GuideBoot: Guided Bootstrap for Deep Contextual Bandits in Online Advertising

Feiyang Pan*

Institute of Computing Technology,
Chinese Academy of Sciences.
Beijing, China

Wei Wang

Tencent
Shenzhen, China

Haoming Li*

Institute of Computing Technology,
Chinese Academy of Sciences
Beijing, China

Yanrong Kang

Tencent
Shenzhen, China

Xiang Ao*[†]

Institute of Computing Technology,
Chinese Academy of Sciences
Beijing, China

Ao Tan

Tencent
Shenzhen, China

Qing He*

Institute of Computing Technology,
Chinese Academy of Sciences
Beijing, China

ABSTRACT

The exploration/exploitation (E&E) dilemma lies at the core of interactive systems such as online advertising, for which contextual bandit algorithms have been proposed. *Bayesian* approaches provide guided exploration via uncertainty estimation, but the applicability is often limited due to over-simplified assumptions. Non-Bayesian *bootstrap* methods, on the other hand, can apply to complex problems by using deep reward models, but lack a clear guidance to the exploration behavior. It still remains largely unsolved to develop a practical method for complex deep contextual bandits.

In this paper, we introduce Guided Bootstrap (GuideBoot), combining the best of both worlds. GuideBoot provides explicit guidance to the exploration behavior by training multiple models over both real samples and noisy samples with fake labels, where the noise is added according to the predictive uncertainty. The proposed method is *efficient* as it can make decisions on-the-fly by utilizing only one randomly chosen model, but is also *effective* as we show that it can be viewed as a non-Bayesian approximation of Thompson sampling. Moreover, we extend it to an online version that can learn solely from streaming data, which is favored in real applications. Extensive experiments on both synthetic tasks and large-scale advertising environments show that GuideBoot achieves significant improvements against previous state-of-the-art methods.

*At Key Lab of Intelligent Information Processing of Chinese Academy of Sciences (CAS). Also at University of Chinese Academy of Sciences (UCAS).

[†] Correspondence to: Xiang Ao <aoxiang@ict.ac.cn>.

This paper is published under the Creative Commons Attribution 4.0 International (CC-BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

WWW '21, April 19–23, 2021, Ljubljana, Slovenia

© 2021 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC-BY 4.0 License.

ACM ISBN 978-1-4503-8312-7/21/04.

<https://doi.org/10.1145/3442381.3449987>

CCS CONCEPTS

• Computing methodologies → Machine learning; • Information systems → Online advertising.

KEYWORDS

Deep Contextual Bandits, Bootstrap, Online Advertising

ACM Reference Format:

Feiyang Pan, Haoming Li, Xiang Ao, Wei Wang, Yanrong Kang, Ao Tan, and Qing He. 2021. GuideBoot: Guided Bootstrap for Deep Contextual Bandits in Online Advertising. In *Proceedings of the Web Conference 2021 (WWW '21)*, April 19–23, 2021, Ljubljana, Slovenia. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3442381.3449987>

1 INTRODUCTION

With the progress of deep learning, machine-learning-powered online advertising and recommender systems have achieved significant success in recent years for its ability to model complex deep dependencies between users, ads, and contexts. However, for cost-sensitive decision-making problems such as the cold-start problem, it is still challenging to develop practical methods for minimizing the cumulative regret.

Online advertising can be viewed as a massive decision-making problem, where the ultimate goal is to deliver the right ads to the right users. The trade-off between exploration and exploitation (E&E) is crucial to solve this problem, as the decision-maker can receive feedback only through interaction with the environment. Such a problem can be formulated as contextual bandits, for which many algorithms have been proposed and proved successful either theoretically or empirically in specific settings.

There are two major tracks for exploration in contextual multi-armed bandits: Bayesian and non-Bayesian methods. Bayesian methods provide *guided* exploration by estimating the model uncertainty with Bayesian posterior inference and make decisions thereby. For example, for linear bandits where the dependencies between inputs and rewards are linear functions, LinUCB [22] and Thompson Sampling [4] rely on Bayesian Linear Regression (BLR)

Figure 1: Inference and training of GuideBoot. The solid lines are feed-forward passes for inference, making use of one model out of K models. The dashed lines are for training, including a data generation process for each model (red dashed lines) and back-propagation to train each model (green dashed lines).

to infer the posterior reward distribution. For Bernoulli contextual bandits, GLM-UCB [8] uses Bayesian logistic regression to derive the policy. There are also studies to use Gaussian Processes [19] and Neural Processes [11] for posterior inference in non-linear settings. However, these methods are difficult to scale because exact probabilistic inference is often intractable in complex problems with high-dimensional inputs and non-linear dependencies between features and outcomes.

In contrast, non-Bayesian methods do not require a specific functional form of the underlying model, thus are suitable for complex problems. A typical choice is the *bootstrap* heuristic, i.e., to train multiple reward models (i.e., deep neural networks) with different initialization [24] as well as on randomized subsets of samples [31], and to randomly use one of them to make a decision. However, the behavior of such a bootstrap exploration method is not explicitly guided. The sub-policy induced by each single model is essentially greedy, so there is a big chance that the overall policy eventually degenerates to a sub-optimal greedy policy. Therefore, it still remains largely unsolved to develop practical methods for complex deep contextual bandits.

In this paper, we propose a novel contextual bandit algorithm named Guided Bootstrap (GuideBoot for short), which combines the best of both Bayesian and non-Bayesian methods. The high-level idea of GuideBoot is to provide explicit uncertainty-based guidance to the bootstrap resampling behavior.

Similar to traditional bootstrap-based methods, our method trains a bag of reward models and randomly chooses one to make a decision at each step. In contrast to previous methods, we not only use different subsets of collected samples to train the models, but also randomly generate a small number of fake samples while training the models, so as to enable distributional predictions and preclude overfitting. Particularly, every fake sample is generated randomizing the reward signal of a real sample, and the probability of generating such a fake sample is in proportion to the predictive uncertainty of the reward. In this way, unfamiliar contexts and rarely chosen actions can have chance to be explored.

An essential benefit of GuideBoot is that it can always make decisions on the fly with minimal computational complexity during inference. As shown in Figure 1, it leaves the time-consuming uncertainty estimation to the training stage rather than the inference (decision-making) stage. Therefore, GuideBoot can be incorporated with arbitrary reward models (i.e., modern deep neural networks for click-through rate prediction) and any type of uncertainty estimation approaches ranging from simple non-parametric count-based methods to sophisticated Bayesian methods.

To better adapt to real applications of high-frequency decision-making such as online advertising and recommender systems, we put forward an extended version of GuideBoot that can handle streaming data efficiently, namely Online GuideBoot. The biggest challenge for developing the online algorithm is that the agent can only learn from streaming data without accessing a complete history. Our algorithm replaces the bootstrap resampling step with a shuffle-and-split step which operates solely on an online buffer, which is easy to implement and empirically powerful.

We conduct on systematical experiments on synthetic tasks and two large-scale real-world tasks on industrial applications of online advertising. The experimental results demonstrate that our methods perform better and are more stable than the state-of-the-arts in both classic settings and online settings.

The remainder of this paper is organized as follows. Section 2 introduces the problem settings and backgrounds. Section 3 details GuideBoot in a standard setting where the agent maintains a complete history. Section 4 illustrates how to deploy our algorithm in an online setting with streaming data, where the agent only learns from recent experiences without accessing to the whole history. Section 5 demonstrates the experiments.

2 PRELIMINARIES

2.1 Contextual multi-armed bandits

We consider a standard contextual multi-armed bandit problem [21, 22]. At each step, the environment reveals a *context* (i.e., the user attributes in a recommender system) and m candidate actions (i.e., items in the platform). Without loss of generality, we write $x_{t,a} \in \mathcal{X} \subset \mathbb{R}^d$ as input features of the a -th action at step t , including the relevant context information. The environment also defines the expected reward for each action by a stationary function $y_{t,a} \triangleq \mathbb{E}[r_t(a)] = f^*(x_{t,a})$, where the reward function $f^*(\cdot)$ is unknown to the agent. The agent is required to choose one action a_t (i.e., an item to recommend) out of m candidates $\{1, \dots, m\}$. For simplicity, we write $x_t \triangleq x_{t,a_t}$. Then, the environment draws a reward according to x_t . As many applications on the web involve optimizing binary outcomes (i.e., clicks or conversions), we assume the reward is drawn from a Bernoulli distribution, i.e., $r_t \sim \text{Bernoulli}(y_t)$, where $y_t \triangleq f^*(x_t) \in [0, 1]$ is the expected reward. We put a general assumption that $f^*(\cdot)$ is stationary when deriving the GuideBoot algorithm in Section 3. This stationary assumption can be relaxed when extends to the online version in Section 4.

2.2 Bayesian approaches

Bayesian methods for contextual bandits rely on accurate probabilistic inference of the expected reward so as to obtain an uncertainty

estimation for each action. Let $\mathcal{D}_{t-1} = \{(x_\tau, r_\tau)\}_{\tau=1}^{t-1}$ be the collected samples before time step t . A Bayesian method not only learns an approximate reward model $\hat{y}_a = f_\theta(x_a)$, but also learns the posterior distribution of the parameters $p(\theta|\mathcal{D}_{t-1})$. By using the posterior, UCB-type algorithms can derive the upper confidence bound $\hat{y}_a^{\text{UCB}}$ of the expected reward, i.e., $p(y_a < \hat{y}_a^{\text{UCB}}) > 1 - \delta$, and choose the action with the highest upper confidence bound

$$a_t = \arg \max_a \hat{y}_a^{\text{UCB}}.$$

Thompson sampling, on the other hand, randomly samples a set of parameters from the posterior $\theta' \sim p(\theta|\mathcal{D}_{t-1})$ to predict the reward, and make a decision thereby,

$$a_t = \arg \max_a f_{\theta'}(x_a).$$

The drawbacks for these Bayesian methods are two-folds. First, exact probabilistic inference is often intractable for real problems. When the reward function is assumed as a (generalized) linear model, the posterior can be inferred with Bayesian linear (logistic) regression. For complex problems, it is infeasible especially when one would like to use a deep neural net to predict the reward.

Second, even if we assume the posterior is known, there is a substantial increase in computing burden required for making a decision than a simple greedy policy. For linear bandits, LinUCB [22] and GLM-UCB [8] need a time complexity of $O(d^2)$ to compute the upper confidence bounds, an order of magnitude slower than an ordinary linear model to make predictions. Thompson sampling with linear payoffs requires $O(d^3)$ to sample the parameter from the multivariable Gaussian posterior distribution. In non-linear cases, the Gaussian Process bandit [19] requires even more computations to get the kernel matrix. Recent Bayesian neural nets that rely on Monte Carlo Variational Inference [3, 9, 10, 17] also significantly increase the inference time cost. Such slowdown in decision-making is often unacceptable in real applications. In this paper, the proposed method has the same computational complexity as a standard feed-forward neural net, which can make decisions on-the-fly.

2.3 Non-Bayesian approaches

On the contrary, non-Bayesian methods can be applied to problems with arbitrary complex reward functions. For example, the simplest heuristic ϵ -greedy chooses the $\arg \max_a f_\theta(x_a)$ at probability $1 - \epsilon$, and explores randomly otherwise. Boltzmann policy computes the softmax over predicted rewards of candidate actions to derive a stochastic policy, which is also shown competitive in [26]. Bootstrap-based exploration is also shown effective in both reinforcement learning [24, 27] and bandit problems [6, 7, 20, 31, 32], which either maintain multiple bootstrap samples of the history or train multiple reward models from different subsets of data. These methods can be incorporated with deep neural networks, therefore can achieve state-of-the-art performance in deep contextual bandits. Tang et al. [31] use a Poisson distribution to draw duplicated samples in order to mimic the resampling procedure with online streaming data. Elmachet et al. [7] proposed a contextual bandit algorithm that follows the same bootstrapping way but uses decision trees as the base learner. Apart from bootstrapping real samples and updating with real observed rewards, Kveton et al. [20] proposed Giro that designs pseudo rewards for each pull of the arm.

Wang et al. [33] bounded the pseudo rewards with the perturbed residual from the history.

However, these methods lack explicit guidance of an explicit notion of uncertainty or prior, therefore they tend to have sub-optimal regret comparing to the Bayesian methods in non-deep settings. Our work is orthogonal to these unguided exploration methods as our GuideBoot perturbs the bootstrapped samples under an explicit guidance of the uncertainty of the reward, where the guidance is achieved by manipulating fake samples to mimic a data prior.

2.4 Other related work

There are also studies on theoretical justification for bootstrap in the (contextual) multi-armed bandits problem. Osband and Van Roy [25] proposed a bandit algorithm named BootstrapThompson and showed the algorithm approximates Thompson sampling in Bernoulli bandits. Vaswani et al. [32] generalized it to categorical and Gaussian rewards. Hao et al. [13] extended UCB with multiplier bootstrap and derived both problem-dependent and problem-independent regret bounds for the proposed algorithm. Although with strong theoretical properties, these methods often limit to oversimplified settings. In this paper, we also provide a Bayesian point of view to explain our algorithm (in Section 3.4). We show that our GuideBoot can approximate the Bayesian posterior with an informative prior, which is more general and flexible than previous work [25] that is based on a degenerate prior.

3 GUIDEBOOT

3.1 Overview of the method

GuideBoot is a practical method for deep contextual bandits that can make decisions on the fly. The basic inference and training procedure is shown in Figure 1. In this part, we describe the procedure of model updates *over a logged dataset* (aka the *experience buffer*), as in the same setting with previous studies on bootstrap-based contextual bandits [20, 31]. Later in Section 4 we will introduce a more practical implementation of online GuideBoot with online *streaming* data.

A GuideBoot agent maintains K independently trained reward models $f^{(1)}(\cdot), \dots, f^{(K)}(\cdot)$. At the t -th step, when a query comes to the system with a number of candidate actions, the agent randomly selects one model $f^{(k)}(\cdot)$ to predict the scores of the actions and chooses the action with maximum score, i.e.,

$$a_t = \arg \max_a f^{(k)}(x_{t,a}). \quad (1)$$

For training, all models are updated in parallel with gradient descent over different bootstrapped subsets of data consisting of both real samples and randomly generated fake samples. The pseudo-code of GuideBoot is outlined in Algorithm 1.

The key novelty of GuideBoot is the uncertainty-based guidance for the trade-off between exploration and exploitation, achieved by training the bag of models with perturbed datasets including a small fraction of fake samples. In the rest of this chapter, we introduce *what* the fake samples are in Section 3.2, *how* to train the models using the fake samples in Section 3.3, and *why* the fake samples should be used in Section 3.4.

Table 1: Notations

Notation	Description
t	the time step
m	the number of candidate actions
K	the number of candidate models
$\mathcal{H}_{1:t}$	The <i>replay buffer</i> that stores all history samples
$\mathcal{H}_{\text{on}}$	The <i>online buffer</i> that temporarily collects a batch of samples before updating models
$x_{t,a}$	Input features of candidate action a (including the context feature) at step t
$\hat{y}_{t,a}$	The predicted expected reward of action a at step t
a_t	The selected action at step t
x_t	Simplified notation of x_{t,a_t}
$\tilde{\mathcal{B}}_i$	The i -th resampling subset of samples
$f^{(k)}(\cdot)$	The k -th reward model
$g(\cdot)$	The <i>guidance function</i> g that maps an input x to a probability of generating fake samples
$\hat{\rho}(\cdot)$	The unnormalized density estimator

3.2 Guidance from fake samples

We start with an example showing when vanilla bootstrap resampling fails. Consider a cold-start advertising problem where we would like to maximize the click-through rate (CTR). Consider two candidate ads Ad_A and Ad_B, whose groundtruth CTRs are 0.02 and 0.03, respectively. Suppose that both ads have already been impressed 50 times, Ad_A has got 1 click, and Ad_B has got no click. In this case, unfortunately, the bootstrap estimator for the CTR of Ad_B will always be naught in both mean and variance. Although the estimation perfectly fits the collected data, the induced policy would always recommend Ad_A, which is obviously sub-optimal.

To fix this issue, intuitively we should introduce a “prior knowledge” to the agent, e.g., “the CTR for an ad is not likely to be zero”. Then the agent can explore unfamiliar actions following the prior when there is not enough data, and exploit greedily with the learned reward models when there is sufficient data for all candidate actions. Conventionally, parametric Bayesian approaches add the prior distribution to model parameters as a regularization in the training process. In contrast, we achieve this by introducing a prior on the data rather than on the parameters, and the derived algorithm can be successfully incorporated with non-parametric bootstrap techniques.

The data prior is achieved by generating *fake samples* during training to regularize the optimization process. The generation procedure of fake samples is as follows. At step t , if the agent selects the action a_t with input features $x_t \triangleq x_{t,a_t}$ and receives reward r_t , it collects a real sample (x_t, r_t) . Then, we generate fake samples $(x_t, 0)$ and $(x_t, 1)$, which have the same input to the real sample but fake rewards. Next, when training the reward model, each of the two fake samples are getting used at probability $0 < g(x_t) \leq 1$, and are leaved unused otherwise.

By adding the randomized fake samples, the sample distribution can be smoothed, so it is less likely that the agent would be over-confident on its prediction. We call the probability of generating the fake samples $g(\cdot) : \mathcal{X} \rightarrow (0, 1]$ a *guidance function* or a *guidance*,

Algorithm 1 GuideBoot (Experience Replay)

Require: Number of models K , bootstrap batchsize b .

Require: Parameter α for the guidance $g(\cdot)$ in Eq. (5).

```

1: Initialize Replay Buffer  $\mathcal{H} = \emptyset$ ;
2: Initialize  $K$  models  $f^{(1)}, \dots, f^{(K)}$  independently;
3: Initialize the density function  $\hat{\rho}$ .
4: for  $t = 1, 2, \dots$  do
    // Part 1. Decision-making
5:   Observe  $m$  actions  $x_{t,1}, \dots, x_{t,m}$ .
6:   Choose  $k \sim \mathcal{U}\{1, K\}$  as the model index.
7:   Estimate  $\hat{y}_{t,a} \leftarrow f^{(k)}(x_{t,a})$  for all  $a = 1, \dots, m$ .
8:   Take action  $a_t \leftarrow \arg \max_a \hat{y}_{t,a}$  and receive reward  $r_t$ .
9:   Append the history  $\mathcal{H} \leftarrow \mathcal{H} \cup \{(x_{t,a_t}, r_t)\}$ .
    // Part 2. Prepare bootstrap data
10:  Update the density function  $\hat{\rho}$  with  $(x_{t,a_t}, r_t)$ .
11:  for  $k = 1, \dots, K$  do
12:     $\mathcal{B}_k \leftarrow$  resample a batch with size  $b$  from  $\mathcal{H}$ .
13:     $\tilde{\mathcal{B}}_k \leftarrow \mathcal{B}_k$ 
14:    for each record  $(x_j, r_j)$  in  $\mathcal{B}_k$  do
15:       $g_j \leftarrow \min\{\alpha/\hat{\rho}(x_j), 1\}$ 
16:      With probability  $g_j$ :
17:         $\tilde{\mathcal{B}}_k \leftarrow \tilde{\mathcal{B}}_k \cup \{(x_j, 1)\}$ 
18:      With probability  $g_j$ :
19:         $\tilde{\mathcal{B}}_k \leftarrow \tilde{\mathcal{B}}_k \cup \{(x_j, 0)\}$ 
20:    end for
21:  end for
    // Part 3. Update the models
22:  for  $k = 1, \dots, K$  do
23:    Update model  $f^{(k)}$  with gradient descent on  $\tilde{\mathcal{B}}_k$ 
24:  end for

```

which is essentially the salience weight on the data prior given the context x_t . Intuitively, the guidance function represents the uncertainty about the input: the more the agent is familiar with the action, the smaller the probability is given by the guidance function. This can be achieved by using simple count-based methods or sophisticated Bayesian methods. A formal probabilistic explanation as well as the specific choices of the guidance function for real applications are detailed in Section 3.4.

Let’s come back to the example in the beginning of this section. By using the data prior, if we assign a positive guidance function to generate fake samples for both Ad_A and Ad_B, the bootstrapping estimator by resampling from the data for Ad_B can be non-zero with a large probability, so there is always a chance for Ad_B being impressed.

3.3 Training with resampled real samples and generated fake samples

The training of GuideBoot combines vanilla bootstrap resampling from real data and generated fake samples.

We need to train K models (e.g., K feed-forward neural networks) over an experience buffer $\mathcal{H}_{1:t} = \{(x_\tau, r_\tau)\}_{\tau=1}^t$ of all the collected samples up to time step t . During training, all the models are randomly initialized and then updated in parallel with mini-batch

gradient descent. For an update of the k -th model, we first follow the Bag of Little Bootstraps (BLB, [18]) to perform a b -out-of- t bootstrap from $\mathcal{H}_{1:t}$ (i.i.d. resampling with replacements) and get a subset $\mathcal{B}_k$, where b is a hyper-parameter for the size of $\mathcal{B}_k$. Next, for each sample $x \in \mathcal{B}_k$, we compute the guidance function $g(x)$ and generate fake samples $(x, 0)$ or $(x, 1)$ at probabilities $g(x)$. So we can get a mixed set of samples $\tilde{\mathcal{B}}_k$ as the union of $\mathcal{B}_k$ and the set of generated fake samples. Finally, we perform a step of gradient descent to minimize the negative log-likelihood over the samples $\tilde{\mathcal{B}}_k$ to update the model.

The pseudo code of GuideBoot is shown in algorithm 1. As different models are trained with different set of samples, the bootstrap resampling, fake data generation, and gradient descent updates for the bag of models can be done in parallel.

3.4 GuideBoot as a Bayesian approximation

It is known in the statistics literature that “the bootstrap distribution represents an (approximate) nonparametric, *noninformative* posterior” (see chapter 8.4 of [14]), which is also used in [25] to derive a bootstrap-based multi-armed bandit. In this part, we put forward a novel result that our GuideBoot can be viewed as an approximation of an *informative* posterior due to the explicit guidance.

To specifically show the connection between GuideBoot and Bayesian Inference for our bandit task, let’s first consider a simplified multi-armed bandit problem, where we have collected a set of feedbacks $r_{a,1}, \dots, r_{a,n} \sim \text{Bernoulli}(y_a)$ for an action a , with $0 \leq n_s \leq n$ successes out of n . In this case, if we use Bayesian inference to estimate y_a by $\hat{y}_a$ with a symmetric prior distribution of $\text{Beta}(\alpha, \alpha)$, the posterior distribution of $\hat{y}_a$ is

$$\hat{y}_a \mid \{r_{a,1}, \dots, r_{a,n}\} \sim \text{Beta}(\alpha + n_s, \alpha + n - n_s), \quad (2)$$

which is a biased estimator of y_a with a mean of $\frac{\alpha + n_s}{2\alpha + n}$ and a variance of $\frac{(\alpha + n_s)(\alpha + n - n_s)}{(2\alpha + n)^2(2\alpha + n + 1)}$.

On the other hand, the bootstrap mean estimator obtained by resampling n i.i.d. samples with replacement from this data follows a binomial distribution

$$n\hat{y}_a^b \sim \text{Bi}(n, n_s/n) \quad (3)$$

where $\text{Bi}(\cdot, \cdot)$ denotes the binomial distribution. So $\hat{y}_a^b$ is an unbiased estimator of y_a with a mean of n_s/n and a variance of $n_s(n - n_s)/n^3$. So the mean and variance of the posteriors of $\hat{y}_a$ and $\hat{y}_a^b$ are very similar.

Notice that the bootstrap posterior can be viewed as an approximate Bayesian posterior with a noninformative prior $\alpha \rightarrow 0$. However, in bandit algorithms, an informative prior is essential, e.g., classic Thompson Sampling for Bernoulli multi-armed bandits often uses a prior of $\text{Beta}(1, 1)$ for each arm. So now the problem is, how do we modify the bootstrap algorithm to approximate a Bayesian posterior with an informative prior?

The solution is GuideBoot. In the above case, we can achieve this by adding two fake feedbacks 0 and 1 to the data and then run bootstrap resampling. Specifically, the bootstrap weights on the fake rewards are set as α . The resulted estimator $\hat{y}_a^g$ (superscript g for GuideBoot) has a mean of $\frac{\alpha + n_s}{2\alpha + n}$ which is the same as $\hat{y}_a$, and a variance of $\frac{(\alpha + n_s)(\alpha + n - n_s)}{(2\alpha + n)^3}$ which is asymptotically the same as that of $\hat{y}_a$.

Further, as bootstrap resampling is with replacement, we show that $\hat{y}_a^g$ can also be constructed by adding fake samples *after* a standard bootstrap resampling procedure. That is, for each resampled data point, we add the fake samples with rewards 0 or 1, both at a probability α/n . Since n is the number of impression counts of action a , we can rewrite this probability as

$$g(a) = \alpha / \text{Count}(a), \quad (4)$$

which is the guidance function of GuideBoot when applying to non-contextual multi-armed bandits. With the analysis above, such an algorithm can be viewed as a nonparametric approximation to the Bayesian Thompson Sampling with Beta prior. Note that moving the fake data generation procedure after resampling is essential in our algorithm, which enables experience replay without adding fake samples to the buffer. It also enables online training as described in Section 4.

Finally, the methodology described above can be extended to contextual bandits. When each sample has a featured input x and the reward models are deep neural networks, it is non-trivial to use Bayesian methods because neither the prior nor the posterior of the neural nets can be obtained easily. However, our nonparametric GuideBoot is still simple to run. Recall that the $\text{count}(a)$ in Eq. (4) can be viewed as an unnormalized density estimation of action a . With the same intuition, the guidance function for contextual bandit with input x is formulated as a measure of uncertainty which is inversely proportional to an unnormalized density estimator,

$$g(x) = \alpha / \hat{\rho}(x). \quad (5)$$

where $\hat{\rho}(x)$ denotes an unnormalized density estimation for input x in the experience buffer. In particular, we assume that $\hat{\rho}(x) \geq 1$ for any seen input x so that the value of $g(x)$ is strictly bounded.

Now we recommend several specific choices of $\hat{\rho}(\cdot)$ as well as the guidance function. From the analysis above, one can see that the mentioned count-based guidance function in Eq. (4) becomes a special case of Eq. (5) by letting $\hat{\rho}(x_a) = \text{Count}(a)$, where a can be a categorical index inside the context vector x_a . For example, for cold-start advertising where we would like to promote exploration on small ads or new ads, we can simply let

$$\hat{\rho}(x_a) = \text{Count}(\text{ad_ID}), \quad (6)$$

where $\text{Count}(\text{ad_ID})$ is the count of historical impressions of this ad. For inputs that contain multiple categorical fields, one can alternatively use an unscaled harmonic average of counts, i.e.,

$$\hat{\rho}(x) = \left[\sum_{j=1}^J \text{Count}^{-1}(x_j) \right]^{-1}, \quad (7)$$

with J the number of categorical fields and x_j the j -th entry of x . It has shown in previous study on Thompson Sampling [29] that it works as an Precision Diagonal posterior approximation of Bayesian linear regression which yields good performance in a range of tasks.

4 ONLINE GUIDEBOOT

In real applications such as advertising and recommendation, the algorithm has to respond to large number of concurrent requests and update the models with streaming data in real time. The large amount of data and the demand of real-time response makes the

algorithm impossible to do bootstrap resampling from the complete history buffer. Even solely storing the experience buffer can be costly in many applications. To remedy the challenges in online scenario, we propose *Online GuideBoot*, an simple and efficient version of GuideBoot that is more practical for online settings.

4.1 GuideBoot in the online setting

To better describe the online environment, we give a typical outline of the online bandit using greedy policy in Algorithm 2. The setting of the online learning environment with streaming data is described as follows. We suppose that the online agent always collects a batch of samples before updating itself, which is usual in many real systems with concurrent queries. We denote the collected online mini-batch of c samples by $\mathcal{H}_{\text{on}} \triangleq \mathcal{H}_{t-c:t}$, which should be used to update the reward models. To maximize generality, here the batch-size c can be a varied or fixed number, which depends on the actual design of the online serving. For example, if it is a constant (e.g., $c = 512$), it means that the agent updates its reward model only every c time steps. Likely, c can be a variable, e.g., be the number of online samples received within the last 10 seconds.

To use bootstrap techniques in such an online setting when there is no complete history data, one can solely perform resampling on the recent batch $\mathcal{H}_{\text{on}}$, which might degrade the performance. However, fortunately, our GuideBoot algorithm can still work well by constructing a randomized dataflow with online streaming data. Although the dataflow we propose does not completely follow the conventional bootstrap resampling, it is strong in empirical performance and easy to implement.

Recall that the reason of using multiple models trained on different resampling datasets is to keep track of the model uncertainty in the predicted rewards. Considering that the reward models (i.e., logistic regression or neural networks) are usually updated by stochastic gradient descent (SGD) during real online learning, we can directly get stochastic updates on the models by randomizing the sequence of arriving samples. Therefore, we propose to use shuffling and minibatching instead of resampling to construct the batches of real data. After having the subsets of real data, we use the same process of generating fake samples as in the previous section, which is done independently for each model. Therefore, different models can still receive different fake samples while training.

The pseudo code of the derived online algorithm, named Online GuideBoot, is shown as Algorithm 3. In particular, after collecting an online buffer $\mathcal{H}_{\text{on}}$ with c samples, we construct the dataflow for updating each model as follows: First, the samples in the buffer is shuffled and split into n disjoint mini-batches, where n is a pre-defined number of mini-batches. Second, for each mini-batch, we generate fake samples with the same principles described in the previous section and add them into the mini-batches. Finally, the models are updated by gradient descent on these mini-batches one after another. Therefore, the gradients to update each model is computed over different mini-batches with randomly generated fake samples, which enables stochastic optimization.

4.2 Implementation and deployment

Our algorithm can be implemented easily in industrial platforms such as advertising systems [5, 15, 23] or modern industrial RL

Algorithm 2 Conventional greedy algorithm (online)

```

1: Initialize the model  $f$ .
2: while True do
3:   Clear online buffer  $\mathcal{H}_{\text{on}} \leftarrow \emptyset$ 
   // Part 1. Serving
4:   for  $t = 1, \dots, c$  do
5:     Observe  $m$  actions  $x_{t,1}, \dots, x_{t,m}$ 
6:     Estimate  $\hat{y}_{t,a} \leftarrow f(x_{t,a})$  for all  $a = 1, \dots, m$ .
7:     Take action  $a_t \leftarrow \arg \max_a \hat{y}_{t,a}$  and receive reward  $r_t$ .
8:     Append the buffer  $\mathcal{H}_{\text{on}} \leftarrow \mathcal{H}_{\text{on}} \cup \{(x_{t,a_t}, r_t)\}$ .
9:   end for
   // Part 2. Learning
10:  Split  $\mathcal{H}_{\text{on}}$  into mini-batches  $\{\mathcal{B}_i\}_1^n$ 
11:  Update model  $f$  with gradient descent on  $\{\mathcal{B}_i\}_1^n$ 
12: end while

```

Algorithm 3 Online GuideBoot

```

1: Initialize  $K$  models  $f^{(1)}, \dots, f^{(K)}$  independently;
2: Initialize the density function  $\hat{p}$ .
3: while True do
4:   Clear Online Buffer  $\mathcal{H}_{\text{on}} \leftarrow \emptyset$ 
   // Part 1. Serving
5:   for  $t = 1, \dots, c$  do
6:     Observe  $m$  actions  $x_{t,1}, \dots, x_{t,m}$ .
7:     Choose  $k \sim \mathcal{U}\{1, K\}$  as the model index.
8:     Estimate  $\hat{y}_{t,a} \leftarrow f^{(k)}(x_{t,a})$  for all  $a = 1, \dots, m$ .
9:     Take action  $a_t \leftarrow \arg \max_a \hat{y}_{t,a}$  and receive reward  $r_t$ .
10:    Append the buffer  $\mathcal{H}_{\text{on}} \leftarrow \mathcal{H}_{\text{on}} \cup \{(x_{t,a_t}, r_t)\}$ .
11:  end for
   // Part 2. Learning
12:  Update the density function  $\hat{p}$  with  $\mathcal{H}_{\text{on}}$ .
13:  for  $k = 1, \dots, K$  do
14:    Shuffle and split  $\mathcal{H}_{\text{on}}$  into mini-batches  $\{\mathcal{B}_l\}_1^n$ .
15:    for each mini-batch  $\mathcal{B}_l$  do
16:      for each record  $(x_j, r_j)$  in  $\mathcal{B}_l$  do
17:         $g_j \leftarrow \min\{\alpha/\hat{p}(x_j), 1\}$ 
18:        With probability  $g_j$ :
19:           $\mathcal{B}_l \leftarrow \mathcal{B}_l \cup \{(x_j, 1)\}$ 
20:        With probability  $g_j$ :
21:           $\mathcal{B}_l \leftarrow \mathcal{B}_l \cup \{(x_j, 0)\}$ 
22:      end for
23:    end for
24:  end while

```

platforms [12], which is particularly helpful in the cold-start problems [28, 30]. An overview of the learning pipeline on an online advertising system is illustrated in Figure 2. The pipeline consists of two parts: online serving and online learning.

The serving section is responsible for the real-time decision-making for each coming request. The following steps are subsequently processed whenever the system receives a new request with an input context¹.

1. **Retrieval of the candidate set.** A set of candidate actions (i.e., ads or items) are provided according to the context (query), along with the input attributes and features.
2. **Inference.** Each candidate is evaluated by the reward model randomly selected from the bag of K models and gets a predicted score. The scores are usually the expected rewards, such as the eCPM (expect cost per thousand impressions), which need to be estimated according to the predicted click through rates (pCTR). The system will use a specific model to make the predictions on a backend server, which is chosen randomly from the backend cluster of models.
3. **Decision-making.** Send the top ranked ad to the user.

On the other hand, the learning section collects user feedback and updates the models. We consider logging a small online buffer with c steps so that we can run the GuideBoot algorithm.

1. **Logging.** The user's response (e.g., click or not), along with the input context and selected action, is collected and published to the message queue. Records are fetched from the message queue and stored to an online buffer, until there are enough data to train the models.
2. **Updating the uncertainty estimator.** As described in Section 3.4, the guidance function $g(\cdot)$ is inversely proportional to an unnormalized density model, so it can be understood as an uncertainty estimator. After observing the online buffer, the estimator should be updated. For example, when using the count-based guidance in Eq. (4), the system needs to maintain a counting table.
3. **GuideBoot dataflow.** The dataflow is detailed in Section 4.1 and line 12-20 in Algorithm 3. It constructs K series of mini-batches in parallel by shuffling and splitting the online buffer, and generating fake samples with the guidance function.
4. **Gradient-based training.** The learner updates the models with the samples from the GuideBoot dataflow. The models are pushed to the serving section only after the updates are finished. After training, the buffer and the dataflow are cleared and ready for receiving new samples.

5 EXPERIMENTS

In this section, we evaluate our GuideBoot on both synthetic task and large-scale real-world advertising environments. Specifically, we seek to answer the following questions through experiments:

- Q1.** In a standard non-deep Bernoulli setting, how does GuideBoot compare with classic UCB and Thompson Sampling algorithms?
- Q2.** In complex deep contextual bandits, how does GuideBoot compare with recent advanced deep methods?
- Q3.** What contributes to the improvement of GuideBoot?

¹We use an online advertising scenario as an example. The input context usually contains a query from a user, along with the text input of the query, the user's attributes, and the user's viewing history.

Figure 2: Overall pipeline of an online advertising system with Online GuideBoot.

5.1 Experiments on synthetic task

5.1.1 Settings. Although the original motivation of this paper is to build practical algorithms for deep complex contextual bandits, we would like to first check the correctness and effectiveness of GuideBoot in the classic setting. Therefore, we setup standard Bernoulli bandit environments where the reward functions are generalized linear models. In such a setting, it has been extensively studied that a number of UCB and Thompson Sampling based algorithms are provably efficient with sub-linear cumulative regret bounds [8].

Now we describe the detailed setup of the Bernoulli bandits. At each step, we observe $m = 25$ actions, each with three categorical attributes $x_a = [a, x_{a,1}, x_{a,2}]$, where $a = 1, \dots, 25$, $x_{a,1} \sim \mathcal{U}\{1, 5\}$, and $x_{a,2} \sim \mathcal{U}\{1, 5\}$. For each input x_a , the reward is drawn from Bernoulli(y_a), where the oracle expected reward is $y_a = \sigma(\mathbf{w}_0[a] + \mathbf{w}_1[x_{a,1}] + \mathbf{w}_2[x_{a,2}] - 1)$ with σ the sigmoid function, $\mathbf{w}_0 \in \mathbb{R}^{25}$, $\mathbf{w}_1 \in \mathbb{R}^5$, and $\mathbf{w}_2 \in \mathbb{R}^5$ the coefficient vectors for the three fields. Here we use the notation $\mathbf{w}[i]$ denoting the i -th entry of a vector $\mathbf{w}$. We generate 100 randomized environments of this setting, where the coefficients in $\mathbf{w}_0$, $\mathbf{w}_1$ and $\mathbf{w}_2$ in each environment are i.i.d. drawn from $\mathcal{U}(-0.25, 0.25)$.

5.1.2 Compared methods. As the problem is standard Bernoulli contextual bandits, we compare our GuideBoot with the following classic algorithms.

ϵ -greedy: The simplest unguided algorithm, which selects the action with the greatest expected reward with probability $1 - \epsilon$ and explores randomly otherwise. We set $\epsilon = 0.1$ which is a common choices in previous studies.

ϵ -greedy (decay): A variant of ϵ -greedy with decaying ϵ as t increases. Although as simple as the ordinary ϵ -greedy, it is commonly used as an empirically strong baseline method. We let ϵ linearly decay from 0.1 to zero throughout training.

GLM-UCB: The UCB algorithm with generalized linear model, proposed in [8]. It relies on estimating the upper confidence bound of the reward prediction with a Bayesian generalized linear model. The coefficient on the exploration bonus is set as $\alpha_t = \sqrt{\log(t + 1)}$ as suggested in its original paper.

Figure 3: Experimental results on the synthetic tasks with 50 random runs. The average regret over time is reported.

Table 2: Statistics of the Tencent Ads simulator

Dataset	Tencent Ads A	Tencent Ads B
# time steps	4.5M	14.7M
# ads (total)	8.9K	6.5K
# candidates (at each step)	250 ~ 450	100 ~ 200
# fields of context	13	13
# fields of candidate	12	12
# features of context	4.2K	4.3K
# features of ad	12K	9.4K

TS-BLR: Thompson Sampling with Bayesian generalized linear model [1, 4], which is commonly used and is shown to achieve state-of-the-art performance in a recent study [29]. The reward model is learned with the same generalized linear model in GLM-UCB.

Bootstrap: The standard bootstrap resampling method with experience replay, which maintains K reward models trained with resampled data from the history buffer.

OBB: Online Bootstrap Bandit algorithm proposed in [31], which enables online bootstrap with duplicated samples where the number of duplicates is drawn from a Poisson distribution.

Giro: A recent bootstrap-based method by Kveton et al. [20], which trains K models on data resampled from a perturbed history buffer. When each new sample is added to the buffer, a pair of pseudo samples are also added with probability $\alpha = 0.5$.

We test both the experience replay version and the online version of GuideBoot (denoted by **GuideBoot** and **Online Guideboot**, respectively), with $\alpha = 1$ and $\hat{p}(\cdot)$ in Eq. (7). For all the bootstrap-based methods, we set K to 5, which is a sufficient number for bootstrap models as demonstrated in previous work [20, 31]. Among the bootstrap-based methods, Vanilla Bootstrap, Giro, and GuideBoot require storing the whole history, while OBB and Online GuideBoot are online methods without the whole experience buffer.

5.1.3 Results. The averaged regret is shown as Figure 3. From the results, we observe that our GuideBoot works pretty good. Specifically, the experience replay version of GuideBoot performs the best and significantly outperforms the classic Bayesian methods as well as previous bootstrap-based algorithms.

To analyse the results of bootstrap-based algorithms in more detail, we notice that the vanilla bootstrap runs poorly even though it can resample from the whole history. Giro adds a fixed proportion of fake samples to the history and outperforms vanilla bootstrap.

Next, we find that GuideBoot brings a significant improvement in performance, indicating that the fake sample generation along with the guidance function is essential for the algorithm. By comparing GuideBoot with vanilla bootstrap and Giro, it is clear that the uncertainty-based guidance helps, which answers the question Q3.

On the other hand, for the bootstrap-based methods learned with streaming data only, OBB as the online extension of the vanilla bootstrap suffers from a huge loss. However, our Online GuideBoot still works fairly well and its relative loss from GuideBoot is small, indicating the effectiveness of the shuffle-and-split step. Online GuideBoot can be a competitive baseline method when it is difficult to access to the whole history.

5.2 Experiments on online advertising

Next we conduct experiments on real-world online advertising environments, which involves high-dimensional featured inputs and non-stationary user responses with complex dependencies on the features. A large-scale simulator based on historical logs of online advertising platforms is developed for fair comparison due to the infeasibility to deploy online A/B test.

5.2.1 Advertising simulation setup. We collected five days of advertising logs from Tencent’s online advertising systems. Specifically, we subsampled two logging datasets (written as **Tencent Ads A** and **Tencent Ads B**) from two major sites of Tencent display ads and feed ads. The statistics of the datasets are outlined in Table 2.

In the logs, each record is a real request consisting of two parts: the context of the request and a candidate pool. The context includes query and user features such as the user’s gender, age, and device type. The candidate pool is a batch of ads with their attributes (e.g., ad ID and advertiser ID). The simulator has the click-through rate (CTR) of each (context, ad) pair given by a sophisticated black-box user model, which is unknown to the agent. When the agent chooses an ad, the simulator samples a binary feedback using the user model. The goal of the agent is to achieve a high average reward throughout learning.

For each simulator, we construct two sets of experiments: one using experience replay where the model is updated with mini-batches sampled from the whole history, and another using online streaming data without storing all the history. We run all the experiments over 10 random seeds and report the average reward.

5.2.2 Model structure and compared methods. The reward function is modeled by a neural network with an embedding layer and two hidden layers. The embedding layer transforms each raw field into a 8-dimensional dense vector. The embeddings of all the input fields are concatenated into a single vector, which is then fed into a ReLU

Figure 4: The boxplot of the experimental results on advertising tasks with 10 random runs.

network with two 128-dimensional fully connected layers. The neural net is updated by the Adam optimizer [16] to minimize the Log-Loss (aka negative log-likelihood).

The tested methods are listed as follows. In addition to ϵ -greedy, **Bootstrap**, **OBB** and **Giro** algorithms described in Section 5.1.2, we also compare the following Bayesian methods that can be incorporated with deep models:

Deep-UCB1: A common method in the industry that combines deep models with UCB1 [2]. The uncertainty is estimated from the impression count $\alpha \sqrt{2 \log(t) / \text{Count}(\text{ad_ID})}$. We set an empirical choice of $\alpha = 0.1$.

Deep-TS-Beta: Another practical method in the industry that combines Thompson sampling with count-based Beta distributions. The prediction for each input is estimated with $\text{Beta}(a, b)$ s.t. $a + b = \text{Count}(\text{ad_ID})$ and $a/(a + b) = f(x_a)$. Following [4], we use posterior shaping of $\alpha = 0.25$.

MC-Dropout: To use Monte-Carlo Dropout during both training and inference, which can be understood as variational Bayes [9, 17]. Using dropout for decision-making can be seen as an approximation of Thompson Sampling [9, 26]. The Dropout is applied only on the last hidden layer [29] and the Dropout rate is set to 0.1.

For bootstrap-based methods, we again used 5 rewards models. For Giro, we found that large α led to unacceptable performance in the advertising tasks, so we searched over $[0.2, 0.1, 0.05, 0.01]$ and chose $\alpha = 0.01$ in this experiment. For our GuideBoot and Online GuideBoot, in order to ensure fair comparison with the count-based Deep-UCB1 and Deep-TS-Beta, we use the count-based density estimator in Eq. (4) to construct the guidance function. Note that vanilla Bootstrap and Giro uses the whole history, so they are only tested in the experience replay group. OBB uses online streaming data, so it is tested in the online group.

5.2.3 Results. The comparisons are illustrated in Figure 4, and the detailed results are listed in Table 3 and Table 4, respectively.

From the results, we first observe that our GuideBoot and Online GuideBoot have significant and consistent superiority in average cumulative reward comparing to all other baselines. Meanwhile, the

Table 3: Results on Tencent Ads (Experience Replay)

Method	Tencent Ads A		Tencent Ads B	
	reward	std	reward	std
ϵ -greedy ($\epsilon = 0.1$)	0.0248	(0.0010)	0.0225	(0.0003)
ϵ -greedy (decay)	0.0260	(0.0009)	0.0231	(0.0006)
Deep-TS-Beta	0.0262	(0.0007)	0.0238	(0.0002)
Deep-UCB1	0.0222	(0.0005)	0.0220	(0.0002)
MC Dropout	0.0245	(0.0010)	0.0221	(0.0006)
Bootstrap	0.0264	(0.0010)	0.0239	(0.0003)
Giro	0.0249	(0.0012)	0.0230	(0.0003)
GuideBoot	0.0280	(0.0006)	0.0248	(0.0003)

Table 4: Results on Tencent Ads (Online)

Method	Tencent Ads A		Tencent Ads B	
	reward	std	reward	std
ϵ -greedy ($\epsilon = 0.1$)	0.0289	(0.0006)	0.0254	(0.0002)
ϵ -greedy (decay)	0.0293	(0.0012)	0.0261	(0.0002)
Deep-TS-Beta	0.0295	(0.0004)	0.0266	(0.0001)
Deep-UCB1	0.0243	(0.0004)	0.0237	(0.0001)
MC-Dropout	0.0297	(0.0010)	0.0258	(0.0002)
OBB	0.0302	(0.0009)	0.0262	(0.0002)
Online GuideBoot	0.0317	(0.0004)	0.0270	(0.0002)

standard deviations of GuideBoots are relatively small, indicating that the improvement of our method is robust in all the experiments.

To provide more insights from the experiments, we find some simple methods are considerably effective, while some sophisticated methods do not always work well. For example, despite its simplicity, the ϵ -greedy algorithm with decaying ϵ is shown to be strong and can often outperform MC-Dropout. For the Bayesian methods, Deep-UCB1 is much worse than Deep-TS-Beta, most likely because it requires a gradually increasing coefficient for exploration (i.e., the $\sqrt{\log(t)}$ factor) which may cause over-exploration.

Additionally, in these two simulations, the online algorithms have higher average click-through rate comparing to the experience replay algorithms, which is in contrast with the previous experiments on synthetic data. It is most likely because that the synthetic task has a stationary environment, while in advertising the distributions of inputs and outputs shift over time. Therefore, the online algorithms that updates the models only from recent samples are capable of fast adaptation to the distributional shift.

Overall, from the experiments on both the synthetic task and complex deep contextual bandits, we find that our GuideBoot and Online GuideBoot methods enjoy both efficiency and effectiveness.

6 CONCLUSION AND DISCUSSION

In this paper, we propose GuideBoot, an efficient and effective algorithm for addressing the trade-off between exploration and exploitation in deep contextual bandits. GuideBoot trains multiple reward models with resampled subsets of data and guides the exploration behavior by adding fake samples to the resampled subsets. We show that when the probability of generating fake samples is in proportion to the predictive uncertainty of the reward, GuideBoot can be understood as a non-Bayesian approximation to Thompson Sampling with an informative prior. For online applications where resampling from the whole history is infeasible, we propose an online version of GuideBoot that solely learns from the online streaming data, which is efficient and easy to implement in modern platforms. Experiments over synthetic tasks and real-world large-scale advertising show that our method consistently and significantly outperforms previous state-of-the-art methods.

ACKNOWLEDGMENTS

This work is supported by 2019 Tencent Marketing Solution Rhino-Bird Focused Research Program. The research work is also supported by the National Key Research and Development Program of China under Grant No. 2017YFB1002104, the National Natural Science Foundation of China under Grant No. 92046003, 61976204, U1811461, the Project of Youth Innovation Promotion Association CAS and Beijing Nova Program Z201100006820062.

REFERENCES

- [1] Shipra Agrawal and Navin Goyal. 2013. Thompson sampling for contextual bandits with linear payoffs. In *ICML*. 127–135.
- [2] Peter Auer, Nicolò Cesa-Bianchi, and Paul Fischer. 2002. Finite-time analysis of the multiarmed bandit problem. *Machine learning* 47, 2-3 (2002), 235–256.
- [3] Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. 2015. Weight Uncertainty in Neural Network. In *International Conference on Machine Learning*. 1613–1622.
- [4] Olivier Chapelle and Lihong Li. 2011. An empirical evaluation of thompson sampling. In *Advances in neural information processing systems*. 2249–2257.
- [5] Heng-Tze Cheng, Levent Koc, Jeremiah Harmsen, Tal Shaked, Tushar Chandra, Hrishi Aradhye, Glen Anderson, Greg Corrado, Wei Chai, Mustafa Ipsir, et al. 2016. Wide & deep learning for recommender systems. In *Proceedings of the 1st workshop on deep learning for recommender systems*. 7–10.
- [6] Dean Eckles and Maurits Kaptein. 2014. Thompson sampling with the online bootstrap. *arXiv preprint arXiv:1410.4009* (2014).
- [7] Adam N Elmachet, Ryan McNellis, Sechan Oh, and Marek Petrik. 2017. A Practical Method for Solving Contextual Bandit Problems Using Decision Trees. In *UAI*.
- [8] Sarah Filippi, Olivier Cappé, Aurélien Garivier, and Csaba Szepesvári. 2010. Parametric bandits: The generalized linear case. In *NIPS*. 586–594.
- [9] Yarín Gal and Zoubin Ghahramani. 2016. Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*. 1050–1059.
- [10] Yarín Gal, Jiri Hron, and Alex Kendall. 2017. Concrete dropout. In *Advances in neural information processing systems*. 3581–3590.
- [11] Marta Garnelo, Jonathan Schwarz, Dan Rosenbaum, Fabio Viola, Danilo J Rezende, SM Eslami, and Yee Whye Teh. 2018. Neural processes. *arXiv preprint arXiv:1807.01622* (2018).
- [12] Jason Gauci, Edoardo Conti, Yitao Liang, Kittipat Virochsiri, Yuchen He, Zachary Kaden, Vivek Narayanan, Xiaohui Ye, Zhengxing Chen, and Scott Fujimoto. 2018. Horizon: Facebook’s open source applied reinforcement learning platform. *arXiv preprint arXiv:1811.00260* (2018).
- [13] Botao Hao, Yasin Abbasi-Yadkori, Zheng Wen, and Guang Cheng. 2019. Bootstrapping Upper Confidence Bound. In *NeurIPS*. 12123–12133.
- [14] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. 2009. *The elements of statistical learning: data mining, inference, and prediction*. Springer Science & Business Media.
- [15] Xinran He, Junfeng Pan, Ou Jin, Tianbing Xu, Bo Liu, Tao Xu, Yanxin Shi, Antoine Atallah, Ralf Herbrich, Stuart Bowers, et al. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [16] Durk P Kingma, Tim Salimans, and Max Welling. 2015. Variational dropout and the local reparameterization trick. In *Advances in neural information processing systems*. 2575–2583.
- [17] Ariel Kleiner, Ameet Talwalkar, Purnamrita Sarkar, and Michael I Jordan. 2014. A scalable bootstrap for massive data. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 76, 4 (2014), 795–816.
- [18] Andreas Krause and Cheng S Ong. 2011. Contextual gaussian process bandit optimization. In *Advances in Neural Information Processing Systems*. 2447–2455.
- [19] Branislav Kveton, Csaba Szepesvári, Sharan Vaswani, Zheng Wen, Tor Lattimore, and Mohammad Ghavamzadeh. 2019. Garbage In, Reward Out: Bootstrapping Exploration in Multi-Armed Bandits. In *ICML*. 3601–3610.
- [20] John Langford and Tong Zhang. 2008. The epoch-greedy algorithm for multi-armed bandits with side information. In *NIPS*. 817–824.
- [21] Lihong Li, Wei Chu, John Langford, and Robert E Schapire. 2010. A contextual-bandit approach to personalized news article recommendation. In *WWW*. 661–670.
- [22] H Brendan McMahan, Gary Holt, David Sculley, Michael Young, Dietmar Ebner, Julian Grady, Lan Nie, Todd Phillips, Eugene Davydov, Daniel Golovin, et al. 2013. Ad click prediction: a view from the trenches. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*. 1222–1230.
- [23] Ian Osband, Charles Blundell, Alexander Pritzel, and Benjamin Van Roy. 2016. Deep exploration via bootstrapped DQN. In *Advances in neural information processing systems*. 4026–4034.
- [24] Ian Osband and Benjamin Van Roy. 2015. Bootstrapped thompson sampling and deep exploration. *arXiv preprint arXiv:1507.00300* (2015).
- [25] Feiyang Pan, Qingpeng Cai, Pingzhong Tang, Fuzhen Zhuang, and Qing He. 2019. Policy Gradients for Contextual Recommendations. In *The World Wide Web Conference* (San Francisco, CA, USA) (WWW ’19). ACM, New York, NY, USA, 1421–1431.
- [26] Feiyang Pan, Jia He, Dandan Tu, and Qing He. 2020. Trust the Model When It Is Confident: Masked Model-based Actor-Critic. In *Advances in neural information processing systems*.
- [27] Feiyang Pan, Shuokai Li, Xiang Ao, Pingzhong Tang, and Qing He. 2019. Warm Up Cold-start Advertisements: Improving CTR Predictions via Learning to Learn ID Embeddings. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 695–704.
- [28] Carlos Riquelme, George Tucker, and Jasper Snoek. 2018. Deep bayesian bandits showdown: An empirical comparison of bayesian deep networks for thompson sampling. *arXiv preprint arXiv:1802.09127* (2018).
- [29] Parikshit Shah, Ming Yang, Sachidanand Alle, Adwait Ratnaparkhi, Ben Shahshani, and Rohit Chandra. 2017. A practical exploration system for search advertising. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 1625–1631.
- [30] Liang Tang, Yexi Jiang, Lei Li, Chunqiu Zeng, and Tao Li. 2015. Personalized recommendation via parameter-free contextual bandits. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 323–332.
- [31] Sharan Vaswani, Branislav Kveton, Zheng Wen, Anup Rao, Mark Schmidt, and Yasin Abbasi-Yadkori. 2018. New insights into bootstrapping for bandits. *arXiv preprint arXiv:1805.09793* (2018).
- [32] Chi-Hua Wang, Yang Yu, Botao Hao, and Guang Cheng. 2020. Residual Bootstrap Exploration for Bandit Algorithms. *arXiv preprint arXiv:2002.08436* (2020).