a sequence, this application is particularly useful in time-series. An advantage of this risk-minimization approach is that

we do not need to sample from the posterior or calculate the likelihood.

We illustrate the proposed approach in several examples.

Key words and phrases: Bayesian deep learning, machine learning, quantile estimation.

1. Introduction

We consider the common Bayesian scenario wherein we have a dataset, X, generated from some parametrically specified distribution $(X\,|\theta,\eta) \sim P_{\theta,\eta}$. Here, we assume that $\theta \in \mathbb{R}$ is a one-dimensional parameter of interest. For simplicity, we assume that the nuisance parameter $\eta \in \mathbb{R}^p$ is finite-dimensional, although this is not essential – in the Discussion, we provide more details on handling an infinite-dimensional η. We further assume that (θ,η) follows a known prior distribution g, i.e., $(\theta,\eta) \sim g$. In analyzing our dataset X, we generally would like to evaluate the posterior distribution for θ:

$$dP\left(\theta \,|\, X\right) = \frac{\int_\eta dP_{\theta,\eta}(X) g(\theta,\eta) d\eta}{\int_{\theta,\eta} dP_{\theta,\eta}(X) g(\theta,\eta) d\theta d\eta} \ .$$

The integral in the denominator is often analytically intractable. When the likelihood function, $dP_{\theta,\eta}(X)$, can be efficiently calculated (up to a normalizing constant) one can draw samples from the posterior $dP\left(\theta \,|\, X\right)$ using stochastic simulation techniques such as Markov Chain Monte Carlo (MCMC) and rejection sampling (Asmussen and Glynn, 2007). From these samples, any posterior

summary can be computed. However, there are a number of situations in which the likelihood is difficult to calculate yet it is relatively straightforward to sample (η, θ) from g, and X from $P_{\theta,\eta}$. A class of methods, known as Approximate Bayesian Computation (ABC) methods, have been developed for inference in this regime (e.g., Marjoram et al., 2003). These rely on simulating data from the joint prior-likelihood distribution and computing the empirical distribution of parameters drawn alongside data that are "close" to the observed data in some sense. As ABC and more classical stochastic simulation techniques both often use some form of rejection sampling, they can suffer in problems that are complex or high dimensional. Furthermore, defining "close" as it pertains to sampled data X, is nontrivial in many cases. For even a moderately large X, informative summary statistics must be constructed to reduce rejection rate, otherwise the problem becomes intractable. If summary statistics are not suitable for the particular setting, the approximation of the posterior will be poor (Prangle, 2015).

In this work, we frame the calculation of posterior quantiles as an optimization problem. We consider the posterior quantile functions as risk minimizers with respect to particular losses. We approximate the function that minimizes this risk by restricting our optimization to the class of functions that can be represented by recurrent neural networks (RNNs) (Rumelhart et al., 1985). Since

we expect the posterior quantile to "update" as more data becomes available, the recurrent structure is advantageous. Moreover, the recurrent framework allows us to evaluate posterior quantiles for datasets of arbitrary length. Provided our network architecture is adequate for the specific problem, our function class will be sufficiently rich, and the risk minimizer over this class will be close to the true posterior quantile function. We discuss how stochastic subgradient optimization can be used to find local optima over this class. In particular, by simulating parameters/data from our prior/likelihood, our neural network can update its parameter estimates to "learn" how to evaluate posterior quantiles. We refer to this as a "statistical meta-learning procedure", as it uses simulated data to learn a Bayesian-optimal statistical rule.

Another obvious advantage of RNNs is their natural compatibility with time-series data, where the contribution of each observation to a parameter of interest can depend on previous or later observations. RNNs share information across recurrent steps, making them naturally suited to handling this type of problem. Furthermore, many Bayesian time-series are difficult due to the intractability of the likelihood (see Ghahramani (2001) for some examples). Avoiding the approximation of the entire posterior distribution is favorable in these scenarios when our target is credible intervals, for example.

Other work has proposed the use of deep learning in similar contexts. The

idea of using RNNs as meta-learners was first introduced by Hochreiter et al. (2001). Wong et al. (2018) use deep neural networks to approximate posterior summaries – in particular they consider approximating the posterior mean of a functional in high-dimensional problems. Creel (2017) build on this work, and apply this approach to particular econometric models, using feedforward neural networks to estimate posterior means. Chan et al. (2018) use exchangeable neural networks to model population genetic data. They train on the fly, simulating new data each time they update the weights of their neural network. However, they approximate the entire posterior using a parametric approach.

2. Reframing as an Optimization Problem

For any $t \in (0,1)$, let $Q^t(X)$ denote the t^{th} quantile of the posterior of θ under prior g and likelihood $P_{\theta,\eta}$. Our goal is to learn the quantile functions $\{\Phi(t) := Q^t(\cdot)$ for all $t \in \mathcal{T}\}$, where $\mathcal{T} \subseteq (0,1)$ either contains finitely many elements or is equal to $(0,1)$.

The quantity $Q^t(X)$ is the solution in q to

$$P\left(\theta \leq q \mid X\right) = t \ ,$$

where we assume for simplicity that this posterior is continuous.

Our approach hinges on the fact that Q^t can be written as the solution to

an optimization problem. In particular,

$$Q^t(X) = \text{argmin}_Q \, \mathrm{E}\left[\rho_t\left(\theta - Q(X)\right)\right], \tag{2.1}$$

where $\rho_t(u) = u\left(t - I\left\{u < 0\right\}\right)$ is the asymmetric L_1 norm (or "pinball loss", as referred to by Koenker and Hallock, 2001), and the minimization is taken over all functions that are measurable with respect to $\sigma(X)$ (the sigma-algebra generated by the data). Note that in the case where $t = 0.5$, this loss is precisely a scaled L_1 loss.

Directly solving (2.1) is intractable in general over the space of all measurable functions. Instead we restrict the problem to a rich finite-dimensional subclass: those functions that can be represented by deep recurrent neural networks. This is known to be a very rich class (Cybenko, 1989; Bach, 2017). In this case, our optimization problem consists of finding

$$\begin{aligned} \beta^* &= \text{argmin}_\beta \, \mathrm{E}\left[\rho_t\left(\theta - Q_\beta(X)\right)\right] \\ &= \text{argmin}_\beta \int_{\theta,\eta,\mathcal{X}} \rho_t\left(\theta - Q_\beta(X)\right) dP_{\theta,\eta}(X) g(\theta,\eta) d\theta d\eta \ , \end{aligned} \tag{2.2}$$

where Q_β is a deep recurrent neural network that takes as input each of the n observations in X, and outputs a single value; and β is the vector of weight and bias parameters in the neural network. For this network to be an effective estimator, we may need a large number of parameters: the network needs to learn both how to combine inputs to get the output, as well as what values to store in memory across recurrent steps.

2.1 Relation to Quantile Regression

The asymmetric L_1 norm described above is most commonly used in quantile regression. However, unlike quantile regression, which is generally engaged with in the frequentist context, where the goal is to estimate a quantile or distribution of a conditional outcome given a set of features, the outcomes and covariates aren't fixed. In fact, the only similarity between our proposed method and quantile regression is the use of pinball loss. We are attempting to learn one or many posterior quantiles as functions of data simulated from the prior and likelihood, $Q^t : \mathcal{X} \to \mathbb{R}$, as a minimizer of (2.1) over the joint distribution of (X, θ), $\int_\eta dP_{\theta,\eta}(X)g(\theta,\eta)d\eta$. We do this without ever engaging with a single, fixed, observed dataset, but rather through engaging with many simulated datasets. Reframing the derivation of the posterior quantile in this way is the key that allows us to leverage deep learning for posterior quantile approximation.

2.2 Motivation for Recurrent Structure

2.2.1 Time Series and Sequential Models

RNNs are advantageous for posterior inference in sequential models because they pass information sequentially by design. For illustrative purposes, consider

the second-order moving-average model described in Section 3.5:

$$X_j \sim Z_j + \theta_1 Z_{j-1} + \theta_2 Z_{j-1}.$$

The observed sequence X depends on the parameters of interest, $\{\theta_1, \theta_2\}$, only through an unobserved, latent sequence Z.

For most choices of likelihood over Z and choices of prior over θ, the likelihood $P(X \mid \theta)$ is not tractable. Furthermore, the contribution of X_j to the posterior distribution depends on the values of X_{j-1}, X_{j-2}. Because we cannot sample from the posterior directly for non-trivial choices of likelihood and prior, and because the data are sequential, an RNN is a good choice for posterior estimation. Later, in Section 3.5, we show that our proposed method outperforms several others in credible interval estimation.

2.2.2 Canonical Exponential Family

To motivate the choice of a recurrent neural network, as opposed to a simpler multilayer perceptron, we consider a d-parameter canonical exponential family indexed by θ with a conjugate prior on θ, namely

$$p(x|\theta) = h(x)\exp(\theta^\top T(x) - g(\theta))$$

$$p_{\eta_0,\gamma_0}(\theta) = k(\eta_0,\gamma_0)\exp(\eta_0(\gamma_0^\top\theta - g(\theta)))$$

for $\theta \in \mathbb{R}^d$. It follows that, if $X = \{x_1, \cdots, x_n\}$ where $X_1, \cdots, X_n \overset{\text{iid}}{\sim} p(\cdot|\theta)$,

$$p(\theta|X) = p_{\eta', \gamma(X)}(\theta)$$

with $\eta' = n+\eta_0$, and $\gamma(X) = \frac{\sum_{i=1}^n T(x_i)}{n+\eta_0} + \frac{\eta_0}{n+\eta_0}\gamma_0$. By the factorization theorem, in this simple setting, there exists a $(d+1)$-parameter sufficient statistic for $\theta|X$, namely $[\sum_{i=1}^n T(X_i), n]$. Therefore, we can write Q^t recurrently as

$$Q^t(X) = h\left[x_n, f(x_{n-1}, f(\ldots f(x_1, \beta)\ldots))\right]$$

for some fixed $\beta \in \mathbb{R}^{d+1}$, $h : \mathbb{R}^{d+2} \to \mathbb{R}$, and some $f : \mathbb{R}^{d+2} \to \mathbb{R}^{d+1}$. Here $f(x_i, \sum_{j=1}^{i-1} T(x_j), i-1) = \{\sum_{j=1}^{i} T(x_j), i\}$. We can think of f as the "memory" function that tracks the sufficient statistic across recurrent steps, while $h(\sum_{i=1}^n T(X_i), n) = F^{-1}(t)$, where F^{-1} is the quantile function of $p_{\eta', \gamma'}$, so h calculates $Q^t(X)$ from the sufficient statistic. Note that increasing d is equivalent to increasing the size of the memory across recurrent steps.

While it is not true that we can write $Q^t(X)$ in this way for a general likelihood and prior, we hope that for some sufficiently large memory size there exists an h, f such that

$$Q^t(X) \approx h\left[x_n, f(x_{n-1}, f(\ldots f(x_1, \beta)\ldots))\right].$$

In fact, note that if the size of the memory exceeds the number of observations, the set of order statistics is sufficient.

An advantage of the recurrent structure is that it allows us to get an approximation of a posterior quantile for a dataset of arbitrary length. In the case where the data are independent and identically distributed, this is particularly useful, as we are learning how to "update" the posterior as the number of observations grows. Using a feedforward architecture, one could only obtain posterior quantile approximations for datasets of a fixed size, one which is compatible with the input layer of the network.

This proposed recurrent neural network can take in a dataset and return an approximation to a single posterior quantile. It is also possible to design a network which learns more than a single quantile, and potentially all quantiles. These extensions are described in what follows.

2.3 Discretized Multi-task Learning

We first extend our procedure to simultaneously approximate $\{Q^t\}$ for a discrete collection of t with a single network. Let $\mathcal{T}$ denote that discrete set of quantiles and write $m_{\mathcal{T}} = |\mathcal{T}|$. We use a multi-task learning network architecture (Ruder, 2017) with shared hidden nodes and internal weights, and a set of $m_{\mathcal{T}}$ task-specific output nodes leading out of the shared, final hidden layer for approximating the $m_{\mathcal{T}}$ quantiles. Let β_{in} denote the shared internal weights; and $\beta_{out}(t)$ denote the task-specific output weights for estimating Q^t. In this

case, to find the optimal weights, β^*_{in} and $\beta^*_{out}(t)$, $t \in \mathcal{T}$, we must find

$$\text{argmin}_{\beta_{in},\beta_{out}(t):t\in\mathcal{T}} \sum_{t\in\mathcal{T}} w_t \, \mathrm{E}\left[\rho_t\left(\theta - Q_{\beta_{in},\beta_{out}(t)}(X)\right)\right] \tag{2.3}$$

with $w_t > 0$ denoting predefined costs — for example, one might take $w_t = 1$ for all $t \in \mathcal{T}$. This *hard sharing* multi-task learning framework that shares internal nodes across tasks seems sensible, as intermediate representations of the data that are useful for posterior calculations are likely not quantile-specific. However, one might consider modifying this to allow a "final stage" with several hidden layers specific to each quantile.

2.4 Continuous Quantile Learning

We further extend this to the construction of a single network that approximates all quantiles of the posterior simultaneously. To begin, we consider the optimization problem

$$\mathcal{Q} = \text{argmin}_Q \int_0^1 w(t) \, \mathrm{E}\left[\rho_t\left(\theta - Q(t, X)\right)\right] dt \ , \tag{2.4}$$

where w is any positive integrable function on $(0, 1)$. The solution to (2.4), $\mathcal{Q}$, is a function that takes as input both t and X, and returns the t^{th} posterior quantile, i.e., $\mathcal{Q}(t, X) = Q^t(X)$, as follows directly from (2.1). By approximating the problem in (2.4) using neural networks, we can build a single network to simultaneously approximate all quantiles.

Without loss of generality, we can assume that $\int_0^1 w(t)dt = 1$. Then, for random T drawn from density w, we can rewrite (2.4) as

$$\mathcal{Q} = \text{argmin}_Q \, \mathrm{E}\left[\rho_T\left(\theta - Q(T, X)\right)\right], \tag{2.5}$$

where the expectation is now over T, θ and X.

As noted above, we approximate the solution to (2.4) using a deep neural network by solving the optimization problem

$$\beta^* = \text{argmin}_\beta \, \mathrm{E}\left[\rho_T\left(\theta - Q_\beta(T, X)\right)\right], \tag{2.6}$$

where Q_β is a network that takes as input the dataset and the selected quantile, and returns a single output. The architecture of this network is similar to the single quantile problem described in the beginning of Section 2, except that T is included as an additional input at every recurrent step.

2.5 Optimization

Let ℓ_1, ℓ_M and ℓ_C denote the losses in criteria (2.2), (2.3), and (2.6), respectively. Optimization of all three can be conducted using standard stochastic subgradient-based methods (Nedic and Bertsekas, 2001). In particular, for the loss (2.2), we see that the sub-differential of our loss, $\nabla\ell_1\left(\beta\right)$, consists of

$$\nabla\ell_1\left(\beta\right) \ni -\mathrm{E}\left[\dot{\rho}_t\left(\theta - Q_\beta(X)\right)\nabla Q_\beta(X)\right],$$

where $\nabla Q_\beta(X)$ is a subgradient of our network; and $\dot{\rho}$, defined pointwise as $\dot{\rho}_t(u) = tI(u > 0) + (1-t)I(u < 0)$, is a subgradient of the quantile loss function. There are many subgradients based on how one defines $\dot{\rho}_t(0)$. Stochastic subgradients can be calculated by sampling a (θ, η) pair from g and a dataset X from $P_{\theta,\eta}$, and then plugging these into

$$\hat{\nabla}\ell_1(\beta) = \dot{\rho}_t\left(\theta - Q_\beta(X)\right)\nabla Q_\beta(X)\ . \tag{2.7}$$

By linearity, stochastic subgradients can be similarly calculated for the multi-task learning criterion, ℓ_M.

For continuous quantile learning, our expectation-based formulation, ℓ_C, allows us to easily calculate stochastic subgradients. In particular, given a sampled (θ, η) pair, a dataset X sampled from $P_{\theta,\eta}$, and a T sampled from w, we can calculate

$$\hat{\nabla}\ell_C(\beta) = \dot{\rho}_T\left(\theta - Q_\beta(X)\right)\nabla Q_\beta(X) \tag{2.8}$$

and note that $\mathrm{E}[\hat{\nabla}\ell_C(\beta)]$ is in the sub-differential of ℓ_C at β.

The full approximation procedure for 2.2 is described in Algorithm 1. We repeatedly sample (θ, X) pairs, and use those pairs to train our neural network. Note that steps 3 and 4 are generally done by back propagation, and that this algorithm can be extended using linearity for the multi-task learning criterion ℓ_M and having the network output one prediction for each quantile of interest.

For the optimization described by 2.6, we need only sample a T from w before we calculate the gradient in step 3. While simple stochastic gradient descent is described here, extension to more complex stochastic optimization techniques that use adaptive learning rates/ momentum (such as Kingma and Ba, 2014; Smith and Topin, 2017) can be and are actually employed in simulations in practice.

Algorithm 1 Learn $Q^t(\cdot)$

Require: $m =$ batch size, $k =$ learning rate

for $0 \leq i \leq$ total iterations **do**

1. Simulate $(\theta_1, \eta_1), \ldots (\theta_m, \eta_m) \overset{\text{iid}}{\sim} g(\theta, \eta)$
2. Simulate $X_1 \sim P_{\theta_1,\eta_1}, \ldots, X_m \sim P_{\theta_m,\eta_m}$
3. $\hat{\nabla}\ell(\beta) \leftarrow \frac{1}{m}\sum_{j=1}^{m} \dot{\rho}_t\left(\theta_j - Q_\beta(X_j)\right) \nabla Q_\beta(X_j)$
4. $\beta \leftarrow \beta - k\ \hat{\nabla}\ell(\beta)$
5. $i \leftarrow i + 1$

end for

3. Examples

In this section, we consider three examples. The first is simple, namely the estimation of the median of iid gaussian observations using a gaussian prior, the multi-task estimation of posterior deciles, as well as learning posterior quan-

tiles continuously. The posterior quantiles here have a simple closed form, so evaluating performance is very straightforward. For the second example, we approximate the posterior median for the maximum component in a mixture of finite mixtures (MFM). Finally, we approximate the posterior median for the basic reproduction number in a Bayesian Stochastic SIR model.

For all examples, optimization was performed using `Adam` (Kingma and Ba, 2014) as implemented in `TensorFlow` (Abadi et al., 2016), with subgradients calculated using mini-batches. Performance was evaluated in all simulation settings based on mean loss, or estimated risk, over a held out test set. Error bars at the conclusion of training represent a 95% confidence interval of the risk of the neural network based on the sample standard error of a held out test set. Where appropriate, we compare our neural networks to Stan, another general-purpose posterior estimation tool, in addition to comparing to a specialized, problem-specific tool.

3.1 Gaussian Example

We consider a simple example with observations $x_1, \ldots, x_n$ drawn iid from a $N(\theta, 1)$ distribution, with θ drawn from a $N(0, \sigma^2)$ prior. Here, we would like posterior quantiles for θ given our data. It is well known, in this case, that

$$\theta \,|\, x_1, \ldots, x_n \;\sim\; N\left(\frac{\bar{x}\sigma^2}{1/n + \sigma^2}, \left(n + \sigma^{-2}\right)^{-1}\right)$$

where $\bar{x} = n^{-1} \sum_i x_i$. Thus,

$$Q^t(X) = \frac{\bar{x}\sigma^2}{1/n + \sigma^2} + \frac{\Phi^{-1}(t)}{\sqrt{n + \sigma^{-2}}} \ .$$

Furthermore, letting $X_j = \{x_1, \ldots, x_j\}$, we have, for $j \geq 2$

$$Q^t(X_j) = \frac{(\frac{j-1}{j}\bar{x}_{j-1} + x_j)\sigma^2}{1/j + \sigma^2} + \frac{\Phi^{-1}(t)}{\sqrt{j + \sigma^{-2}}} \ .$$

So, in this simple case, there exists an exact recurrent update step that uses a two-dimensional sufficient statistic, $\{\bar{x}_j, j\}$. While this case is quite basic, it is illustrative, and because we have access to the exact value of $Q^t(X)$, it is trivial to evaluate the performance of our approximated posterior quantile.

We consider approximating the posterior median. In this specific case, we can assess the quality of our network by comparing the pinball loss of our network to the pinball loss of the true posterior median on a held-out test set.

We ran simulations using $n = 100$ observations, where θ has prior variance $\sigma^2 = 1/100$ (equal to the conditional variance of $\bar{x}$). We use mini-batch `Adam` for optimization, where each mini-batch contains 100 datasets, with a learning rate of 10^{-2}. Our network has 32 nodes per hidden layer and 4 hidden layers, with ReLU activation functions. We evaluate the risk on a held-out test set of 500 datasets. Performance is summarized in Figure 1.

We also ran simulations using an identical prior and likelihood, but this time attempting to simultaneously learn all 9 posterior deciles, with uniform weight

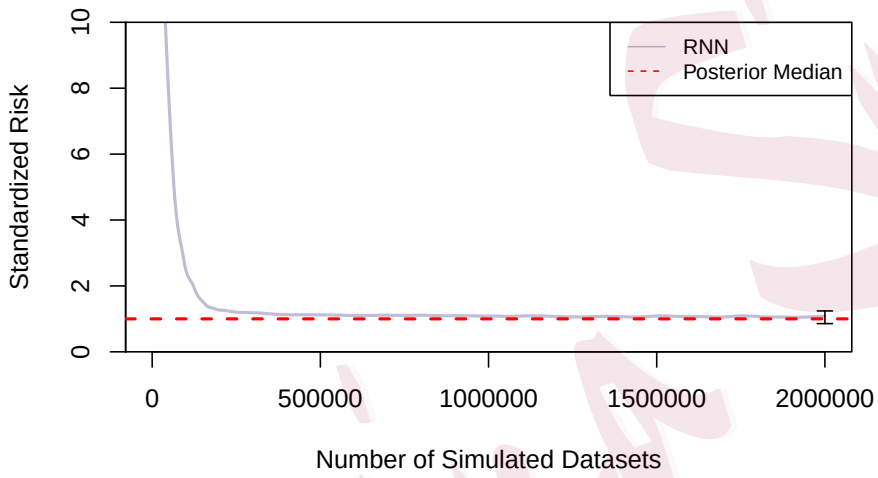

Figure 1: Risk curve for estimation of the posterior median in a Gaussian prior and Gaussian likelihood simulation scenario. The risk of the true minimizer, the posterior median, is standardized to be equal to one (red). Here, the posterior median has closed-form. Our estimator has about 2% excess risk when compared to the Bayes estimator.

on the loss over the 9 tasks as described in 2.3. Performance is summarized in Figure 2. We compare to the risk of the true posterior deciles, and are able to obtain a risk which is close to the true posterior deciles. We evaluate performance on a test set of 500 held-out datasets.

Finally, we ran simulations using the same prior and likelihood as above, but the target of estimation was a random quantile which was sampled uniformly on the unit interval, using the loss described in 2.6. We included the target quantile as an input to our neural network at each recurrent step. We evaluated performance on a test set of 500 held-out datasets and 500 held-out random quantiles. We compare the average loss over this test set to the average loss of the corresponding posterior quantiles, which have closed form. The results are summarized in Figure 3.

3.2 MFM Example

We consider a slightly more complex case in which observations $x_1, \dots, x_n$ are drawn iid from a $\frac{1}{k}\sum_{i=1}^{k} N(\theta_i, 0.01)$ distribution, with θ drawn from a $N(0, 0.25)$ prior, and k drawn from a Pois(4) prior shifted to have a minimum of 1. The choice of variance for the prior and the likelihood were to ensure some separation between the components.

Our target of estimation here is the posterior median of $\theta_{(k)}$, the maximum

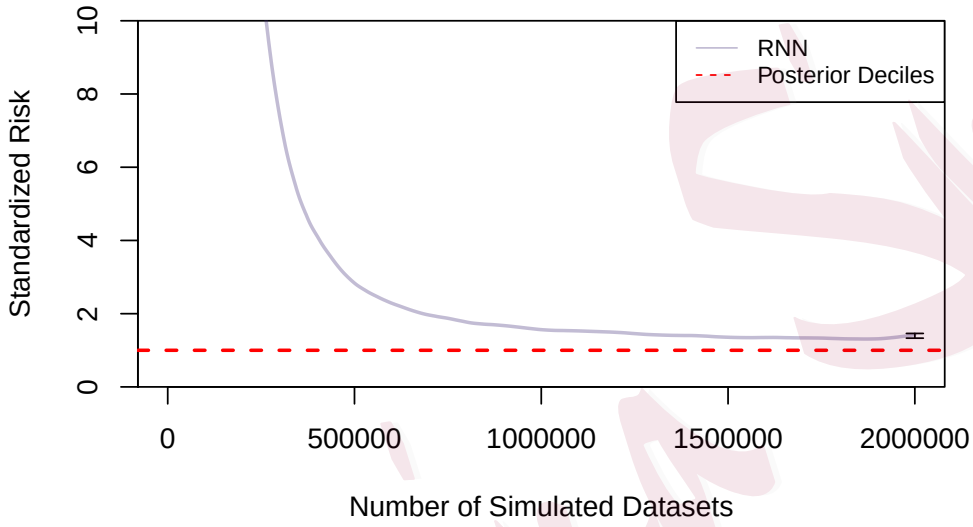

Figure 2: Risk curve for estimation of the posterior deciles in a Gaussian prior and Gaussian likelihood simulation scenario. The risk of the true minimizer, the posterior deciles, is in set to one (red). Here, the posterior deciles have closed-form. Our estimator has about 30% excess risk when compared to the Bayes estimator.

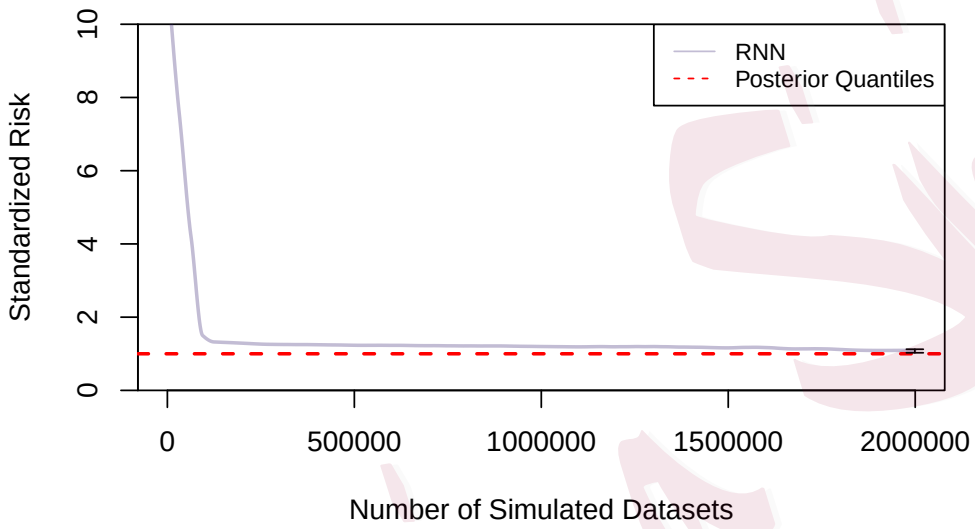

Figure 3: Risk curve for estimation of the posterior quantiles in a Gaussian prior and Gaussian likelihood simulation scenario. The quantile to estimate for each of the 1500 simulated datasets was sampled uniformly at random from the unit interval. The risk of the true minimizer, the posterior quantiles, is in set to one (red). Here, the posterior quantiles have closed-form. Our estimator has about 5.6% excess risk when compared to the Bayes estimator.

component. Note that this parameter is sensitive to the number of components, with the distribution of $\theta_{(k)}|X, k$ depending heavily on k, making procedures that try to determine the number of clusters, such as BIC, inaccurate.

We ran simulations using $n = 250$ observations. We used mini-batch `Adam` for optimization, where each mini-batch contains 150 datasets, with a learning rate of 10^{-4}. We used a network with 32 nodes per hidden layer, and 4 hidden layers.

The reason for the larger mini-batch in this simulation setting is the increased variability in the gradient due to the variable number of components. A larger mini-batch size makes the stochastic optimization more stable. Performance is summarized in Figure 4. We compared to a Julia package specialized for obtaining approximate posteriors for this model class (Miller and Harrison, 2018), as well as Stan (Carpenter et al., 2016), using two strategies to estimate the number of components. One strategy was to use BIC to estimate the number of clusters, and the other was to use the prior mean number of clusters. Stan then proceeded with the number of clusters fixed. As our parameter of interest is sensitive to the number of clusters, the Stan posterior estimation framework does poorly. Our neural network approaches the risk of the approximate posterior median given by a specialized package within about $15,000$ simulated datasets. We evaluated performance on a held-out test set

of 500 datasets.

3.3 Stochastic SIR Example

We also consider a stochastic SIR model. In this setting, we observe a disease epidemic, and would like to estimate the posterior median for the basic reproduction number, R_0. A stochastic SIR model is a continuous time stochastic process that models how a disease interacts with a population of size N. At each time t, there is a number of susceptible individuals, $S(t)$, a number of infected individuals, $I(t)$, and a number of recovered individuals, $R(t)$. The disease is modeled according to the following differential equations

$$\begin{aligned}\frac{\partial S(t)}{\partial t} &= -\frac{\beta S(t)I(t)}{N} + \sqrt{\frac{\beta S(t)I(t)}{N}}\omega_1(\partial t)\\ \frac{\partial I(t)}{\partial t} &= \frac{\beta S(t)I(t)}{N} - \gamma I(t) - \\ &\qquad \sqrt{\frac{\beta S(t)I(t)}{N}}\omega_1(\partial t) + \sqrt{\gamma I(t)}\omega_2(\partial t),\end{aligned}$$

where β is the infection rate, γ is the recovery rate, and ω_1, ω_2 are standard Wiener processes (Allen, 2017). The basic reproduction number, R_0, is the number of expected new infections an individual generates over the course of their disease in a fully susceptible population, β/γ. This parameter, and its posterior quantiles, are of interest in this problem (Clancy et al., 2008), since $R_0 > 1$ means the disease is likely to become an epidemic.

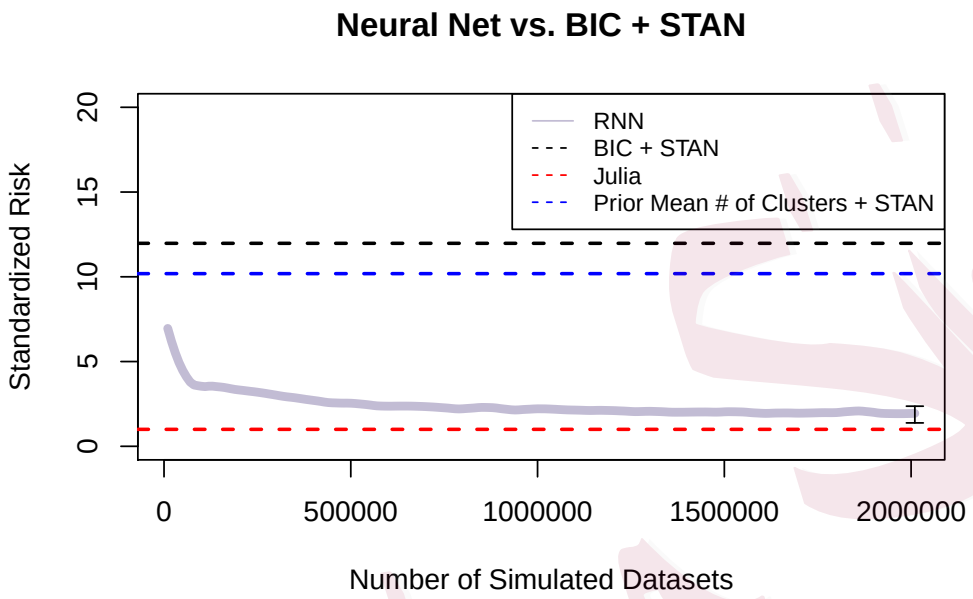

Figure 4: Risk curve for estimation of the posterior median in a Gaussian prior and Gaussian mixture likelihood simulation scenario, with a prior on the number of mixture components. We compare to Stan, with 2 strategies for selecting the number of components, as well as a specialized Julia package, BayesianMixtures. Our estimator has 87% excess risk when compared to BayesianMixtures, though it is much improved over our Stan implementations, and it is possible better architecture and more training would lead to better results.

We simulate from $t = 0$ to $t = 100$ using finite differences with $\partial t = 0.01$. However, we only observe integer values of t, for a total of 101 observations. In this setting, we observe both $S(t)$ and $I(t)$, though it is simple to reduce the observed data to only $R(t)$, for example. We specify the following priors for β and γ,

$$\beta \sim \Gamma(9, 0.05)$$

$$\gamma \sim \Gamma(3, 0.05).$$

We selected these priors based on Clancy et al. (2008), tuning the values so that sample paths expressed a variety of disease dynamics. We used mini-batch `Adam` for optimization, where each mini-batch contains 100 datasets, with a learning rate of 10^{-4}. We used a network with 32 nodes per hidden layer, and 4 hidden layers.

Performance is summarized in Figure 5. We compared to Stan, where we have attempted to solve this problem using finite differences. Our neural network approaches the risk of Stan within about $30,000$ simulated datasets, though this problem is slightly more difficult for us. We evaluated performance on a held-out test set of 2000 datasets.

To see where our method can improve over Stan in this setting, consider the same data-generating mechanism, but instead of observing $S(t)$ and $I(t)$ for all integer values of t, we observe only $t \in \{0, 20, 40, 60, 80, 100\}$. In the

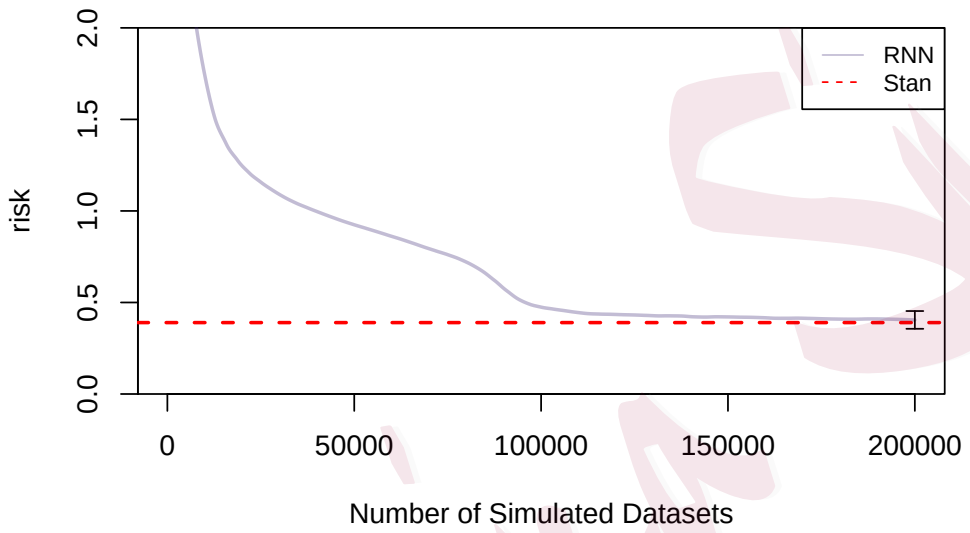

Figure 5: Risk curve for estimation of the posterior median of the basic reproduction rate in the stochastic SIR simulation setting. We compare to a solution produced in Stan, and both methods observe a grid of 101 timepoints. Our estimator achieves a risk which is about 3% greater than the Stan implementation.

case of our Stan implementation, we can only use finite differences over the 6 observations, since we do not marginalize out the drift and the stochastic component over a finer, unobserved grid. Since our observed grid is sparse, this leads to a very poor approximation of the true data generating mechanism. However, for our RNN, we can simulate over a grid of arbitrary granularity, but only train on the 6 observed time-points. The results of this can be seen in Figure 6. We see that, due to the misspecification of Stan, our method has lower risk after a small number of observed datasets. This is an example of how stochastic simulation techniques can fail when calculating the likelihood exactly is computationally intractable.

3.4 HMM Example

We also consider a Hidden Markov Model (HMM) example. In this setting, we observe a sequence of values, $y_1, \ldots, y_{100}$. These values are emissions from hidden states $x_1, \ldots, x_{100}$ which are not observed. There are three possible states at each x_i, $\{s_1, s_2, s_3\}$. We selected the following prior and likelihood for these simulations:

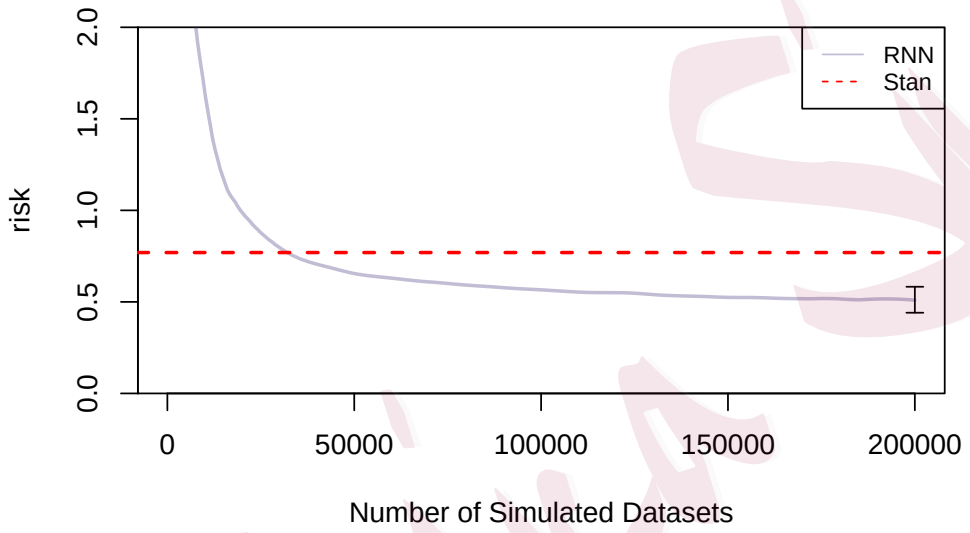

Figure 6: Risk curve for estimation of the posterior median of the basic reproduction rate in the stochastic SIR simulation setting. We compare to a solution produced in Stan, and both methods observe only 6 timepoints. Our estimator had about a 44% reduction in risk over the Stan implementation.

$$\theta \sim N(0, 1)$$

$$Z_1, Z_2, Z_3 \sim N(\theta, 1)$$

$$X_0 \sim U(\{s_1, s_2, s_3\})$$

$$y_i | X_i = s_j \sim N(Z_j, 1)$$

$$P(X_{i+1} = s_j | X_i = s_k) = \frac{e^{|z_j - z_k|}}{\sum_{j=1}^{3} e^{|z_j - z_k|}}$$

So, the hidden states represent unobserved means from which observed Gaussian random variables are drawn. The transition probability between states is proportional to the exponentiated ℓ_1-distance between those states. Our goal is to make inference on the posterior distribution of θ, the center of these unobserved means, $\theta | y_1, \ldots, y_{100}$. This model is a simplified version of the type of continuous-emission HMMs used in speech recognition (Ananthi and Dhanalakshmi, 2015).

We estimated 90% credible intervals on a held out test set of 2000 sequences. We compared to standard importance sampling ABC with Gaussian kernel weights. As Y is too high-dimensional to be computationally tractable, we instead used the deciles of the Y sequence as a summary statistic. We generated 1000 importance samples per observed sequence, and estimated the

0.05 and 0.95 posterior quantiles from this sample with a bandwidth of 1.

Results can be seen in Figure 7. In this simulation scenario, it is difficult to choose sensible summary statistics. The recurrent network had a marginal coverage of 88.95% and an average interval width of 2.11. ABC with summary statistics had a marginal coverage of 50.35% and an average interval width of 0.92. The prior quantiles, as a reminder, have an interval width of 3.29 and marginal coverage of 90%. So, ABC with these summary statistics is more conservative than the prior, while the RNN leverages the information to provide narrower intervals. The RNN is ideal in this scenario because of the difficulty of the problem, and the sequential nature of the observed data.

3.5 Order Two Moving Average Example

Finally, we consider another latent variable model, an order-two moving-average model. Suppose we observe a sequence X with

$$X_j \sim Z_j + \theta_1 Z_{j-1} + \theta_2 Z_{j-1}$$

where Z_j are latent Gaussian random variables. We would like to obtain a 90% credible interval on θ_2. Note that if the Z_j follow a non-Gaussian distribution, the likelihood $P(X \mid \theta)$ is intractable. However here, we can draw

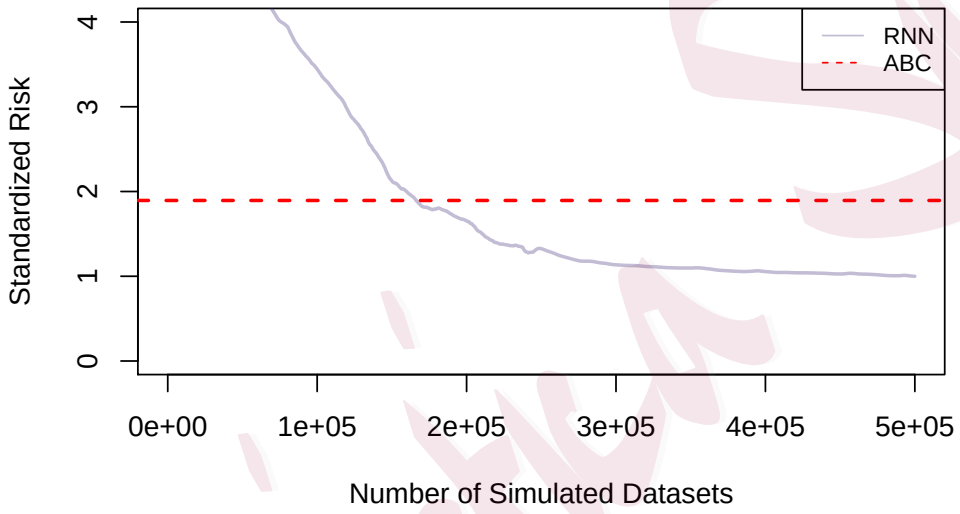

Figure 7: Standardized risk of the estimated 90% credible interval for the center of unobserved hidden state means.

comparisons to the exact posterior distribution.

The likelihood and prior are identical to (Wong et al., 2018). The Z_i are standard Gaussian, and θ_1, θ_2 are uniform over a specific triangular region such that they are identifiable.

We compare our method to the semi-automatic ABC procedure described in (Fearnhead and Prangle, 2011), transforming the input data into the vector of 1-gap products, $X_j X_{j+2}$. We chose this transformation because the sample mean of $X_j X_{j+2}$ is a consistent estimator of θ_2. Using a polynomial basis expansion of this transformed X, we then use linear regression to estimate the posterior mean, and use that estimate as a summary statistic. Fearnhead and Prangle call the method "semi-auto ABC", as although the procedure can be automatic, the choice of transformation in this problem is not. We also compare to a method proposed in (Wong et al., 2018). In that work, the authors use a Deep Neural Network to estimate the posterior mean. Then, they use that estimate as a summary statistic for ABC. In both settings, we use rejection sampling to obtain the approximate posterior. The rejection scheme, as well as the number of samples drawn, were calibrated to have approximately the same runtime as our RNN method. Methods were tested on a held-out set of 1000 sequences.

Loss results can be seen in 8. As both comparison methods use an estimate

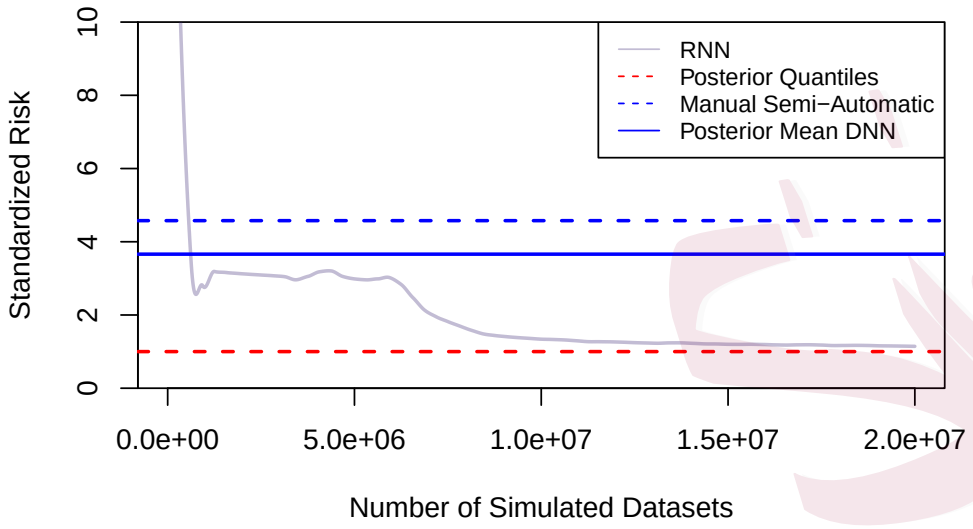

Figure 8: Standardized risk of the estimated 90% credible interval for the second order coefficient in a moving average model.

of the posterior mean as a summary statistic, it is not surprising that their performances are poor in the tails. Our method directly targets the 90% credible interval, and has approximately correct coverage while maintaining relatively narrow interval widths.

3.6 Summary of Coverage Results

For all of the simulation settings above, we estimated 90% posterior credible intervals. A comparison of our coverage results to various other estimators can be seen in Table 1. Our RNN estimators outperform or have comparable performance to STAN in every simulation setting. We maintain approximately correct marginal coverage, while having much narrower intervals than the prior.

4. Discussion

In this manuscript we have proposed a method for using deep recurrent neural nets to approximate posterior quantiles of a univariate parameter of interest from a possibly multivariate Bayesian problem. To fit this neural net, it is only necessary to sample from the prior and likelihood – posterior samples never need to be drawn, and, so long as we can sample from $P_{\theta,\eta}$, the likelihood itself never needs to be calculated. We propose three types of networks: a simple network for estimating a single, pre-specificed, conditional quantile; a multi-task network for estimating a finite set of pre-specified quantiles; and a slightly more complex network for estimating the entire conditional quantile function. We show, in increasingly complex settings, that a recurrent neural network can approximate single posterior quantiles approximately as accurately as specific

Gaussian			
Method	Coverage	Interval Width	Loss
RNN	0.906	0.241	0.0149
Exact Posterior	0.8958	0.233	0.0146
Prior Quantiles	0.903	0.323	0.0203
Mixture of Finite Mixtures			
Method	Coverage	Interval Width	Loss
RNN	0.942	0.339	0.019
Stan (BIC)	0.130	0.283	0.377
BayesianMixtures	0.921	0.147	0.009
Prior Quantiles	0.900	1.225	0.079
Stochastic SIR			
Method	Coverage	Interval Width	Loss
RNN	0.879	4.12	0.321
Stan	0.878	2.773	0.198
Prior Quantiles	0.900	10.17	0.935
Sparse Stochastic SIR			
Method	Coverage	Interval Width	Loss
RNN	0.9105	3.64	0.362
Stan	0.586	2.55	0.926
Prior Quantiles	0.900	10.17	0.935
Hidden Markov Model			
Method	Coverage	Interval Width	Loss
RNN	0.890	2.11	0.137
ABC	0.504	0.92	0.260
Prior Quantiles	0.900	3.29	0.203
Order Two Moving Average			
Method	Coverage	Interval Width	Loss
Exact Posterior	0.90	0.417	0.029
RNN	0.865	0.527	0.035
Semi-Auto ABC	0.440	0.476	0.135
Post-Mean DNN	0.466	0.400	0.108

Table 1: Coverage of 90% credible intervals using various posterior estimators in several simulation settings. Estimators exhibiting poor marginal coverage are highlighted in red.

state-of-the-art methods that attempt to sample from the posterior directly.

To simplify the exposition, in Section 1, we assumed that the nuisance parameter, η, belongs to a finite-dimensional space. In fact, we did not make use of this assumption in this manuscript – we only require that one can sample parameters from the prior, and subsequently sample data from the corresponding likelihood. Therefore, in nonparametric Bayes problems, where an infinite-dimensional parameter can be sampled from a prior and the data can be sampled from the likelihood, the proposed method can be immediately used to estimate a univariate summary of this infinite-dimensional parameter.

Compared to other generalized tools for posterior approximation, such as Stan, our method is advantageous in cases where the likelihood is not easily calculable. In the second stochastic SIR model, the results of which are in Figure 6, we are able to achieve a lower loss than a natural implementation in Stan. We reach a lower loss because our Stan implementation must necessarily misspecify the likelihood in order to get posterior estimates.

It is interesting to draw parallels to the ABC methodology and link the contrast to different forms of smoothing methods used in classical nonparametric problems. Both the ABC methodology and our meta-learning procedure aim to approximate a quantile using many draws from the prior-likelihood. In ABC, the function mapping from the observed data to a quantile of the posterior

distribution is approximated via a local smoother that only uses information from simulation replicates near the data that was originally observed. This is analogous to kernel smoothing methods in nonparametric regression settings, where a regression function estimate at a single predictor value x only relies on observations with predictors near x – for bounded kernels, a perturbation of distant observed predictor values has no impact on the regression fit at a distant point. Our neural network approach, on the other hand, uses all simulated data points to learn a rich neural network approximation to the function mapping from the observed data to a posterior quantile. This is analogous to series estimators from nonparametric statistics, where linear models are fit using a collection of basis functions applied to the observed predictors – here, a perturbation of distant observed predictor values can have a nontrivial impact on the regression fit at a given point.

Our proposed method is also quite similar to other Bayesian deep learning methodologies, such as the exchangeble network proposed in Chan et al. (2018). However, our networks exhibit good performance in both exchangeable and non-exchangeable cases. One critical advantage of our recurrent networks is that we are able to get predictions for sequences or datasets of arbitrary length. This is not possible with an exchangeble network. However, as noted in Vaswani et al. (2017), the memory in a recurrent network is a bottleneck in a number of

contexts. Therefore, as dependence between datapoints which are far apart in our sequences grows, the size of the memory in our recurrent nodes must grow as well. The fixed memory size of recurrent nodes could also be beneficial in an online learning setting, however. As new data are acquired, our RNN estimator could quickly calculate new posterior estimates without the need for retraining. Further investigation into the potential applications to online inference should be conducted.

As with other posterior sampling and approximation methodologies, we anticipate that our method will suffer from the curse of dimensionality. For MCMC methods, the curse of dimensionality typically occurs when the parameters are high-dimensional. In contrast, for our method, we do not expect that having a high-dimensional nuisance parameter η will negatively impact the performance of our method since we only estimate the posterior of a univariate parameter θ. However, our method is susceptible to the curse of dimensionality in the data structure, since the method is implicitly smoothing across possible data realizations. One approach to try to improve performance of our method in the presence of high-dimensional data would be to increase the richness of our network. Because we can simulate an unlimited number of prior-likelihood draws for our stochastic gradient steps, using a very rich network will not lead to overfitting, though the optimization scheme may require more draws to converge.

Therefore, when the data are high-dimensional, we expect that our method will benefit greatly from the dramatic recent progress in neural network software that enables users to leverage massive computational power to quickly optimize deep networks.

5. Appendix

5.1 Feedforward Comparison

As an illustrative example of how an RNN outperforms a Feedforward network in posterior quantile estimation, we performed simulations which make a direct comparison. In Figure 9, we show that a feedforward network with an identical internal structure (same number of hidden nodes per layer and number of layers) is outperformed by a recurrent counterpart in estimating a 90% credible interval for the maximum component of a mixture of finite mixtures. Furthermore, the majority of the cases where our method is best-suited involve series data where the likelihood is not easily calculable, making the RNN an even more preferable option.

5.2 Rate of Convergence

According to the Bernstein-von Mises theorem, any posterior quantile should concentrate around the true parameter at a rate of $O(1/\sqrt{n})$. For the Gaus-

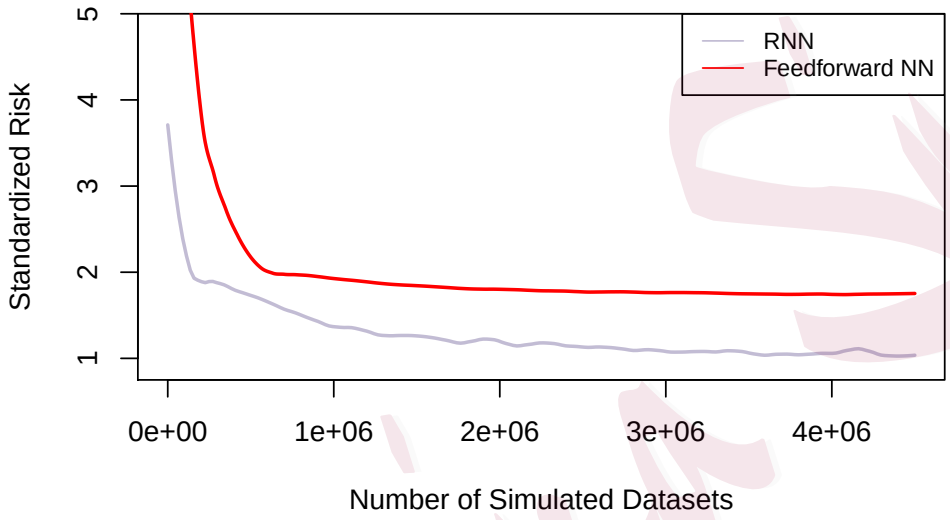

Figure 9: Standardized risk of a feedforward neural network vs. a recurrent neural network. Targets of estimation are the 0.05 and 0.95 posterior quantiles of the maximum component in a mixture of finite mixtures with a prior on the number of components. Both networks have 3 hidden layers and 10 nodes per hidden layer.

sian conjugate prior simulation setting, we trained several networks with varying sample sizes to predict the 0.05 quantile. Averaging the distance from these estimates to the true values over a large test set give us the approximate convergence rate. In Figure 10, we see a log-log plot of the 0.05 quantile distance from the true value vs. sample size. The slope of the imposed line of best fit is -0.48. Therefore, we see that the 0.05 quantile concentrates around the true parameter at approximately a rate of $O(1/\sqrt{n})$, though we admit this relationship is not perfectly linear. While ideally we would also show this convergence in more complicated simulation settings, doing so requires training a separate neural network for each value of n, which is computationally expensive for more complex examples. Note also that as n increases, the optimal architecture and training regime changes. However, all of the networks use to produce this convergence plot were identical in their architecture and training regimen. This is one reason why the convergence rate is not exact.

5.3 Distribution Function Analysis

In the continuous estimation setting, where the quantile to be estimated is an imput to the RNN, it is possible to obtain an estimate for the entire distribution function by inputting a grid of quantiles. In the Gaussian conjugate prior scenario, it is also possible to compare this estimated distribution function with

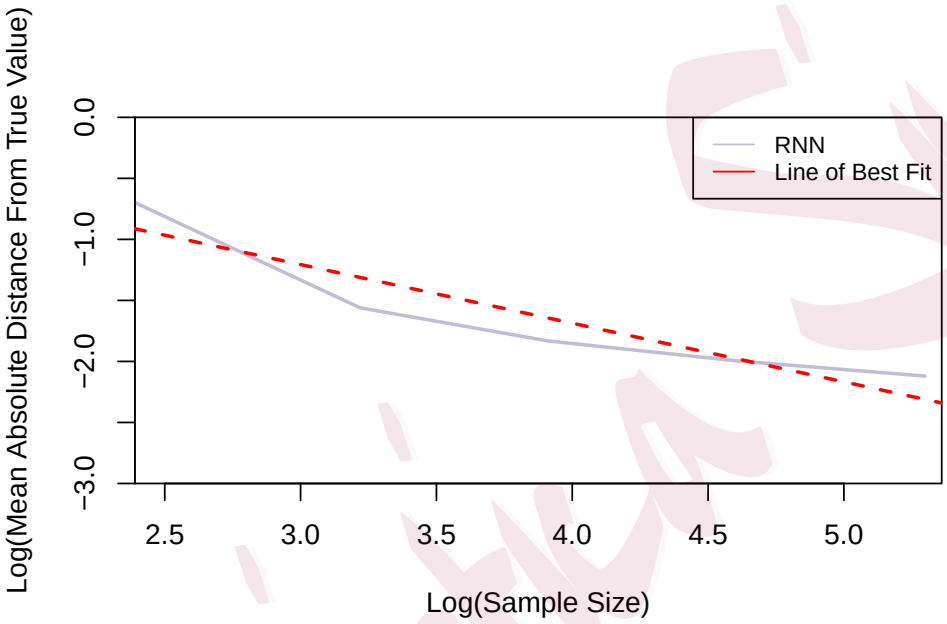

Figure 10: Log(Mean absolute difference) between estimated 0.05 quantile and true value. The line of best fit is superimposed.

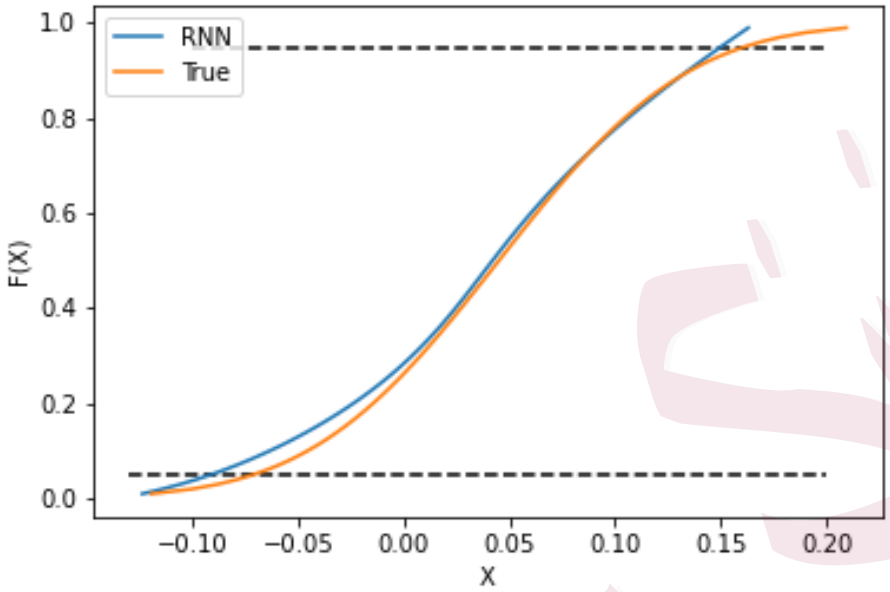

Figure 11: True posterior distribution functions vs. RNN estimated distribution function. Horizontal lines represent the 0.05 and 0.95 quantiles.

the true posterior distribution, since it has closed-form. One example of this can be seen in Figure 11. While the distribution function is easier to estimate than the density function, one could imagine binning predicted quantile values to achieve a density estimate.

5.4 Conditional Coverage Analysis

Risk is not an ideal measure of performance, but it is the best option in some cases. Other measures worth considering, such as marginal coverage and interval width, are reported in Table 1. However, conditional coverage is still of interest. In the Gaussian conjugate prior setting, where we have direct access to the posterior distribution, we can stratify by the sufficient statistic, the sample mean, to obtain a measure of conditional coverage, seen in Figure 12. We see that, apart from extreme values of the sufficient statistic, we obtain approximately correct conditional coverage as well.

5.5 Permutation Invariance

In the first two simulation settings presented, the posterior distribution of our parameter of interest is invariant to permutations of the input data. In the first of these examples, the Gaussian conjugate prior scenario, we analyzed the performance of our model by averaging predictions from permutations of the input data. Two examples of this can be seen in Figure 13.

Another potential solution to exchangeability is sorting the input data. When we tested coverage in Table 1, we sorted the input data for the Mixture of Finite Mixture example, which improved performance. In this specific set-

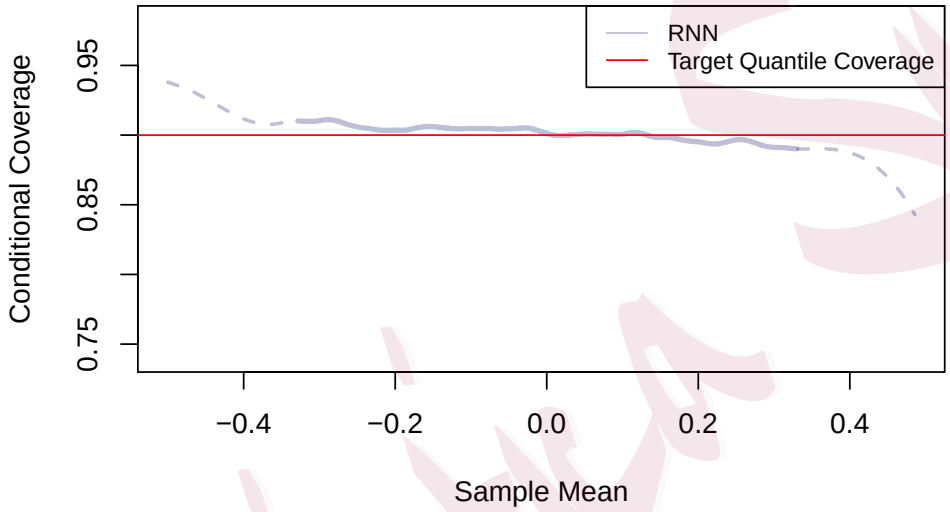

Figure 12: Conditional coverage of RNN posterior quantile estimates, stratified by sample mean. Dotted lines represent values in the lowest or highest 1% of our test set.

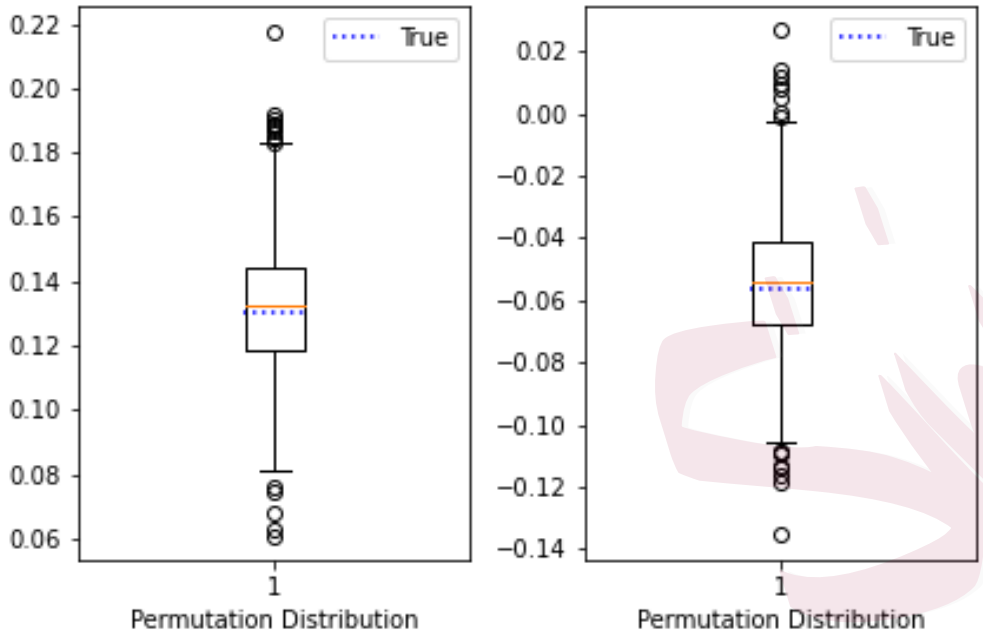

Figure 13: Boxplot of predicted values under permutation with superimposed true posterior value for two example datasets.

ting, because we are interested in the posterior distribution of the maximum component, sorting is an intuitive solution.

References

Abadi, M., P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, et al. (2016). Tensorflow: A system for large-scale machine learning. In *OSDI*, Volume 16, pp. 265–283.

Allen, L. J. (2017). A primer on stochastic epidemic models: Formulation, numerical simulation, and analysis. *Infectious Disease Modelling 2*(2), 128–142.

Ananthi, S. and P. Dhanalakshmi (2015). Svm and hmm modeling techniques for speech recognition using lpcc and mfcc features. pp. 519–526.

Asmussen, S. and P. W. Glynn (2007). *Stochastic simulation: algorithms and analysis*, Volume 57. Springer Science & Business Media.

Bach, F. (2017). Breaking the curse of dimensionality with convex neural networks. *Journal of Machine Learning Research 18*(19), 1–53.

Carpenter, B., A. Gelman, M. Hoffman, D. Lee, B. Goodrich, M. Betancourt, M. A. Brubaker, J. Guo, P. Li, A. Riddell, et al. (2016). Stan: A probabilistic programming language. *Journal of Statistical Software 20*(2), 1–37.

Chan, J., V. Perrone, J. P. Spence, P. A. Jenkins, S. Mathieson, and Y. S. Song (2018, Feb). A likelihood-free inference framework for population genetic data using exchangeable neural networks.

Clancy, D., P. D. O'Neill, et al. (2008). Bayesian estimation of the basic reproduction number in

stochastic epidemic models. *Bayesian Analysis 3*(4), 737–757.

Creel, M. (2017). Neural nets for indirect inference. *Econometrics and Statistics 2*, 36–49.

Cybenko, G. (1989). Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems 2*(4), 303–314.

Fearnhead, P. and D. Prangle (2011). Constructing summary statistics for approximate bayesian computation: Semi-automatic abc.

Ghahramani, Z. (2001). An introduction to hidden markov models and bayesian networks. In *Hidden Markov models: applications in computer vision*, pp. 9–41. World Scientific.

Hochreiter, S., A. S. Younger, and P. R. Conwell (2001). Learning to learn using gradient descent. In *International Conference on Artificial Neural Networks*, pp. 87–94. Springer.

Kingma, D. P. and J. Ba (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

Koenker, R. and K. F. Hallock (2001). Quantile regression. *Journal of economic perspectives 15*(4), 143–156.

Marjoram, P., J. Molitor, V. Plagnol, and S. Tavaré (2003). Markov chain monte carlo without likelihoods. *Proceedings of the National Academy of Sciences 100*(26), 15324–15328.

Miller, J. W. and M. T. Harrison (2018). Mixture models with a prior on the number of components. *Journal of the American Statistical Association 113*(521), 340–356.

Nedic, A. and D. P. Bertsekas (2001). Incremental subgradient methods for nondifferentiable opti-

mization. *SIAM Journal on Optimization 12*(1), 109–138.

Prangle, D. (2015). Summary statistics in approximate bayesian computation.

Ruder, S. (2017). An overview of multi-task learning in deep neural networks. *arXiv preprint arXiv:1706.05098*.

Rumelhart, D. E., G. E. Hinton, and R. J. Williams (1985). Learning internal representations by error propagation. Technical report, California Univ San Diego La Jolla Inst for Cognitive Science.

Smith, L. N. and N. Topin (2017). Super-convergence: Very fast training of neural networks using large learning rates.

Vaswani, A., N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin (2017). Attention is all you need.

Wong, W., B. Jiang, T.-y. Wu, and C. Zheng (2018). Learning summary statistic for approximate bayesian computation via deep neural network. *Statistica Sinica*.