
Thompson Sampling via Local Uncertainty

Zhendong Wang¹ Mingyuan Zhou¹

Abstract

Thompson sampling is an efficient algorithm for sequential decision making, which exploits the posterior uncertainty to address the exploration-exploitation dilemma. There has been significant recent interest in integrating Bayesian neural networks into Thompson sampling. Most of these methods rely on global variable uncertainty for exploration. In this paper, we propose a new probabilistic modeling framework for Thompson sampling, where local latent variable uncertainty is used to sample the mean reward. Variational inference is used to approximate the posterior of the local variable, and semi-implicit structure is further introduced to enhance its expressiveness. Our experimental results on eight contextual bandit benchmark datasets show that Thompson sampling guided by local uncertainty achieves state-of-the-art performance while having low computational complexity.

1. Introduction

There has been significant recent interest in employing deep neural networks to better solve sequential decision-making problems, such as these in reinforcement learning (Mnih et al., 2013; 2015; 2016; Arulkumaran et al., 2017; François-Lavet et al., 2018) and contextual bandits (Riquelme et al., 2018; Russo et al., 2018). In a typical setting, by sequentially interacting with the environment, the agent or algorithm needs to learn how to take a sequence of decisions in order to maximize the expected cumulative reward. These problems are frequently encountered in various practical applications, ranging from clinical trials to recommender systems to anomaly detection (Djallel & Irina, 2019). Using a deep neural network as a powerful function approximator, whose task is to learn the mapping from an observed contextual feature vector to the hidden reward distributions, has

¹McCombs School of Business, The University of Texas at Austin, Austin, TX 78712, USA. Correspondence to: Mingyuan Zhou <mingyuan.zhou@mcombs.utexas.edu>.

become a common practice. Since the model training and data collection usually happen at the same time, the model needs to not only accurately approximate the distribution of the observed data, but also gain enough flexibility to predict that of the future data.

Addressing the exploration-exploitation dilemma is a vital part of sequential decision making. To maximize the expected cumulative reward, the agent needs to balance its effort in exploration, which chooses actions that may potentially increase its understanding of the environment, and its effort in exploitation, which takes the action that is expected to be the best given existing information. Typically, under-exploration will possibly trap the agent at a bad local optimal solution, while over-exploration could lead to a significant exploration cost. Various strategies have been proposed to tackle the exploration-exploitation dilemma, such as ϵ -greedy (Sutton & Barto, 1998), upper-confidence bound (Auer, 2002), Boltzmann exploration (Cesa-Bianchi et al., 2017; Sutton & Barto, 1998), and Thompson sampling (Thompson, 1933). More recently, carefully adding random noise to model parameters (Plappert et al., 2017; Fortunato et al., 2017; Gal & Ghahramani, 2016) or bootstrap sampling (Osband et al., 2016) before decision making also provide effective ways to encourage exploration.

Thompson sampling (TS) (Thompson, 1933), an elegant and widely used exploration strategy, is known for both its simplicity and good practical performance (Chapelle & Li, 2011; Agrawal & Goyal, 2012; 2013; Russo et al., 2018). TS will keep updating the posteriors of the parameters of the hidden reward distributions, and take actions according to the posterior predictive distributions of the rewards. Relying on posterior uncertainty to do exploration is the promising point of TS. Unfortunately, the exact posteriors are tractable for only a few models with limited representation power. Therefore, significant effort has been dedicated to posterior approximation. A recent development along this direction is empowering TS with Bayesian neural networks (Hinton & Van Camp, 1993; Bishop, 2006; Graves, 2011; Neal, 2012; Hernández-Lobato & Adams, 2015), and relying on the posteriors of the neural network weights to perform exploration under the TS framework (Riquelme et al., 2018).

Various inference methods have been employed to capture the uncertainty of the neural network weights, including

Bayes by backprop (BBB) (Blundell et al., 2015), stochastic gradient Markov chain Monte Carlo (MCMC) (Welling & Teh, 2011; Li et al., 2016; Mandt et al., 2016), point estimate combined with random sampling (Riquelme et al., 2018), and interactive particles based approximation (Zhang et al., 2019). All these methods are focused on modeling the uncertainty of the global variables (*i.e.*, weights of the deep neural network) to maintain flexible approximation to the posterior predictive distributions.

In this paper, differing from all aforementioned methods, we propose to sample from a local latent variable distribution to model the uncertainty of the mean rewards of actions given the contextual input. Our framework uses a latent variable model to model the reward distribution given a contextual input, and encodes this input to approximate the posterior distribution of the local latent variable given both the context, which has already been observed, and reward, which is yet to be observed. To further improve the expressiveness of the latent distribution, we introduce a semi-implicit variational distribution structure into the framework. We test our framework on contextual bandits, a classical task in sequential decision making, to verify its effectiveness. Experimental results show that the proposed local uncertainty guided TS algorithms achieve state-of-the-art performance, while having low computational complexity.

2. TS via Global Uncertainty

Below we first briefly review contextual bandits and TS.

2.1. Contextual Bandits and TS

In a contextual bandit problem, we denote $\mathbf{x} \in \mathbb{R}^d$ as the d dimensional context (state) given by the environment, $a \in \mathcal{A} = \{1, \dots, C\}$ as the action in a finite discrete space of size C , and $r \in \mathbb{R}$ as the scalar reward. Commonly, the agent will interact with the environment sequentially for T times. At each time $t = 1, \dots, T$, the agent observes a new context \mathbf{x}_t , chooses action $a_t \in \mathcal{A}$ based on the information provided by \mathbf{x}_t , and receives reward r_t provided by the environment. The reward r_t can be a deterministic mapping $r_t = f(\mathbf{x}_t, a_t)$ or a more complicated stochastic mapping $r_t = f(\mathbf{x}_t, a_t, \epsilon_t)$, where ϵ_t represents random noise. The interactions (\mathbf{x}_t, a_t, r_t) at different times are independent from each other. The objective of the agent is to maximize the expected cumulative reward $\mathbb{E}[\sum_{t=1}^T r_t]$, or equivalently to minimize the expected cumulative regret as

$$\text{CR}(T) = \sum_{t=1}^T \mathbb{E} \left[\max_{a_t \in \mathcal{A}} \mathbb{E}[r(\mathbf{x}_t, a_t)] - r_t \right]. \quad (1)$$

TS is a widely-used classical algorithm that has been shown to be effective for bandit problems both in practice (Chapelle & Li, 2011) and theory (Agrawal & Goyal, 2012). Unlike

the ϵ -greedy algorithm that use parameter ϵ to control exploration and upper-confidence bound (UCB) that uses variance approximation to encourage exploration (Sutton & Barto, 1998), TS uses the uncertainty from the posterior samples of the model parameters for solving exploration-exploitation dilemma. If the model is not confident about its parameters, there will be large variations among the posterior samples, which will force the model to explore more to help better approximate the true underlying distribution.

For contextual bandits, vanilla TS maintains a posterior distribution $p_t(\boldsymbol{\theta})$ for global model parameter $\boldsymbol{\theta}$. For $t = 1, \dots, T$, it samples $\boldsymbol{\theta}$ from its current posterior $p_{t-1}(\boldsymbol{\theta})$ and uses the sampled $\boldsymbol{\theta}$ to transform the context-action pairs (\mathbf{x}_t, a) to estimate the mean rewards with $\hat{r}(a) = f(\mathbf{x}_t, a; \boldsymbol{\theta})$; after that, it greedily chooses the best action $a_t = \text{argmax}_{a \in \mathcal{A}} \hat{r}(a)$, receives reward r_t from the environment, and uses the observed data to update its posterior on $\boldsymbol{\theta}$ via Bayes' rule. Vanilla TS faces the difficulty of balancing the complexity of the mapping function f and tractability of posterior inference for $\boldsymbol{\theta}$, as discussed below.

2.2. Existing Global Uncertainty Guided TS

In this section, we describe representative TS based algorithms for contextual bandits, which all share the same strategy of relying on the uncertainty of the global parameters (*e.g.*, neural network weights) that are shared across all observations to perform exploration. They differ from each other on how complex the mapping function f is and how the posterior inference on $\boldsymbol{\theta}$ is implemented.

Linear Method: This method uses Bayesian linear regression with closed-form Gibbs sampling update equations, which relies on the posteriors of the regression coefficients for TS updates and maintains computational efficiency due to the use of conjugate priors. It assumes that at time t , the reward y_t of an action given contextual input \mathbf{x}_t is generated as $y_t = \mathbf{x}_t^T \boldsymbol{\beta} + \epsilon_t$, where $\boldsymbol{\beta}$ is the vector of regression coefficients and $\epsilon_t \sim \mathcal{N}(0, \sigma^2)$ is the noise. Note to avoid cluttered notation, here we omit the action index. This method places a normal prior on $\boldsymbol{\beta}$ and inverse gamma prior on σ^2 . At time t , given \mathbf{x}_t and the current random sample of $\boldsymbol{\beta}$, it takes the best action under TS and receives reward y_t ; with $\mathbf{x}_{1:t}$ and $y_{1:t}$, it samples σ^2 from its inverse gamma distributed conditional posterior, and then samples $\boldsymbol{\beta}$ from its Gaussian distributed conditional posterior; it proceeds to the next time and repeats the same update scheme under TS.

While this linear method accurately captures the posterior uncertainty of the global parameters $\boldsymbol{\beta}$ and σ^2 , its representation power is limited by both the linear mean and Gaussian distribution assumptions on reward y given context \mathbf{x} . In practice, the linear method often provides surprisingly competitive results, thanks to its ability to provide accurate uncertainty estimation. However, when its assumptions do

not hold well in practice, such as when there are complex nonlinear dependencies between the rewards and contextual vectors, the linear method, even though with accurate posterior estimation, may not be able to converge to a good local optimal solution. Following [Riquelme et al. \(2018\)](#), we refer to this linear method as “LinFullPost” in what follows.

Neural Linear: To enhance the representation power of LinFullPost while maintaining closed-form posterior sampling, [Riquelme et al. \(2018\)](#) propose the “Neural Linear” method, which feeds the representation of the last layer of a neural network as the covariates of a Bayesian linear regression model. It models the reward distribution of an action conditioning on x as $y \sim \mathcal{N}(\beta^T z_x, \sigma^2)$, where z_x is the output of the neural network given x as the input. It separates representation learning and uncertainty estimation into two parts. The neural network part is responsible for finding a good representation of x , while the Bayesian linear regression part is responsible for obtaining uncertainty estimation on β and making the decision on which action to choose under TS. The training for the two parts can be performed at different time-scales. It is reasonable to update the Bayesian linear regression part as soon as a new data arrives, while to update the neural network part only after collecting a sufficient number of new data points.

As Neural Linear transforms context x into latent space z via a deterministic neural network, the model uncertainty still all comes from sampling the global parameters β and σ from their posteriors under the Bayesian linear regression part. Hence, this method relies on the uncertainty of global model parameters to perform TS.

Bayes By Backprop (BBB): This method uses variational inference to perform uncertainty estimation on the neural network weights ([Blundell et al., 2015](#)). In order to exploit the reparameterization trick for tractable variational inference ([Kingma & Welling, 2013](#)), it models the neural network weights with independent Gaussian distributions, whose means and variances become the network parameters to be optimized. However, the fully factorized mean-field variational inference used by BBB is well-known to have the tendency to underestimate posterior uncertainty ([Jordan et al., 1999](#); [Blei et al., 2017](#)). Moreover, it is also questionable whether the weight uncertainty can be effectively translated into reward uncertainty given context x ([Bishop, 2006](#); [Sun et al., 2019](#)), especially considering that BBB makes both the independent and Gaussian assumptions on its network weights. For TS, underestimating uncertainty often leads to under exploration. As the neural network weights are shared across all observations, BBB also relies on the uncertainty of global parameters to perform TS.

Particle-Interactive TS via Discrete Gradient Flow (π -TS-DGF): The π -TS-DGF method of [Zhang et al. \(2019\)](#) casts posterior approximation as a distribution optimization

problem under a Wasserstein-gradient-flow framework. In this setting, posterior sampling in TS can be considered as a convex optimization problem on the space of probability measures. For tractability, it maintains a set of particles that interact with each other and evolve over time to approximate the posterior. For contextual bandits, each particle corresponds to a set of neural network weights, and the algorithm uniformly at random chooses one particle at each time and uses it as a posterior sample of the neural network weights. A benefit of π -TS-DGF is that it imposes no explicit parametric assumption on the posterior distribution. However, it faces an uneasy choice of setting the number of particles. Maintaining a large number of particles means training many sets of neural network weights at the same time, which is considerably expensive in computation, while a small number might lead to bad uncertainty estimation due to inaccurate posterior approximation. The computational cost prevents π -TS-DGF from using large-size neural networks. Similar to BBB, π -TS-DGF also relies on the uncertainty of global parameters to perform exploration.

3. TS via Local Uncertainty

Vanilla TS has several limitations. Its performance is sensitive to the accuracy of the mapping function $f(x, a; \theta)$ and maintaining the exact posteriors for all model parameters is often infeasible. Utilizing global parameter uncertainty to capture the posterior uncertainty of the mean rewards is challenging: first, the number of global parameters is often large, making it difficult to model their uncertainty under limited data without imposing strong assumptions; second, the model size is often constrained by the computational cost and training stability; third, the uncertainty on the global parameters may not be well translated into the uncertainty of the mean rewards by the mapping function.

To overcome these aforementioned limitations of vanilla TS, we propose TS via local uncertainty (LU). Rather than following the convention to impose uncertainty on global parameter θ to model the mean reward uncertainty, we apply uncertainty on local latent variables to balance exploration and exploitation under TS. We first construct a neural network powered latent variable model to model the mean reward distribution, and then introduce a contextual variational distribution to model the pre-posterior uncertainty on the mean rewards, which is used to guide the selection of actions. We first consider a contextual variational distribution using a diagonal Gaussian construction, and then another one using a semi-implicit construction.

3.1. Local Variable based Mean Reward Estimation

In a contextual bandit problem, the agent needs to continuously update its estimate of the unknown mean reward distributions through its interactions with the environment.

In this case, given context \mathbf{x} , taking different actions a will heavily impact the accumulated data for rewards, which predominantly influences the approximation of the mean reward distribution. TS via global uncertainty relies on the posterior sample of the global parameter to capture the uncertainty of the mean reward of an action at time t as

$$\mathbb{E}[r_t | \mathbf{x}_t, a_t, \beta], \quad \beta \sim p(\beta | \mathbf{x}_{1:t-1}, a_{1:t-1}, r_{1:t-1}). \quad (2)$$

It chooses the action whose mean reward given \mathbf{x}_t and β is the largest, receives reward from the environment, and then updates the posterior of the global parameter β before taking another contextual vector.

By contrast, denoting $\mathbf{r}_t \in \mathbb{R}^{|\mathcal{A}|}$ as the rewards of all actions in \mathcal{A} , we use a latent variable model to approximate the distribution of \mathbf{r}_t given \mathbf{x}_t as

$$\mathbf{r}_t \sim p(\mathbf{r}_t | \mathbf{x}_t, \mathbf{z}_t), \quad \mathbf{z}_t \sim p(\mathbf{z}).$$

This provides a flexible marginal distribution, whose density is often intractable, to model \mathbf{r}_t given \mathbf{x}_t as

$$p(\mathbf{r}_t | \mathbf{x}_t) = \mathbb{E}_{\mathbf{z}_t \sim p(\mathbf{z})}[p(\mathbf{r}_t | \mathbf{x}_t, \mathbf{z}_t)]. \quad (3)$$

To maximize the likelihood of this intractable marginal, we resort to variational inference (Jordan et al., 1999; Bishop & Tipping, 2000; Blei et al., 2017). More specifically, related to auto-encoding variational Bayes (Kingma & Welling, 2013), we introduce contextual variational distribution $q(\mathbf{z}_t | \mathbf{x}_t)$ to approximate the posterior $p(\mathbf{z}_t | \mathbf{x}_t, \mathbf{r}_t)$ by minimizing the Kullback–Leiber (KL) divergence as $\text{KL}(q(\mathbf{z}_t | \mathbf{x}_t) || p(\mathbf{z}_t | \mathbf{x}_t, \mathbf{r}_t))$. Since one may show that $\log p(\mathbf{r}_t | \mathbf{x}_t) = \mathcal{L}_t + \text{KL}(q(\mathbf{z}_t | \mathbf{x}_t) || p(\mathbf{z}_t | \mathbf{x}_t, \mathbf{r}_t))$ and the KL divergence is non-negative, where the evidence lower bound (ELBO) \mathcal{L}_t is expressed as

$$\mathcal{L}_t = \mathbb{E}_{\mathbf{z}_t \sim q(\cdot | \mathbf{x}_t)} \left[\log p(\mathbf{r}_t | \mathbf{x}_t, \mathbf{z}_t) + \log \frac{p(\mathbf{z}_t)}{q(\mathbf{z}_t | \mathbf{x}_t)} \right], \quad (4)$$

minimizing $\text{KL}(q(\mathbf{z}_t | \mathbf{x}_t) || p(\mathbf{z}_t | \mathbf{x}_t, \mathbf{r}_t))$ becomes the same as maximizing the ELBO \mathcal{L}_t . We note if $p(\mathbf{r}_t | \mathbf{x}_t, \mathbf{z}_t)$ in (4) is simplified as $p(\mathbf{r}_t | \mathbf{z}_t)$, then it becomes related to the variational hetero-encoder of Zhang et al. (2020).

Note for TS via LU, distinct from the usual auto-encoding variational inference (Kingma & Welling, 2013), we are facing an online learning problem, where we need to draw \mathbf{z}_t from variational posterior $q(\mathbf{z}_t | \mathbf{x}_t)$ to estimate the actions mean rewards before we are able to choose an action and hence observe its reward. Moreover, we only observe the chosen actions reward but not the other actions. Therefore, the contextual variational distribution $q(\mathbf{z}_t | \mathbf{x}_t)$ is not amortized over the actions’ rewards at time t . At time t , before optimizing the ELBO, we need to first sample from the mean reward distribution using

$$\mathbb{E}[r_t | \mathbf{x}_t, a_t, \mathbf{z}_t], \quad \mathbf{z}_t \sim q(\cdot | \mathbf{x}_t), \quad (5)$$

where $\mathbb{E}[r_t | \mathbf{x}_t, a_t, \mathbf{z}_t] = \int r_t \cdot p(r_t | \mathbf{x}_t, a_t, \mathbf{z}_t) dr_t$; we then choose the action whose mean reward given \mathbf{z}_t and \mathbf{x}_t is the largest, observe the true reward r_t returned by the environment, and optimize the parameters of p and q to maximize the ELBO in (4).

Note a key difference between TS via LU and TS via global uncertainty is that to estimate the mean rewards of all actions, a random sample from the contextual variational distribution of the local variable \mathbf{z}_t , as shown in (5), has replaced the role of a random sample from the posterior distribution of the global variable β , as shown in (2). In other words, rather than approximating the posterior of global parameters, our model estimates the posterior of local latent variable \mathbf{z}_t , and utilizes its uncertainty to perform exploration under TS. As global parameter β often has a very high dimension (e.g., when a neural network is used in $p(\mathbf{r}_t | \mathbf{x}_t, \beta)$), one often has to impose strong assumptions (e.g., independent Gauss with bounded variance) on its variational posterior for stable inference. By contrast, \mathbf{z}_t often has low dimension (e.g., 50), which can be well modeled with flexible variational distribution. We use a neural network to define the deterministic mapping from \mathbf{x}_t and \mathbf{z}_t to the mean rewards of all actions, and train the network parameter according to the ELBO in (4). We describe two different versions of TS via LU, as will be discussed in detail, in Algorithms 2 and 3 in the Appendix, respectively.

In summary, the change from relying on global uncertainty to replying on local uncertainty brings several potential benefits. First, the neural network mapping \mathbf{x}_t and \mathbf{z}_t to the mean rewards can be made as complex as needed, without the need to worry about the tractability of posterior inference on the global parameters, the number of which is often so large that uncertainty estimation on them becomes possible only under strong distributional assumptions. Second, the uncertainty comes from the input feature space rather than from the weight space, leading to more direct influence on the uncertainty of the mean rewards. Third, \mathbf{z}_t often has a much lower dimension than β , making it much more computationally efficient when the dimension of β is high.

3.2. Local Uncertainty Modeling with Gaussian Variational Posterior

We model both the rewards conditioning on \mathbf{x}_t and \mathbf{z}_t and the prior using diagonal Gaussian distributions as

$$\begin{aligned} p(\mathbf{r}_t | \mathbf{x}_t, \mathbf{z}_t) &= \mathcal{N}(\boldsymbol{\mu}_{\mathbf{r}_t}, \boldsymbol{\Sigma}_{\mathbf{r}}), \quad \boldsymbol{\mu}_{\mathbf{r}_t} = \mathcal{T}_{\theta}([\mathbf{x}_t, \mathbf{z}_t]), \\ p(\mathbf{z}_t) &= \mathcal{N}(0, \boldsymbol{\Sigma}_{\mathbf{z}}), \end{aligned} \quad (6)$$

where \mathcal{T}_{θ} is a neural networks parameterized by θ that maps the $[\mathbf{x}_t, \mathbf{z}_t]$ concatenation to the mean rewards of all actions as $\boldsymbol{\mu}_{\mathbf{r}_t} \in \mathbb{R}^{|\mathcal{A}|}$; both $\boldsymbol{\Sigma}_{\mathbf{r}} \in \mathbb{R}^{|\mathcal{A}| \times |\mathcal{A}|}$ and $\boldsymbol{\Sigma}_{\mathbf{z}} \in \mathbb{R}^{|\mathbf{z}| \times |\mathbf{z}|}$ are diagonal covariance matrix. Under this construction, the

estimated mean rewards of all actions can be expressed as

$$\mathbb{E}[r_t | \mathbf{x}_t, \mathbf{z}_t] = \mathcal{T}_\theta([\mathbf{x}_t, \mathbf{z}_t]). \quad (7)$$

We define the contextual variational distribution as a diagonal Gaussian distribution as

$$q(\mathbf{z}_t | \mathbf{x}_t) = \mathcal{N}(\boldsymbol{\mu}_{\mathbf{z}_t}, \boldsymbol{\Sigma}_{\mathbf{z}_t}), \\ \boldsymbol{\mu}_{\mathbf{z}_t} = \mathcal{T}_{\phi_1}(\mathbf{h}), \quad \boldsymbol{\Sigma}_{\mathbf{z}_t} = \mathcal{T}_{\phi_2}(\mathbf{h}), \quad \mathbf{h} = \mathcal{T}_{\phi_0}(\mathbf{x}_t), \quad (8)$$

where ϕ_0 , ϕ_1 , and ϕ_2 are neural network parameters. We describe how to address contextual bandit problems in Algorithm 2 in the Appendix, a limitation of which is that $q(\mathbf{z}_t | \mathbf{x}_t)$ is restricted to be a Gaussian distribution with a diagonal covariance matrix, which might not be flexible enough to well model the true posterior that may exhibit multi-modality, skewness, heavy tails, and dependencies between different dimensions. To improve its expressiveness, below we leverage semi-implicit variational inference (SIVI) of Yin & Zhou (2018) that mixes a distribution, which is simple to sample from but not required to be explicit, with an explicit and reparameterizable distribution to make the resulted hierarchical distribution more flexible, while maintaining tractable inference.

3.3. Local Uncertainty Modeling with Semi-Implicit Variational Posterior

Keeping likelihood $p(r_t | \mathbf{x}_t, \mathbf{z}_t)$ and prior $p(\mathbf{z}_t)$ the same as in (6), we model the contextual variational distribution using a semi-implicit construction as

$$q(\mathbf{z}_t | \mathbf{x}_t) = \int q(\mathbf{z}_t | \boldsymbol{\psi}_t, \mathbf{x}_t) q(\boldsymbol{\psi}_t | \mathbf{x}_t) d\boldsymbol{\psi}_t \quad (9)$$

where the first-layer explicit distribution is defined as

$$q(\mathbf{z}_t | \boldsymbol{\psi}_t, \mathbf{x}_t) = \mathcal{N}(\boldsymbol{\psi}_t, \boldsymbol{\Sigma}_{\mathbf{z}_t}), \quad \boldsymbol{\Sigma}_{\mathbf{z}_t} = \mathcal{T}_{\phi_2}(\mathbf{x}_t), \quad (10)$$

and the mean $\boldsymbol{\psi}_t$ is drawn from an implicit distribution, which generates its random samples by nonlinearly transforming random noise $\boldsymbol{\epsilon}_t \sim p(\boldsymbol{\epsilon})$ as

$$\boldsymbol{\psi}_t = \mathcal{T}_{\phi_1}([\mathbf{x}_t, \boldsymbol{\epsilon}_t]), \quad \boldsymbol{\epsilon}_t \sim p(\boldsymbol{\epsilon}). \quad (11)$$

We choose $p(\boldsymbol{\epsilon}) = \mathcal{N}(0, 4\mathbf{I})$ in this paper. Note while the probability density function of $p(\boldsymbol{\epsilon})$ is analytic, that of $\boldsymbol{\psi}_t$ is implicit if the transformation \mathcal{T}_{ϕ_1} is not invertible.

While given $\boldsymbol{\epsilon}_t$ and \mathbf{x}_t , the local latent variable \mathbf{z}_t follows a diagonal Gaussian distribution, the marginal distribution $q(\mathbf{z}_t | \mathbf{x}_t)$, obtained by integrating out the random noise $\boldsymbol{\epsilon}_t$, becomes an implicit distribution that is no longer restricted to follow a diagonal Gaussian as in Section 3.2. Thus, given the same mean reward mapping function as in (7), we can better capture the uncertainty on the mean rewards by sampling the local latent variable \mathbf{z}_t from a more flexible contextual variational distribution $q(\mathbf{z}_t | \mathbf{x}_t)$ as in (9).

While the original ELBO becomes intractable given an implicit contextual variational distribution, as in Yin & Zhou (2018) and Molchanov et al. (2019), we can optimize a lower bound of the ELBO that is amenable to direct optimization via stochastic gradient descent (SGD): we sample $K + 1$ $\boldsymbol{\psi}_t$'s, use only one of them to sample \mathbf{z}_t from the conditional distribution $q(\mathbf{z}_t | \boldsymbol{\psi}_t, \mathbf{x}_t)$, and combine them for computing a lower bound of the ELBO as

$$\underline{\mathcal{L}}_{K,t} = \mathbb{E}_{\boldsymbol{\epsilon}_t^{(0)}, \dots, \boldsymbol{\epsilon}_t^{(K)} \stackrel{iid}{\sim} p(\boldsymbol{\epsilon})} \mathbb{E}_{\mathbf{z}_t \sim q(\cdot | \boldsymbol{\psi}_t^{(0)}, \mathbf{x}_t)} \left[\ln p(r_t | \mathbf{x}_t, \mathbf{z}_t) + \ln \frac{p(\mathbf{z}_t)}{\frac{1}{K+1} \sum_{k=0}^K q(\mathbf{z}_t | \boldsymbol{\psi}_t^{(k)}, \mathbf{x}_t)} \right],$$

where $\boldsymbol{\psi}_t^{(k)} := \mathcal{T}_{\phi_1}([\mathbf{x}_t, \boldsymbol{\epsilon}_t^{(k)}])$. Different from related works (Ranganath et al., 2016; Maaløe et al., 2016) that also employ a hierarchical variational distribution, SIVI allows $q_\phi(\boldsymbol{\psi})$ to follow an implicit distribution (Huszár, 2017; Tran et al., 2017) and directly optimizes a surrogate ELBO.

We describe TS guided by semi-implicit LU in Algorithm 3 in the Appendix. The value of K is related to how close $\underline{\mathcal{L}}_{K,t}$ is to \mathcal{L}_t . A moderate value of $K = 50$ is found to be sufficient for neural network training, which does not bring much extra computational cost. Benefiting from the expressiveness improvement of using the semi-implicit structure, the distribution of $\mathbb{E}[r_t | \mathbf{x}_t, \mathbf{z}_t]$ under $q(\mathbf{z}_t | \mathbf{x}_t)$ becomes more flexible and can fit more complicated mean reward distributions. While this added flexibility may slightly degrade the performance for problems with simple reward distributions, overall, semi-implicit local uncertainty is found to work better than Gaussian local uncertainty.

4. Experiments

We evaluate Gaussian variational LU guided TS, referred to as LU-Gauss, and semi-implicit variational LU guided TS, referred to as LU-SIVI, on the contextual bandits benchmark used in Riquelme et al. (2018). We consider eight different datasets from this benchmark, including Mushroom, Financial, Statlog, Jester, Wheel, Covertypes, Adult, and Census, which exhibit a wide variety of statistical properties. Details on these datasets are provided in Table 3. For both LU-Gauss and LU-SIVI, we choose the Adam optimizer with the learning rate set as 10^{-3} . Python (TensorFlow 1.14) code for both LU-Gauss and LU-SIVI is available at <https://github.com/Zhendong-Wang/Thompson-Sampling-via-Local-Uncertainty>

We compare the proposed LU-Gauss and LU-SIVI to Lin-FullPost, BBB, and Neural Linear implemented in Riquelme et al. (2018), and π -TS-DGF of Zhang et al. (2019). The details on the neural network structures are provided in the Appendix. Since each experiment involves randomly sampling a subset of contextual vectors from the full dataset, all results are averaged over 50 independent random trials. In

each random trial, we rerun the code of LinFullPost, BBB, and Neural Linear provided by Riquelme et al. (2018) and the code of π -TS-DGF provided by Zhang et al. (2019). To ensure a fair comparison, these 50 random sequences for each dataset are made the same across all the aforementioned algorithms. Note we have also considered making comparison with functional variational Bayesian neural networks (FBNNs) of Sun et al. (2019). However, as the code for FBNNs is too computationally expensive for us to run as many as 50 independent random trials for each of the eight benchmark datasets, we defer the details of an informal comparison with FBNNs to the Appendix.

4.1. Exploratory Analysis

To analyze how the performance is impacted by the expressiveness of the variational distribution $q(z_t | x_t)$ that is used to model the local uncertainty, we choose the Mushroom dataset as an example, in which the stochastic rewards exhibit multi-modality. The Mushroom dataset has two distinct classes: Poisonous and Safe, and the agent has two possible actions: Eat or Not Eat. Eating a safe mushroom will be awarded +5, while eating a poisonous one will be awarded either +5 or -35, which are equally likely to occur. If the agent chooses to not eat the mushroom, it will receive 0 reward. Thus, given a poisonous mushroom, the true reward distribution of Action Not Eat has a single mode at zero, while that of Action Eat has two modes, +5 and -35, with mean -15. For this reason, given a poisonous mushroom, a variational distribution that is not flexible enough will face the risk of concentrating the high density region of its mean reward distribution of Action Eat around +5, which is a local mode, leading to the wrong action; by contrast, a sufficiently flexible variational distribution could escape the local mode at +5 even if it initially concentrates its mean reward distribution around it, leading to better exploration.

Running both LU-Gauss and LU-SIVI on the same random sequence of contextual vectors selected from Mushroom, we perform two different exploratory analyses: 1) We pick one poisonous mushroom, whose contextual feature vector x_p is being regarded as edible at the beginning of training by both LU-Gauss and LU-SIVI, and visualize how the histogram of the sampled mean rewards of x_p at each time step changes as the training progresses; 2) We visualize how the histogram of the sampled mean rewards of all poisonous mushroom at each step changes as the training progresses.

More specifically, for the first analysis, for $s = 1, \dots, S$, we use (8) to sample $z_{t,p}^{(s)} | x_p \sim q(z_{t,p} | x_p)$ for LU-Gaussian, or use (10) and (11) to sample $z_t^{(s)} | x_p \sim q(z_t | \psi_t^{(s)}, x_p)$, $\psi_t^{(s)} | x_p \sim q(\psi_t | x_p)$ for LU-SIVI, and then use (7) to compute a mean reward given x_p and $z_{t,p}^{(s)}$ as $\mathbb{E}[r_{t,p}^{(s)} | x_p, z_{t,p}^{(s)}] = \mathcal{T}_\theta([x_p, z_{t,p}^{(s)}])$. We draw $S = 2000$

mean reward samples for each t to form a histogram for that training step, and visualize these histograms over training steps as a heatmap, as shown in Figure 1.

For Action Not Eat, as shown in Figures 1 (a) and (c), the empirical distribution of the sampled mean rewards for this poisonous mushroom becomes more and more concentrated around zero for both LU-Gauss and LU-SIVI. For Action Eat, as shown in Figures 1 (b) and (d), while both LU-Gauss and LU-SIVI initially consider this poisonous mushroom x_p as edible with a positive reward, their mean reward distributions given x_p become more and more different as the training progresses. In particular, LU-Gauss first becomes less certain and then more certain, with its high density region of the mean reward empirical distribution first shifting below zero, then moving back above zero, and eventually concentrating around +5, which means it mistakenly treats this poisonous mushroom as edible at many training steps. By contrast, LU-SIVI gradually increases its uncertainty and then maintains it around the same level, with its high density region of the mean reward empirical distribution quickly shifting below zero and remaining below zero at most of the training steps, which means it correctly identifies this poisonous mushroom most of the time.

For the second analysis, we take a single random sample of the mean reward for each mushroom at t , use the sampled mean rewards over all poisonous mushrooms as $\{r_{t,p}\}_p$ to form a histogram, and visualize the histograms over time as a heatmap, as shown in Figure 2. For Action Not Eat, as shown in Figures 2 (a) and (c), both LU-Gauss and LU-SIVI quickly capture the underlying reward distribution, concentrating the histograms around zero. For Action Eat, as shown in Figure 2 (b), under LU-Gauss, the sampled mean rewards of all poisonous mushrooms gradually concentrate around two density modes, with one clearly below zero and the other clearly above zero, suggesting that LU-Gauss will make a large number of mistakes in treating poisonous mushrooms as edible. By contrast, as shown in Figure 2 (d), under LU-SIVI, the sampled mean rewards of all poisonous mushrooms quickly move down their high density region below zero and then maintain a single density mode around -30, suggesting that LU-SIVI will correctly identify poisonous mushrooms most of the time.

These analyses suggest that for the Mushroom dataset, whose true reward distribution of eating a poisonous mushroom exhibits two density modes, the variational distribution of LU-Gauss shown in (8) underperforms that of LU-SIVI shown in (9) in exploration and faces a greater risk to concentrate its mean reward distribution around the undesired local mode, leading to poorer performance. LU-SIVI, which introduces a more flexible variational distribution that is also amenable to optimization via SGD, achieves a better balance between exploration and exploitation.




Figure 1. The distribution of mean reward on a single poisonous mushroom




Figure 2. Convergence of the distribution of mean reward on all poisonous mushrooms

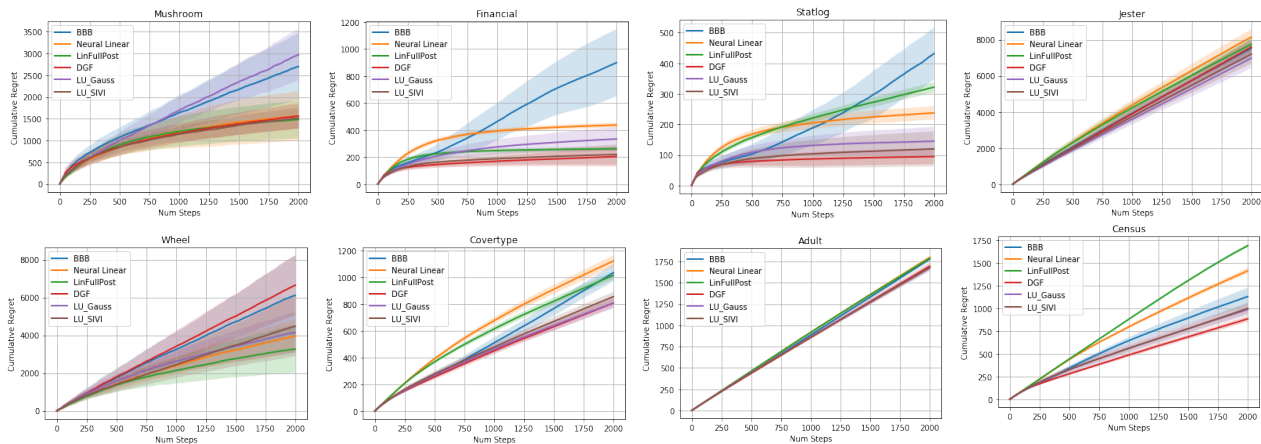


Figure 3. Comparison of Cumulative Regrets over eight different datasets. The solid line is the average performance over 50 random seeds, with the shaded area representation \pm one standard error.

4.2. Performance Comparison

Following the same experimental settings of the contextual bandits benchmark in Riquelme et al. (2018), we evaluate the proposed LU-Gaussian and LU-SIVI and compare them to representative TS algorithms relying on global uncertainty. Neither LU-Gaussian nor LU-SIVI include noise injection to the global parameters, dropout layer, and bootstrapping techniques discussed in Riquelme et al. (2018). These techniques, designed to better capture global uncertainty to guide TS, can potentially be combined with LU-Gaussian and LU-SIVI to further improve their performance. We leave that for future study.

In Figure 3, for each algorithm on a dataset, we plot the mean (a colored line) and standard error (line shade) of its accumulative regret against the training step over 50 random trials. Using the performance of the Uniform algorithm, which uniformly at random chooses its actions from \mathcal{A} , as

the reference, we show in Table 1 the normalized cumulative regrets and in Figure 4 their boxplots.

We first examine the performance of various global uncertainty guided TS algorithms, including BBB, LinFullPost, Neural Linear, and π -TS-DGF. We find that BBB has the worst overall performance and exhibits large variance on its cumulative regrets across different random trials, suggesting that using a diagonal Gaussian variational distribution on the global parameters (neural network weights) leads to poor exploration. LinFullPost performs well in some datasets. E.g., it works very well on the Mushroom dataset, which is likely because using Gibbs sampling on global variables makes it simple to move its reward distribution of eating a poisonous mushroom away from the bad local mode of $+5$. However, it clearly suffers from the lack of representation power on datasets that violate the linear assumption, such as Covtype, Census, Adult, and Statlog. Neural Linear involves feature extraction using a neural network and keeps

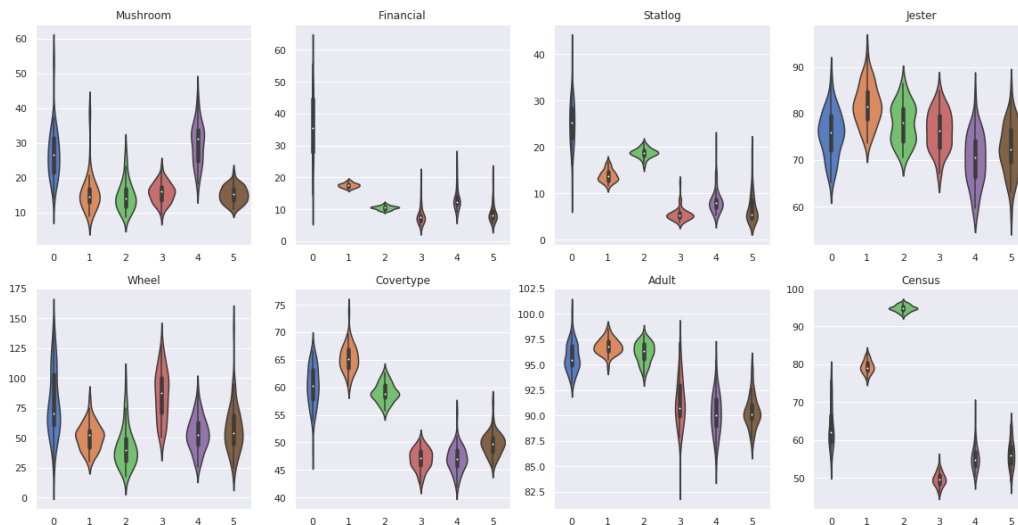


Figure 4. Boxplots of Normalized Cumulative Regrets on eight different datasets. The algorithms are ordered as: 0. BBB, 1. Neural Linear, 2. LinFullPost, 3. π -TS-DGF, 4. LU-Gauss, 5. LU-SIVI.

the exact linear posterior updates for TS, but its representation learning relies heavily on the number of training steps that have been taken, which might be the reason for its relatively poor performance. Since π -TS-DGF of Zhang et al. (2019) can provide better uncertainty estimation than BBB on the global parameters and more representation power than LinFullPost, it works quite well on some datasets and overall outperforms BBB and LinFullPost. Nevertheless, it does not perform that well on Mushroom and Jester and performs poorly on Wheel, which requires heavy exploration in order to minimize the cumulative regrets.

We then examine the performance of the proposed LU guided TS algorithms, including LU-Gauss and LU-SIVI. As shown by Figures 3 and 4 and Table 1, LU-Gauss in general performs well, except that it provides poor performance on Mushroom, which, as analyzed in detail in Section 4.1, is likely because the diagonal Gaussian variational distribution is not flexible enough to encourage sufficient exploration. To be more specific, as in Figures 1(b) and 2(b), it prevents the sampled mean rewards of eating some poisonous mushrooms from moving away from a bad local density mode. By contrast, LU-SIVI performs well across all eight benchmark datasets. In particular, for Mushroom that LU-Gauss performs poorly on, as shown in Figures 1(d) and 2(d), LU-SIVI places most of its sampled mean rewards of eating poisonous mushrooms clearly below zero. This can be explained by the improved ability of LU-SIVI in exploration due to its use of a semi-implicit variational distribution that is more flexible but remains simple to optimize.

To better compare the overall performance of different algorithms, as shown in Table 1, for each algorithm, we follow Riquelme et al. (2018) and Sun et al. (2019) to compute both the Mean Rank and Mean Value of normalized cumu-

lative regrets over all eight benchmark datasets. The Mean Rank and Mean Value suggest that LU-SIVI has the best overall performance, followed by LU-Gauss and π -TS-DGF. In Appendix D, we further compare various algorithms in term of the Simple Regret metric (Riquelme et al., 2018).

4.3. Ablation Study

We introduce an ablation study to demonstrate the effectiveness of using local uncertainty. In the ablation baseline, everything stays the same as LU-Gauss/SIVI except that $\mathbf{z} \sim q(\mathbf{z})$ becomes a global variable that no longer depends on \mathbf{x}_t in q , and $\text{KL}(q(\mathbf{z})||p(\mathbf{z}))$ will be the single KL term shared by all observed data and hence needs to be appropriately rescaled in the objective function for stochastic variational inference. Specifically, in our framework, we change \mathbf{z} from via local uncertainty to via global uncertainty by replacing the input \mathbf{x} of $q(\mathbf{z}|\mathbf{x})$ with a constant vector of ones, which leads to \mathbf{z} shared by all data. We show the results of ablation baselines in Table 2. In general, LU-Gauss outperforms LU-Gauss-Ablation and LU-SIVI outperforms LU-SIVI-Ablation, suggesting the advantages of using local uncertainty over global uncertainty in our applications.

4.4. Runtime comparison

We report the time cost based on an Nvidia 1080-TI GPU. Note when the contextual vector dimension is low, both LU-Gauss and LU-SIVI take more time to run than LinFullPost. For example, on Mushroom whose contextual vector dimension is 22, LU-Gauss and LU-SIVI take about 18 and 30 seconds for 2000 training steps, while LinFullPost and Neural Linear take about 15 and 22 seconds. However, the computational complexity of LinFullPost increases cubically with the dimension of the contextual vector, due to

Table 1. Comparison of Normalized Cumulative Regret between various methods, with the normalization performed with respect to the Cumulative Regret of Uniform. For each dataset, the same set of 50 random contextual sequences are used for all algorithms. For each algorithm on a given dataset, we report its mean and standard error over these 50 independent random trials.

Algorithms	Mean Rank	Mean Value	Mushroom	Financial	Statlog	Jester	Wheel	Covertime	Adult	Census
Uniform	7	100.00	100.00 \pm 4.95	100.00 \pm 11.63	100.00 \pm 1.04	100.00 \pm 7.67	100.00 \pm 5.73	100.00 \pm 0.88	100.00 \pm 0.50	100.00 \pm 0.63
BBB	4.75	56.53	26.08 \pm 7.34	34.31 \pm 9.72	25.93 \pm 4.62	75.58 \pm 4.06	71.12 \pm 18.51	60.37 \pm 3.91	95.64 \pm 1.29	63.15 \pm 4.29
Neural Linear	4.75	52.17	15.77 \pm 6.84	17.15 \pm 0.71	14.00 \pm 1.28	81.89 \pm 3.32	46.51 \pm 10.80	65.37 \pm 1.92	97.06 \pm 1.01	79.57 \pm 1.99
LinFullPost	3.75	51.15	13.66 \pm 3.77	10.16 \pm 0.65	18.77 \pm 0.79	78.03 \pm 3.18	38.38 \pm 13.05	58.83 \pm 1.62	96.05 \pm 0.97	95.26 \pm 0.77
π -TS-DGF	2.5	46.63	15.24 \pm 2.87	8.00 \pm 3.32	6.10 \pm 2.75	75.98 \pm 3.71	81.16 \pm 22.20	46.81 \pm 2.31	90.77 \pm 2.00	49.02 \pm 1.96
LU-Gauss	2.75	46.03	31.40 \pm 6.74	13.10 \pm 3.51	8.73 \pm 2.80	70.19 \pm 4.38	52.90 \pm 15.54	47.33 \pm 2.56	89.30 \pm 2.04	55.28 \pm 2.83
LU-SIVI	2.5	45.42	14.84 \pm 2.87	8.28 \pm 3.16	7.62 \pm 4.08	71.64 \pm 4.52	63.06 \pm 21.84	49.86 \pm 2.25	89.88 \pm 1.96	58.18 \pm 5.45

Table 2. Analogous table to Table 1 for Ablation Study.

Algorithms	Mean Rank	Mean Value	Mushroom	Financial	Statlog	Jester	Wheel	Covertime	Adult	Census
Uniform	5	100.00	100.00 \pm 4.95	100.00 \pm 11.63	100.00 \pm 1.04	100.00 \pm 7.67	100.00 \pm 5.73	100.00 \pm 0.88	100.00 \pm 0.50	100.00 \pm 0.63
LU-Gauss	2.375	46.03	31.40 \pm 6.74	13.10 \pm 3.51	8.73 \pm 2.80	70.19 \pm 4.38	52.90 \pm 15.54	47.33 \pm 2.56	89.30 \pm 2.04	55.28 \pm 2.83
LU-Gauss-Ablation	2.5	46.55	21.28 \pm 10.57	12.35 \pm 5.09	7.45 \pm 2.83	69.98 \pm 4.63	64.03 \pm 14.14	49.85 \pm 3.12	91.10 \pm 2.11	56.37 \pm 2.71
LU-SIVI	2.5	45.42	14.84 \pm 2.87	8.28 \pm 3.16	7.62 \pm 4.08	71.64 \pm 4.52	63.06 \pm 21.84	49.86 \pm 2.25	89.88 \pm 1.96	58.18 \pm 5.45
LU-SIVI-Ablation	2.625	47.81	20.32 \pm 11.35	9.01 \pm 2.71	6.34 \pm 3.53	71.10 \pm 4.44	79.33 \pm 23.50	49.90 \pm 2.18	92.40 \pm 2.17	54.07 \pm 3.68

Table 3. Details of the benchmark datasets and comparison of runtime (in seconds) between various methods. The reported runtime values are approximated with a single run.

		Mushroom	Financial	Statlog	Jester	Wheel	Covertime	Adult	Census
Dataset	Context dimension	22	21	16	32	2	54	94	389
Information	Number of actions	2	8	7	8	5	7	14	9
Algorithms	Uniform	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04
	BBB	13	13	12	13	12	13	13	15
	Neural Linear	22	35	33	35	28	33	50	40
	LinFullPost	14	7	4	9	3	13	48	710
	π -TS-DGF	270	107	82	128	62	168	246	750
	LU-Gauss	18	18	18	18	18	18	19	21
	LU-SIVI	32	30	30	30	30	31	31	40

the need to perform both Cholesky decomposition of the covariance matrix and matrix inversion when sampling the regression coefficient vector β . Neural Linear involves a global latent layer z with a fixed dimension to do runtime control. In large contextual dimension case, for example, on Census whose contextual vector dimension is 389, it takes LU-Gauss and LU-SIVI about 20 and 40 seconds, respectively, to run 2000 training steps, while it takes LinFullPost and Neural Linear 710 and 40 seconds, respectively. For π -TS-DGF, it takes about 270 seconds on Mushroom and 750 seconds on Census, due to its large number of particles. Details of runtime comparison are shown in Table 3.

5. Conclusion

To address the problem of contextual bandits, we propose Thompson sampling (TS) guided by local uncertainty (LU). This new TS framework models the reward distribution given the context using a latent variable model, and uses a pre-posterior contextual variational distribution to approximately capture the uncertainty of the local latent variable, whose true posterior depends on both the observed context

and the reward that is yet to be observed. Under this framework, we introduce both LU-Gauss, which uses a diagonal Gaussian contextual variational distribution, and LU-SIVI, which uses a semi-implicit one. In comparison to LU-Gauss, LU-SIVI has a more flexible variational distribution that enhances its ability of exploration, leading to improved performance on datasets with complex reward distributions. Experimental results on eight different contextual bandit datasets demonstrate that both LU-Gauss and LU-SIVI well model the uncertainty and provide good performance, exhibiting reliability and robustness across all datasets. In particular, both LU-Gauss and LU-SIVI perform competitively to the particle-interactive TS algorithm of Zhang et al. (2019), the current state-of-the-art method, but have clearly lower computational complexity. The improved expressiveness, robustness, and computational complexity is the reason for the proposed local uncertainty guided TS method to claim a spot among the state-of-the-art. An interesting topic for future research is to investigate how to integrate both global parameter uncertainty and local latent variable uncertainty under TS to achieve further improved performance.

Acknowledgements

The authors thank the anonymous reviewers, whose invaluable comments and suggestions have helped us to improve the paper. This research was supported in part by Award IIS1812699 from the U.S. National Science Foundation. The authors acknowledge the Texas Advanced Computing Center (TACC) at The University of Texas at Austin for providing HPC resources that have contributed to the research results reported within this paper. URL: <http://www.tacc.utexas.edu>

References

- Agrawal, S. and Goyal, N. Analysis of Thompson sampling for the multi-armed bandit problem. In *Conference on Learning Theory*, pp. 39–1, 2012.
- Agrawal, S. and Goyal, N. Thompson sampling for contextual bandits with linear payoffs. In *International Conference on Machine Learning*, 2013.
- Arulkumaran, K., Deisenroth, M. P., Brundage, M., and Bharath, A. A. Deep reinforcement learning: A brief survey. *IEEE Signal Processing Magazine*, 34(6):26–38, 2017.
- Auer, P. Using confidence bounds for exploitation-exploration trade-offs. In *Journal of Machine Learning Research*, 2002.
- Bishop, C. M. *Pattern Recognition and Machine Learning*. Springer, 2006.
- Bishop, C. M. and Tipping, M. E. Variational relevance vector machines. *UAI*, pp. 46–53, 2000.
- Blei, D. M., Kucukelbir, A., and McAuliffe, J. D. Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112(518):859–877, 2017.
- Blundell, C., Cornebise, J., Kavukcuoglu, K., and Wierstra, D. Weight uncertainty in neural networks. In *International Conference on Learning Representations*, 2015.
- Cesa-Bianchi, N., Gentile, C., Lugosi, G., and Neu, G. Boltzmann exploration done right. In *Advances in Neural Information Processing Systems*, 2017.
- Chapelle, O. and Li, L. An empirical evaluation of Thompson sampling. In *Advances in neural information processing systems*, pp. 2249–2257, 2011.
- Djallel, B. and Irina, R. A survey on practical applications of multi-armed and contextual bandits. *arXiv preprint arXiv:1904.10040*, 2019.
- Fortunato, M., Gheshlaghi Azar, M., Piot, B., Menick, J., Osband, I., Graves, A., Mnih, V., Munos, R., Hassabis, D., Pietquin, O., Blundell, C., and Legg, S. Noisy networks for exploration. In *arXiv:1706.10295*, 2017.
- François-Lavet, V., Henderson, P., Islam, R., Bellemare, M. G., Pineau, J., et al. An introduction to deep reinforcement learning. *Foundations and Trends® in Machine Learning*, 11(3-4):219–354, 2018.
- Gal, Y. and Ghahramani, Z. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *International Conference on Machine Learning*, 2016.
- Graves, A. Practical variational inference for neural networks. In *Advances in Neural Information Processing Systems*, pp. 2348–2356, 2011.
- Hernández-Lobato, J. M. and Adams, R. Probabilistic back-propagation for scalable learning of Bayesian neural networks. In *International Conference on Machine Learning*, pp. 1861–1869, 2015.
- Hinton, G. and Van Camp, D. Keeping neural networks simple by minimizing the description length of the weights. In *in Proc. of the 6th Ann. ACM Conf. on Computational Learning Theory*. Citeseer, 1993.
- Huszár, F. Variational inference using implicit distributions. *arXiv preprint arXiv:1702.08235*, 2017.
- Jordan, M. I., Ghahramani, Z., Jaakkola, T. S., and Saul, L. K. An introduction to variational methods for graphical models. *Machine learning*, 37(2):183–233, 1999.
- Kingma, D. P. and Welling, M. Auto-encoding variational Bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Li, C., Chen, C., Carlson, D., and Carin, L. Preconditioned stochastic gradient langevin dynamics for deep neural networks. In *Association for the Advancement of Artificial Intelligence*, 2016.
- Maaløe, L., Sønderby, C. K., Sønderby, S. K., and Winther, O. Auxiliary deep generative models. In *ICML*, pp. 1445–1453, 2016.
- Mandt, S., Hoffman, M. D., and Blei, D. M. A variational analysis of stochastic gradient algorithms. In *International Conference on Machine Learning*, 2016.
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., and Riedmiller, M. Playing Atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. Human-level control

- through deep reinforcement learning. *Nature*, 518(7540): 529, 2015.
- Mnih, V., Badia, A. P., Mirza, M., Graves, A., Lillicrap, T., Harley, T., Silver, D., and Kavukcuoglu, K. Asynchronous methods for deep reinforcement learning. In *International Conference on Machine Learning*, pp. 1928–1937, 2016.
- Molchanov, D., Kharitonov, V., Sobolev, A., and Vetrov, D. Doubly semi-implicit variational inference. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 2593–2602, 2019.
- Neal, R. M. *Bayesian learning for neural networks*, volume 118. Springer Science & Business Media, 2012.
- Osband, I., Blundell, C., Pritzel, A., and Van Roy, B. Deep exploration via bootstrapped DQN. In *Advances in Neural Information Processing Systems*, 2016.
- Plappert, M., Houthoofd, R., Dhariwal, P., Sidor, S., Y. Chen, R., Chen, X., Asfour, T., Abbeel, P., and Andrychowicz, M. Parameter space noise for exploration. In *arXiv:1706.01905*, 2017.
- Ranganath, R., Tran, D., and Blei, D. Hierarchical variational models. In *ICML*, pp. 324–333, 2016.
- Riquelme, C., Tucker, G., and Snoek, J. Deep Bayesian bandits showdown: An empirical comparison of Bayesian deep networks for Thompson sampling. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=SyYe6k-CW>.
- Russo, D. J., Van Roy, B., Kazerouni, A., Osband, I., and Wen, Z. A tutorial on Thompson sampling. *Foundations and Trends® in Machine Learning*, 11(1):1–96, 2018.
- Sun, S., Zhang, G., Shi, J., and Grosse, R. Functional variational Bayesian neural networks. In *International Conference on Learning Representations*, 2019.
- Sutton, R. S. and Barto, A. G. *Introduction to reinforcement learning*, volume 2. MIT press Cambridge, 1998.
- Thompson, W. R. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933.
- Tran, D., Ranganath, R., and Blei, D. Hierarchical implicit models and likelihood-free variational inference. In *Advances in Neural Information Processing Systems*, pp. 5523–5533, 2017.
- Welling, M. and Teh, Y. W. Bayesian learning via stochastic gradient Langevin dynamics. In *ICML*, pp. 681–688, 2011.
- Yin, M. and Zhou, M. Semi-implicit variational inference. In *International Conference on Machine Learning*, 2018.
- Zhang, H., Chen, B., Tian, L., Wang, Z., and Zhou, M. Variational hetero-encoder randomized GANs for joint image-text modeling. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=H1x5wRVtvS>.
- Zhang, R., Wen, Z., Chen, C., and Carin, L. Scalable Thompson sampling via optimal transport. In *Artificial Intelligence and Statistics*, 2019.