
Online Posterior Sampling with a Diffusion Prior

Branislav Kveton **Boris N. Oreshkin** **Youngsuk Park** **Aniket Deshmukh** **Rui Song**
Adobe Research* Amazon AWS AI Labs AWS AI Labs Amazon

Abstract

Posterior sampling in contextual bandits with a Gaussian prior can be implemented exactly or approximately using the Laplace approximation. The Gaussian prior is computationally efficient but it cannot describe complex distributions. In this work, we propose approximate posterior sampling algorithms for contextual bandits with a diffusion model prior. The key idea is to sample from a chain of approximate conditional posteriors, one for each stage of the reverse diffusion process, which are obtained by the Laplace approximation. Our approximations are motivated by posterior sampling with a Gaussian prior, and inherit its simplicity and efficiency. They are asymptotically consistent and perform well empirically on a variety of contextual bandit problems.

1 Introduction

A *multi-armed bandit* [27, 6, 30] is an online learning problem where an agent sequentially interacts with an environment over n rounds with the goal of maximizing its rewards. In each round, it takes an *action* and receives its *stochastic reward*. The mean rewards of the actions are unknown *a priori* and must be learned. This leads to the *exploration-exploitation dilemma*: *explore* actions to learn about them or *exploit* the action with the highest estimated reward. Bandits have been successfully applied to problems where uncertainty modeling and adaptation are beneficial, such recommender systems [32, 54, 25, 35] and hyper-parameter optimization [34].

Contextual bandits [29, 32] with linear [13, 1] and *generalized linear models (GLMs)* [17, 33, 2, 26] have become popular due to their flexibility and efficiency. The features in these models can be hand-crafted or learned from historic data [40], and the models can be also updated incrementally [1, 24]. While the original algorithms for linear and GLM bandits were based on *upper confidence bounds (UCBs)* [13, 1, 17], *Thompson sampling (TS)* is more popular in practice [11, 3, 42, 44]. The key idea in TS is to explore by sampling from the posterior distribution of model parameter θ_* . TS uses the prior knowledge about θ_* to speed up exploration [11, 40, 36, 9, 21, 20, 5]. When the prior is a multivariate Gaussian, the posterior of θ_* can be updated and sampled from efficiently [11]. This prior has a limited expressive power, because it cannot even represent multimodal distributions. To address this, we study posterior sampling with a diffusion prior. The main benefit of such priors is that they can represent complex distributions and be learned from data.

We make the following contributions. First, we propose novel posterior sampling approximations for linear models and GLMs with a diffusion model prior. The key idea is to sample from a chain of approximate conditional posteriors, one for each stage of the reverse process, which are estimated in a closed form. In linear models, each conditional is a product of two Gaussians, representing prior knowledge and diffused evidence (Theorem 2). In GLMs, each conditional is obtained by a Laplace approximation, which mixes prior knowledge and evidence (Theorem 4). Our approximations are motivated by posterior sampling with Gaussian priors, and inherit its simplicity and efficiency. Prior works (Section 7) sampled from the posterior using the likelihood score, and their approximations become unstable when the score is high. We combine the likelihood with conditional priors, in each

*The work was done at AWS AI Labs.

stage of the diffusion model, using the Laplace approximation. The resulting posterior concentrates at a single point and can be sampled from efficiently even if the likelihood score is high. We prove that this approximation is asymptotically consistent.

Our second contribution is in theory. We properly derive our posterior approximations (Theorems 2 and 4) and show their asymptotic consistency (Theorem 3). The key idea in the proof of Theorem 3 is that the conditional posteriors concentrate at a scaled unknown model parameter as the number of observations increases. While this claim is asymptotic, it is an expected property of a posterior distribution. Many prior works, such as Chung et al. [12], do not propose asymptotically consistent approximations. All of our main results rely on a novel approximation of clean samples by scaled diffused samples (Section 4.3). The most challenging part of the analysis is Theorem 3, where we analyze an asymptotic behavior of a chain of T dependent random vectors.

Our last contribution is an empirical evaluation on contextual bandits. We focus on bandits because the ability to represent all levels of uncertainty precisely is critical for exploration. Our experiments show that a score-based method fails to do so (Section 6.2). Note that our posterior approximations are general and not restricted to bandits.

2 Setting

We start with introducing our notation. Random variables are capitalized, except for Greek letters like θ . We denote the marginal and conditional probabilities under probability measure p by $p(X = x)$ and $p(X = x | Y = y)$, respectively. When the random variables are clear from context, we write $p(x)$ and $p(x | y)$. We denote by $X_{n:m}$ and $x_{n:m}$ a collection of random variables and their values, respectively. For a positive integer n , we define $[n] = \{1, \dots, n\}$. The indicator function is $\mathbb{1}\{\cdot\}$. The i -th entry of vector v is v_i . If the vector is already indexed, such as v_j , we write $v_{j,i}$. We denote the maximum and minimum eigenvalues of matrix $M \in \mathbb{R}^{d \times d}$ by $\lambda_1(M)$ and $\lambda_d(M)$, respectively.

The posterior sampling problem can be formalized as follows. Let $\theta_* \in \Theta$ be an unknown *model parameter* and $\Theta \subseteq \mathbb{R}^d$ be the space of model parameters. Let $h = \{(\phi_\ell, y_\ell)\}_{\ell \in [N]}$ be the *history* of N noisy observations of θ_* , where $\phi_\ell \in \mathbb{R}^d$ is the feature vector for $y_\ell \in \mathbb{R}$. We assume that

$$y_\ell = g(\phi_\ell^\top \theta_*) + \varepsilon_\ell, \quad (1)$$

where $g: \mathbb{R} \rightarrow \mathbb{R}$ is the *mean function* and ε_ℓ is an independent zero-mean σ^2 -sub-Gaussian noise for $\sigma > 0$. Let $p(h | \theta_*)$ be the *likelihood* of observations in history h under model parameter θ_* and $p(\theta_*)$ be its *prior probability*. By Bayes' rule, the posterior distribution of θ_* given h is

$$p(\theta_* | h) \propto p(h | \theta_*) p(\theta_*). \quad (2)$$

We want to sample from $p(\cdot | h)$ efficiently when the prior distribution is represented by a diffusion model. As a stepping stone, we review existing posterior formulas for multivariate Gaussian priors. This motivates our solution for diffusion model priors.

2.1 Linear Model

The posterior of θ_* in linear models can be derived as follows.

Assumption 1. Let g in (1) be an identity and $\varepsilon_\ell \sim \mathcal{N}(0, \sigma^2)$. Then the likelihood of h under model parameter θ_* is $p(h | \theta_*) \propto \exp[-\sum_{\ell=1}^N (y_\ell - \phi_\ell^\top \theta_*)^2 / (2\sigma^2)]$.

Let $p(\theta_*) = \mathcal{N}(\theta_*; \theta_0, \Sigma_0)$ be the prior distribution of θ_* , where $\theta_0 \in \mathbb{R}^d$ and $\Sigma_0 \in \mathbb{R}^{d \times d}$ are the prior mean and covariance, respectively. Then $p(\theta_* | h) \propto \mathcal{N}(\theta_*; \hat{\theta}, \hat{\Sigma})$ [10], where

$$\hat{\theta} = \hat{\Sigma} \left(\Sigma_0^{-1} \theta_0 + \sigma^{-2} \sum_{\ell=1}^N \phi_\ell y_\ell \right), \quad \hat{\Sigma} = \left(\Sigma_0^{-1} + \sigma^{-2} \sum_{\ell=1}^N \phi_\ell \phi_\ell^\top \right)^{-1},$$

are the posterior mean and covariance, respectively. In this work, we write them equivalently as

$$\hat{\theta} = \hat{\Sigma}(\Sigma_0^{-1} \theta_0 + \bar{\Sigma}^{-1} \bar{\theta}), \quad \hat{\Sigma} = (\Sigma_0^{-1} + \bar{\Sigma}^{-1})^{-1}, \quad (3)$$

where $\bar{\theta} = \sigma^{-2} \bar{\Sigma} \sum_{\ell=1}^N \phi_\ell y_\ell$ and $\bar{\Sigma}^{-1} = \sigma^{-2} \sum_{\ell=1}^N \phi_\ell \phi_\ell^\top$ are the empirical mean and inverse of its covariance, respectively. Therefore, the posterior of θ_* is a product of two multivariate Gaussians: $\mathcal{N}(\theta_0, \Sigma_0)$ representing prior knowledge about θ_* and $\mathcal{N}(\bar{\theta}, \bar{\Sigma})$ representing empirical evidence.

Algorithm 1 IRLS: Iteratively reweighted least squares.

- 1: **Input:** Prior parameters θ_0 and Σ_0 , history of observations $h = \{(\phi_\ell, y_\ell)\}_{\ell \in [N]}$
 - 2: Initialize $\hat{\theta} \in \mathbb{R}^d$
 - 3: **repeat**
 - 4: **for** stage $\ell = 1, \dots, N$ **do**
 - 5: $z_\ell \leftarrow \phi_\ell^\top \hat{\theta} + (y_\ell - g(\phi_\ell^\top \hat{\theta})) / \dot{g}(\phi_\ell^\top \hat{\theta})$
 - 6: $\hat{\Sigma} \leftarrow \left(\Sigma_0^{-1} + \sum_{\ell=1}^N \dot{g}(\phi_\ell^\top \hat{\theta}) \phi_\ell \phi_\ell^\top \right)^{-1}$
 - 7: $\hat{\theta} \leftarrow \hat{\Sigma} \left(\Sigma_0^{-1} \theta_0 + \sum_{\ell=1}^N \dot{g}(\phi_\ell^\top \hat{\theta}) \phi_\ell z_\ell \right)$
 - 8: **until** $\hat{\theta}$ converges
 - 9: **Output:** Posterior mean $\hat{\theta}$ and covariance $\hat{\Sigma}$
-

2.2 Generalized Linear Model

Generalized linear models (GLMs) [37] extend linear models (Section 2.1) to non-linear monotone mean functions g in (1). For instance, in logistic regression, $g(u) = 1/(1 + \exp[-u])$ is a sigmoid. The likelihood of observations in GLMs has the following form [26].

Assumption 2. Let $h = \{(\phi_\ell, y_\ell)\}_{\ell \in [N]}$ be a history of N observations under mean function g and the corresponding noise. Then $\log p(h | \theta_*) \propto \sum_{\ell=1}^N y_\ell \phi_\ell^\top \theta_* - b(\phi_\ell^\top \theta_*) + c(y_\ell)$, where c is a real function and b is a function whose derivative is the mean function, $\dot{b} = g$.

The posterior distribution of θ_* in GLMs does not have a closed form in general [10]. Therefore, it is often approximated by the *Laplace approximation*. Let the prior distribution of the model parameter be $p(\theta_*) = \mathcal{N}(\theta_*; \theta_0, \Sigma_0)$, as in Section 2.1. Then the Laplace approximation is $\mathcal{N}(\hat{\theta}, \hat{\Sigma})$, where $\hat{\theta}$ is the *maximum a posteriori (MAP) estimate* of θ_* and $\hat{\Sigma}$ is the corresponding covariance. Note that the Laplace approximation can be applied to non-Gaussian priors.

The MAP estimate $\hat{\theta}$ can be obtained by *iteratively reweighted least squares (IRLS)* [52], which we present in Algorithm 1. IRLS is a Newton-type algorithm that computes $\hat{\theta}$ iteratively (lines 6 and 7). It converges to the optimal solution due to the strong convexity of the problem. The solution has a similar structure to (3). That is, $\mathcal{N}(\hat{\theta}, \hat{\Sigma})$ is a product of two multivariate Gaussians, representing prior knowledge about θ_* and empirical evidence. The new quantities in GLMs are the derivative of the mean function \dot{g} and pseudo-observations z_ℓ (line 5), which play the role of observations y_ℓ in Section 2.1.

2.3 Towards Diffusion Model Priors

The assumption that $p(\theta_*) = \mathcal{N}(\theta_*; \theta_0, \Sigma_0)$ is limiting, for instance because it precludes multimodal priors. We relax it by representing $p(\theta_*)$ by a diffusion model, which we call a *diffusion model prior*. We propose efficient posterior sampling approximations for this prior, where the prior and empirical evidence are mixed similarly to (3) and IRLS. We review diffusion models next.

3 Diffusion Models

Diffusion models [46, 19] are generative models trained by diffusing samples from unknown and hard to represent distributions. They can be viewed in multiple ways [49]. We adopt the probabilistic formulation and presentation of Ho et al. [19]. A *diffusion model* is a graphical model with T stages indexed by $t \in [T]$. Each stage t is associated with a *latent variable* $S_t \in \mathbb{R}^d$. A *sample* from the model is represented by an *observed variable* $S_0 \in \mathbb{R}^d$. We visualize a diffusion model in Figure 1. In the *forward process*, a clean sample s_0 is diffused through a sequence of variables S_1, \dots, S_T . This process is used to learn the *reverse process*, where the clean sample s_0 is generated through a sequence of variables S_T, \dots, S_0 . To sample s_0 from the posterior (Section 4), we add a random variable H that represents partial information about s_0 . We introduce forward and reverse diffusion




Figure 1: Graphical models of the forward and reverse processes in the diffusion model. The variable H represents partial information about S_0 .

processes next. Learning of the reverse process is described in Appendix B. While this is a critical component of diffusion models, it is not necessary to introduce our posterior approximations.

Forward process. In the forward process, a clean sample s_0 is diffused through a chain of latent variables S_1, \dots, S_T (Figure 1). We denote the probability measure under this process by q and define its joint probability distribution as

$$q(s_{1:T} | s_0) = \prod_{t=1}^T q(s_t | s_{t-1}), \quad \forall t \in [T] : q(s_t | s_{t-1}) = \mathcal{N}(s_t; \sqrt{\alpha_t} s_{t-1}, \beta_t I_d), \quad (4)$$

where $q(s_t | s_{t-1})$ is the conditional density of mapping a less diffused s_{t-1} to a more diffused s_t . The diffusion rate is set by parameters $\alpha_t \in (0, 1)$ and $\beta_t = 1 - \alpha_t$. The forward process is sampled from as follows. First, a clean sample s_0 is chosen. Then $S_t \sim q(\cdot | s_{t-1})$ are sampled, from $t = 1$ to $t = T$.

Reverse process. In the reverse process, a clean sample s_0 is generated through a chain of variables S_T, \dots, S_0 (Figure 1). We denote the probability measure under this process by p and define its joint probability distribution as

$$p(s_{0:T}) = p(s_T) \prod_{t=1}^T p(s_{t-1} | s_t), \quad (5)$$

$$p(s_T) = \mathcal{N}(s_T; \mathbf{0}_d, I_d), \quad \forall t \in [T] : p(s_{t-1} | s_t) = \mathcal{N}(s_{t-1}; \mu_t(s_t), \Sigma_t),$$

where $p(s_{t-1} | s_t)$ is the conditional density of mapping a more diffused s_t to a less diffused s_{t-1} . The function μ_t predicts the mean of $S_{t-1} | s_t$ and is learned (Appendix B). As in Ho et al. [19], we keep the covariance fixed at $\Sigma_t = \tilde{\beta}_t I_d$, where $\tilde{\beta}_t = \frac{1 - \bar{\alpha}_{t-1}}{1 - \alpha_t} \beta_t$ and $\bar{\alpha}_t = \prod_{\ell=1}^t \alpha_\ell$. This is known as a *stable diffusion*. We make this assumption only to simplify exposition. All our derivations in Section 4 hold when Σ_t is learned, for instance as in Bao et al. [8].

This process is called reverse because it is learned by reversing the forward process. The reverse process is sampled from as follows. First, an initial diffused sample $S_T \sim p$ is sampled. After that, $S_{t-1} \sim p(\cdot | s_t)$ are sampled, from $t = T$ to $t = 1$.

4 Posterior Sampling

This section is organized as follows. In Section 4.1, we show how to sample from a chain of random variables conditioned on observations. In Sections 4.2 and 4.4, we specialize this to the observation models in Section 2.

4.1 Chain Model Posterior

Let $h = \{(\phi_\ell, y_\ell)\}_{\ell \in [N]}$ denote a *history* of N observations (Section 2) and H be the corresponding random variable. In this section, we assume that h is fixed. The Markovian structure of the reverse process (Figure 1) implies that the joint probability distribution conditioned on h factors as

$$p(s_{0:T} | h) = p(s_T | h) \prod_{t=1}^T p(s_{t-1} | s_t, h).$$

Therefore, $p(s_{0:T} | h)$ can be sampled from efficiently by first sampling from $p(s_T | h)$ and then from T conditional distributions $p(s_{t-1} | s_t, h)$. We derive these next.

Lemma 1. *Let p be a probability measure over the reverse process (Figure 1). Then*

$$p(s_T | h) \propto \int_{s_0} p(h | s_0) p(s_0 | s_T) ds_0 p(s_T),$$

$$\forall t \in [T] \setminus \{1\} : p(s_{t-1} | s_t, h) \propto \int_{s_0} p(h | s_0) p(s_0 | s_{t-1}) ds_0 p(s_{t-1} | s_t),$$

$$p(s_0 | s_1, h) \propto p(h | s_0) p(s_0 | s_1).$$

Algorithm 2 LaplaceDPS: Laplace posterior sampling with a diffusion model prior.

- 1: **Input:** Diffusion model parameters $(\mu_t, \Sigma_t)_{t \in [T]}$, history of observations h
 - 2: Initial sample $S_T \sim \mathcal{N}(\hat{\mu}_{T+1}(h), \hat{\Sigma}_{T+1}(h))$
 - 3: **for** stage $t = T, \dots, 1$ **do**
 - 4: $S_{t-1} \sim \mathcal{N}(\hat{\mu}_t(S_t, h), \hat{\Sigma}_t(h))$
 - 5: **Output:** Posterior sample S_0
-

Proof. The claim is proved in Appendix A.1. □

4.2 Linear Model Posterior

Now we specialize Lemma 1 to the diffusion model prior (Section 3) and linear models (Section 2.1). The prior distribution is the reverse process in (5),

$$p(s_T) = \mathcal{N}(s_T; \mathbf{0}_d, I_d), \quad \forall t \in [T] : p(s_{t-1} | s_t) = \mathcal{N}(s_{t-1}; \mu_t(s_t), \Sigma_t).$$

The term $p(h | s_0)$ is the likelihood of observations in Assumption 1. The main challenge in using the lemma is that the conditional densities of clean samples $p(s_0 | S_T)$ and $p(s_0 | s_t)$ are complex [12]. To get around this, we make an additional assumption, which is discussed in Section 4.3.

Theorem 2. Let p be a probability measure over the reverse process (Figure 1). Let $\bar{\theta}$ and $\bar{\Sigma}^{-1}$ be defined as in (3). Suppose that

$$\int_{s_0} p(h | s_0) p(s_0 | s_t) ds_0 \propto p(h | s_t / \sqrt{\bar{\alpha}_t}) \tag{6}$$

holds for all $t \in [T]$. Then $p(s_T | h) \propto \mathcal{N}(s_T; \hat{\mu}_{T+1}(h), \hat{\Sigma}_{T+1}(h))$, where

$$\hat{\mu}_{T+1}(h) = \hat{\Sigma}_{T+1}(h) \underbrace{(I_d \mathbf{0}_d)}_{\text{Prior}} + \underbrace{\bar{\Sigma}^{-1} \bar{\theta} / \sqrt{\bar{\alpha}_T}}_{\text{Evidence}}, \quad \hat{\Sigma}_{T+1}(h) = \underbrace{(I_d)}_{\text{Prior}} + \underbrace{\bar{\Sigma}^{-1} / \bar{\alpha}_T}_{\text{Evidence}}^{-1}. \tag{7}$$

For $t \in [T]$, we have $p(s_{t-1} | s_t, h) \propto \mathcal{N}(s_{t-1}; \hat{\mu}_t(s_t, h), \hat{\Sigma}_t(h))$, where

$$\hat{\mu}_t(s_t, h) = \hat{\Sigma}_t(h) \underbrace{(\Sigma_t^{-1} \mu_t(s_t))}_{\text{Prior}} + \underbrace{\bar{\Sigma}^{-1} \bar{\theta} / \sqrt{\bar{\alpha}_{t-1}}}_{\text{Evidence}}, \quad \hat{\Sigma}_t(h) = \underbrace{(\Sigma_t^{-1})}_{\text{Prior}} + \underbrace{\bar{\Sigma}^{-1} / \bar{\alpha}_{t-1}}_{\text{Evidence}}^{-1}. \tag{8}$$

Proof. The proof is in Appendix A.2. It has four steps. First, we fix stage t and apply approximation (6). Second, we rewrite the likelihood as in (3). Third, we reparameterize it as a function of s_t . At the end, we combine the likelihood with the Gaussian prior using Lemma 6 in Appendix A.5. □

The algorithm that samples from the posterior distribution in Theorem 2 is presented in Algorithm 2. We call it *Laplace diffusion posterior sampling (LaplaceDPS)* because its generalization to GLMs uses the Laplace approximation. LaplaceDPS samples from a chain of products of two distributions: one distribution represents the pre-trained diffusion model and does not depend on history h , and the other represents the history h . The sampling is implemented as follows. The initial variable S_T is sampled conditioned on h (line 2) from the distribution in (7). This distribution is a product of the h -independent prior $\mathcal{N}(\mathbf{0}_d, I_d)$ and the h -dependent distribution of the diffused evidence up to stage T , $\mathcal{N}(\sqrt{\bar{\alpha}_T} \bar{\theta}, \bar{\alpha}_T \bar{\Sigma})$. Then, for any $t \in [T]$, S_{t-1} is sampled conditioned on s_t and evidence h (line 4) from the distribution in (8). This distribution is a product of the h -independent conditional prior $\mathcal{N}(\mu_t(s_t), \Sigma_t)$, from the pre-trained model, and the h -dependent distribution of the diffused evidence up to stage $t - 1$, $\mathcal{N}(\sqrt{\bar{\alpha}_{t-1}} \bar{\theta}, \bar{\alpha}_{t-1} \bar{\Sigma})$. The last variable S_0 is the clean sample. When compared to Section 2, the prior and evidence are mixed conditionally in a T -stage chain. This increases the computational cost T times, as discussed in Section 8.

4.3 Key Approximation in Theorem 2

Now we motivate our assumption in (6). Simply put, we assume that $s_0 = s_t / \sqrt{\bar{\alpha}_t}$, where s_0 is a clean sample and s_t is the corresponding diffused sample in stage t . This is motivated by the forward process, which relates s_t and s_0 as $s_t = \sqrt{\bar{\alpha}_t} s_0 + \sqrt{1 - \bar{\alpha}_t} \varepsilon_t$, where $\varepsilon_t \sim \mathcal{N}(\mathbf{0}_d, I_d)$ is a standard

Algorithm 3 Contextual Thompson sampling.

- 1: **for** round $k = 1, \dots, n$ **do**
 - 2: Sample $\tilde{\theta}_k \sim p(\cdot | h_k)$, where $p(\cdot | h_k)$ is the posterior distribution in (2)
 - 3: Take action $a_k \leftarrow \arg \max_{a \in \mathcal{A}} r(x_k, a; \tilde{\theta}_k)$ and observe reward y_k
-

Gaussian noise [19]. After rearranging, we get $s_0 = (s_t - \sqrt{1 - \bar{\alpha}_t} \varepsilon_t) / \sqrt{\bar{\alpha}_t}$, and therefore s_0 can be viewed as a random variable with mean $s_t / \sqrt{\bar{\alpha}_t}$. The consequence of (6) is that the likelihood becomes a function of s_t , which yields a closed form when multiplied by the conditional prior, also a function of s_t . Our approximation can be also viewed as the Tweedie’s formula in Chung et al. [12] where the score component is neglected.

Our approximation has several notable properties. First, $\sqrt{(1 - \bar{\alpha}_t) / \bar{\alpha}_t} \rightarrow 0$ as $t \rightarrow 1$. Therefore, it becomes more precise in later stages of the reverse process. Second, in the absence of evidence h , the approximation vanishes, and all posterior distributions in Theorem 2 reduce to the priors in (5). Finally, as the number of observations increases, sampling from the posterior in Theorem 2 is asymptotically consistent.

Theorem 3. Fix $\theta_* \in \mathbb{R}^d$. Let $\tilde{\theta} \leftarrow \text{LaplaceDPS}((\mu_t, \Sigma_t)_{t \in [T]}, h)$, where $h = \{(\phi_\ell, y_\ell)\}_{\ell \in [N]}$ is a history of N observations. Suppose that $\lambda_d(\bar{\Sigma}^{-1}) \rightarrow \infty$ as $N \rightarrow \infty$, where $\bar{\Sigma}$ is defined in (3). Then $\mathbb{P}(\lim_{N \rightarrow \infty} \|\tilde{\theta} - \theta_*\|_2 = 0) = 1$.

Proof. The proof is in Appendix A.3. The key idea is that the conditional posteriors in (7) and (8) concentrate at a scaled unknown model parameter θ_* as the number of observations increases, which we formalize as $\lambda_d(\bar{\Sigma}^{-1}) \rightarrow \infty$. \square

The bound in Theorem 3 can be interpreted as follows. The sampled parameter $\tilde{\theta}$ approaches the true unknown parameter θ_* as the number of observations N increases. To guarantee that the posterior shrinks uniformly in all directions, we assume that the number of observations in all directions grows linearly with N . This is akin to assuming that $\lambda_d(\bar{\Sigma}^{-1}) = \Omega(N)$. This lower bound can be attained in linear models by getting observations according to the D-optimal design [39].

4.4 GLM Posterior

The Laplace approximation in GLMs (Section 2.2) naturally generalizes the exact posterior distribution in linear models (Section 2.1). We generalize Theorem 2 to GLMs along the same lines.

Theorem 4. Let p be a probability measure over the reverse process (Figure 1). Suppose that (6) holds for all $t \in [T]$. Then $p(s_T | h) \propto \mathcal{N}(s_T; \hat{\mu}_{T+1}(h), \hat{\Sigma}_{T+1}(h))$, where

$$\hat{\mu}_{T+1}(h) = \sqrt{\bar{\alpha}_T} \hat{\theta}_{T+1}, \quad \hat{\Sigma}_{T+1}(h) = \bar{\alpha}_T \hat{\Sigma}_{T+1}, \quad \hat{\theta}_{T+1}, \hat{\Sigma}_{T+1} \leftarrow \text{IRLS}(\mathbf{0}_d, I_d / \bar{\alpha}_T, h).$$

For $t \in [T]$, we have $p(s_{t-1} | s_t, h) \propto \mathcal{N}(s_{t-1}; \hat{\mu}_t(s_t, h), \hat{\Sigma}_t(h))$, where

$$\hat{\mu}_t(s_t, h) = \sqrt{\bar{\alpha}_{t-1}} \hat{\theta}_t, \quad \hat{\Sigma}_t(h) = \bar{\alpha}_{t-1} \hat{\Sigma}_t, \quad \hat{\theta}_t, \hat{\Sigma}_t \leftarrow \text{IRLS}(\mu_t(s_t) / \sqrt{\bar{\alpha}_{t-1}}, \Sigma_t / \bar{\alpha}_{t-1}, h).$$

Proof. The proof is in Appendix A.4. It has four steps. First, we fix stage t and apply approximation (6). Second, we reparameterize the prior, from a function of s_t to a function of $s_t / \sqrt{\bar{\alpha}_t}$. Third, we combine the likelihood with the prior using the Laplace approximation. Finally, we reparameterize the posterior, from a function of $s_t / \sqrt{\bar{\alpha}_t}$ to a function of s_t . \square

Similarly to Theorem 2, the distributions in Theorem 4 mix evidence with the diffusion model prior. However, this is done implicitly in IRLS. The posterior can be sampled from using LaplaceDPS, where the mean and covariances would be taken from Theorem 4. Note that Theorem 2 is a special case of Theorem 4 where the mean function g is an identity.

5 Application to Contextual Bandits

Now we apply our posterior sampling approximations (Section 4) to contextual bandits. A *contextual bandit* [29, 32] is a classic model for sequential decision making under uncertainty where the agent takes actions conditioned on context. We denote the *action set* by \mathcal{A} and the *context set* by \mathcal{X} . The *mean reward* for taking action $a \in \mathcal{A}$ in context $x \in \mathcal{X}$ is $r(x, a; \theta_*)$, where $r : \mathcal{X} \times \mathcal{A} \times \Theta \rightarrow \mathbb{R}$ is a *reward function* and $\theta_* \in \Theta$ is a *model parameter* (Section 2). The agent interacts with the bandit for n rounds indexed by $k \in [n]$. In round k , it observes a *context* $x_k \in \mathcal{X}$, takes an *action* $a_k \in \mathcal{A}$, and observes its *stochastic reward* $y_k = r(x_k, a_k; \theta_*) + \varepsilon_k$ with independent noise ε_k . We assume that the noise is zero-mean σ^2 -sub-Gaussian for $\sigma > 0$. The objective of the agent is to maximize its cumulative reward in n rounds, or equivalently to minimize its cumulative regret. We define the *n-round regret* as

$$R(n) = \sum_{k=1}^n \mathbb{E} [r(x_k, a_{k,*}; \theta_*) - r(x_k, a_k; \theta_*)] , \quad (9)$$

where $a_{k,*} = \arg \max_{a \in \mathcal{A}} r(x_k, a; \theta_*)$ is the optimal action in round k .

Arguably the most popular method for solving contextual bandit problems is Thompson sampling [50, 11, 3]. The key idea in TS is to use the posterior distribution of θ_* to explore. This is done as follows. In round k , the model parameter is drawn from the posterior in (2), $\tilde{\theta}_k \sim p(\cdot | h_k)$, where h_k is the *history* of all interactions up to round k . After that, the agent takes the action with the highest mean reward under $\tilde{\theta}_k$. The pseudo-code of this algorithm is given in Algorithm 3.

A *linear bandit* [13] has a linear reward function $r(x, a; \theta_*) = \phi(x, a)^\top \theta_*$, where $\phi : \mathcal{X} \times \mathcal{A} \rightarrow \mathbb{R}^d$ is a *feature extractor*. The feature extractor can be non-linear in x and a . Therefore, linear bandits can be applied to non-linear functions of x and a . The feature extractor can be either learned [40] or hand-crafted. We denote the feature vector of the action in round k by $\phi_k = \phi(x_k, a_k)$. Therefore, the *history* of interactions up to round k is $h_k = \{(\phi_\ell, y_\ell)\}_{\ell \in [k-1]}$. When the prior distribution is a Gaussian, $p(\theta_*) = \mathcal{N}(\theta_*; \theta_0, \Sigma_0)$, the posterior in round k is a Gaussian in (3) for $h = h_k$. When the prior is a diffusion model, we propose sampling from the posterior using

$$\tilde{\theta}_k \leftarrow \text{LaplaceDPS}((\mu_t, \Sigma_t)_{t \in [T]}, h_k) , \quad (10)$$

where $\hat{\mu}_t$ and $\hat{\Sigma}_t$ in LaplaceDPS are computed as in Theorem 2. We call this algorithm **DiffTS**.

A *generalized linear bandit* [17, 24, 33, 26] is an extension of linear bandits to generalized linear models (Section 2.2). When $p(\theta_*) = \mathcal{N}(\theta_*; \theta_0, \Sigma_0)$, the Laplace approximation to the posterior is a Gaussian (Section 2.2). When the prior is a diffusion model, we propose posterior sampling using (10), where $\hat{\mu}_t$ and $\hat{\Sigma}_t$ in LaplaceDPS are computed as in Theorem 4.

6 Experiments

We conduct three experiments: synthetic problems in 2 dimensions (Section 6.2 and Appendix C.1), a recommender system (Section 6.3), and a classification problem (Appendix C.2). In addition, we conduct an ablation study in Appendix C.3, where we vary the number of training samples for the diffusion prior and the number of diffusion stages T .

6.1 Experimental Setup

We have four baselines. Three baselines are variants of contextual Thompson sampling [11, 3]: with an uninformative Gaussian prior (**TS**), a learned Gaussian prior (**TunedTS**), and a learned Gaussian mixture prior (**MixTS**) [22]. The last baseline is diffusion posterior sampling (**DPS**) of Chung et al. [12]. We implement all TS baselines as described in Section 5. The uninformative prior is $\mathcal{N}(\mathbf{0}_d, I_d)$. **MixTS** is used only in linear bandit experiments because the logistic regression variant does not exist. The TS baselines are chosen to cover various levels of prior information. Our implementation of **DPS** is described in Appendix D. We also experimented with frequentist baselines, such as **LinUCB** [1] and the ε -greedy policy. They performed worse than **TS** and thus we do not report them.

Each experiment is set up as follows. First, the prior distribution of θ_* is specified: it can be synthetic or estimated from real-world data. Second, we learn this distribution from 10 000 samples from it. In **DiffTS** and **DPS**, we follow Appendix B. The number of stages is $T = 100$ and the diffusion factor

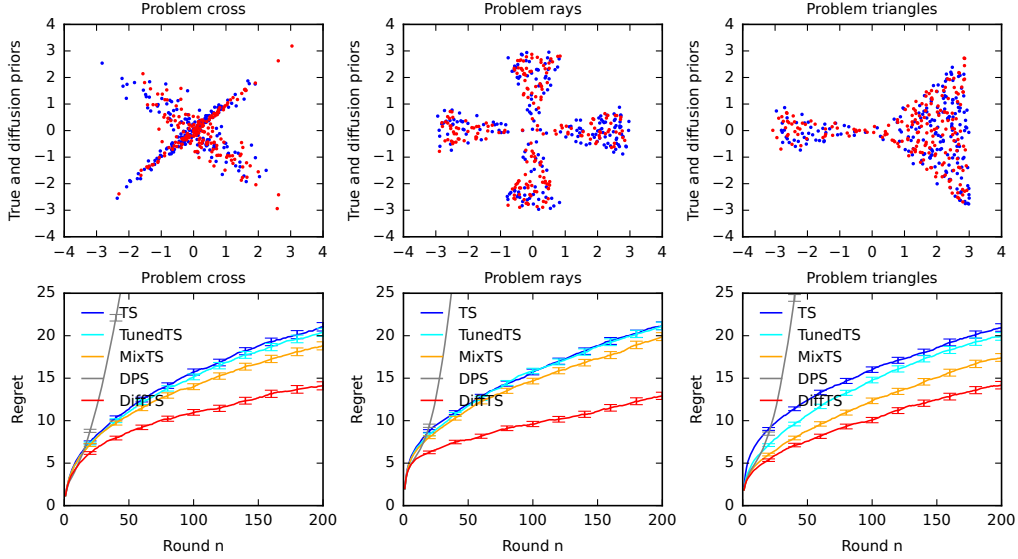


Figure 2: Evaluation of **DiffTS** on three synthetic problems. The first row shows samples from the true (blue) and diffusion model (red) priors. The second row shows the regret of **DiffTS** and the baselines as a function of round n .

is $\alpha_t = 0.97$. Since $0.97^{100} \approx 0.05$, most of the information in the training samples is diffused. The regressor in Appendix B is a 2-layer neural network with ReLU activations. In **TunedTS**, we fit the mean and covariance using maximum likelihood estimation. In **MixTS**, we fit the Gaussian mixture using SCIKIT-LEARN. All algorithms are evaluated on θ_* sampled from the true prior. The regret is computed as defined in (9). All error bars are standard errors of the estimates.

6.2 Synthetic Experiment

The first experiment is on three synthetic problems. Each problem is a linear bandit (Section 5) with $K = 100$ actions in $d = 2$ dimensions. The reward noise is $\sigma = 1$. The feature vectors of actions are sampled uniformly at random from a unit ball. The prior distributions of θ_* are shown in Figure 2. The first is a mixture of two Gaussians and the last can be approximated well by a mixture of two Gaussians. We implement **MixTS** with two mixture components. Therefore, it can represent the first prior exactly and approximate the last one well.

Our results are reported in Figure 2. We observe two main trends. First, samples from the diffusion prior closely resemble those from the true prior. In such cases, **DiffTS** is expected to perform well and even outperforms **MixTS**, because it has a better representation of the prior. We observe this in all problems. Second, **DPS** diverges as the number of rounds increases. This is because **DPS** uses an approximation based on the likelihood score (Section 7), which is unstable when the score is high. This happens despite our best efforts to tune **DPS** (Appendix D). We report results on another three synthetic problems in Appendix C.1.

DiffTS should be T times more computationally costly than **TS** with a Gaussian prior (Section 4.2). We observe this empirically. As an example, the average cost of 100 runs of **DiffTS** on any problem in Figure 2 is 12 seconds. The average cost of **TS** is 0.1 seconds. The computation and accuracy can be traded off, and we investigate this in Appendix C.3. In the cross problem, we vary the number of diffusion stages from $T = 1$ to $T = 300$. We observe that the computational cost is linear in T and the regret drops quickly from 26 at $T = 1$ to 15 at $T = 50$.

6.3 MovieLens Experiment

In the second experiment, we learn to recommend an item to randomly arriving users. The problem is simulated using the MovieLens 1M dataset [28], with one million ratings for 3 706 movies from 6 040 users. We subtract the mean rating from all ratings and complete the sparse rating matrix M by




Figure 3: Evaluation of **DiffTS** on the MovieLens dataset: (a) shows samples from the true (blue) and diffusion model (red) priors, (b) shows regret in the linear bandit, and (c) shows regret in the logistic bandit.

alternating least squares [14] with rank $d = 5$. The learned factorization is $M = UV^\top$. The i -th row of U , denoted by U_i , represents user i . The j -th row of V , denoted by V_j , represents movie j . We use movie embeddings V_j as model parameters and user embeddings U_i as features of the actions. The movies are items.

We experiment with both linear and logistic bandits. In both, an item is initially chosen randomly from V_j and $K = 10$ actions are chosen randomly from U_i in each round. In the linear bandit, the mean reward of item j for user i is $U_i^\top V_j$. The reward noise is $\sigma = 0.75$, and we estimate it from data. In the logistic bandit, the mean reward is $g(U_i^\top V_j)$, where g is a sigmoid.

Our MovieLens results are reported in Figure 3 and we observe similar trends to Section 6.2. First, samples from the diffusion prior closely resemble those from the true prior (Figure 3a). Since the problem is higher dimensional, we visualize the overlap using UMAP [45]. Second, **DiffTS** has a lower regret than all baselines, in both linear (Figure 3b) and logistic (Figure 3c) bandits. Finally, **MixTS** barely outperforms **TunedTS**. We observe this trend consistently in higher dimensions, and this motivated our work on online learning with more complex priors.

7 Related Work

We start with reviewing related works on bandits with diffusion models. Hsieh et al. [23] proposed Thompson sampling with a diffusion model prior for K -armed bandits. There are multiple technical differences from our work. First, the diffusion model in Hsieh et al. [23] is over scalars representing individual arms. Our model is over vectors representing model parameters, and thus can be applied to contextual bandits. Second, the approximations are different. In stage t , Hsieh et al. [23] sample from two distributions: the conditional prior and the distribution of the diffused empirical mean up to stage t . Then they take a weighted sum of the samples. We sample only once, from the posterior distribution that combines the conditional prior in stage t and likelihood. Therefore, the method of Hsieh et al. [23] can be viewed as a non-contextual variant of our method, where posterior sampling is done by weighting samples from the prior and empirical distributions. Finally, Hsieh et al. [23] do not analyze their approximation.

Aouali [4] proposed and analyzed contextual bandits with a linear diffusion model prior: $\mu_t(s_t)$ in (5) is linear in s_t and $q(s_0)$ is a Gaussian. Therefore, this model is a linear Gaussian model and not a general diffusion model, as in our work.

The closest related work on posterior sampling in diffusion models is **DPS** of Chung et al. [12]. The key idea in **DPS** is to sample from the posterior distribution using the likelihood score $\nabla \log p(h | \theta)$, where $p(h | \theta)$ is the likelihood (Assumptions 1 and 2). Note that $\nabla \log p(h | \theta)$ grows linearly in N because the history h in $p(h | \theta)$ involves N terms. Therefore, **DPS** becomes unstable as $N \rightarrow \infty$. We show it empirically in Section 6.2 and discuss the implementation of **DPS** in Appendix D, which was tuned to improve its stability.

Many other posterior sampling methods for diffusion models have been proposed recently: a sequential Monte Carlo approximation for the conditional reverse process [53], a variant of **DPS** with an uninformative prior [38], a pseudo-inverse approximation to the likelihood of evidence [48], and

posterior sampling in latent diffusion models [41]. All of these methods rely on the likelihood score $\nabla \log p(h | \theta)$ and thus become unstable as the number of observations N increases. Our posterior approximations do not have this issue because they are based on the product of prior and evidence distributions (Theorems 2 and 4), and thus gradient-free. They work well across different levels of uncertainty (Section 6) and do not require tuning.

We note that posterior sampling is a special form of inference-time guidance in diffusion models. Other approaches are conditional pre-training [15], a constraint in the reverse process [18], refining the null-space content [51], solving an optimization problem that pushes the reverse process towards evidence [47], and aligning the reverse process with the prompt [7].

8 Conclusions

We propose posterior sampling approximations for diffusion models priors. These approximations are contextual, and can be implemented efficiently in linear models and GLMs. We analyze them and evaluate them empirically on contextual bandit problems. Our method has two main limitations.

Computational cost. The cost of posterior sampling in `LaplaceDPS` with T stages is about T times higher than that of posterior sampling with a Gaussian prior (Section 2). We validate it empirically in Section 6.2. We plot the sampling time as a function of T in Figure 6c (Appendix C.3).

Learning cost and hyper-parameter tuning. In all experiments, the number of diffusion stages is $T = 100$ and the diffusion rate is set such that most of the signal diffuses. The regressor is a 2-layer neural network and we learn it from 10 000 samples from the prior. These settings resulted in stable performance in all our experiments (Section 6). However, they clearly impact the performance. We plot the regret as a function of the number of training samples in Figure 6a and as a function of T in Figure 6b. When T or the number of training samples is small, `DiffTS` performs very similarly to posterior sampling with a Gaussian prior. In summary, there is no benefit in these cases.

Future work. We develop novel posterior approximations rather than bounding their regret. This is because the existing approximations are unstable and may diverge in the online setting (Sections 6.2 and 7). We believe that a proper regret analysis of `DiffTS` is possible and would require bounding two errors. The first error arises because the reverse process does not reverse the forward process exactly (Appendix B). The second error arises because our posterior distributions are approximate (Section 4.3). One possibility is to start with prior works that already showed the utility of complex priors. For instance, Russo and Van Roy [43] proved a $O(\sqrt{\Gamma H(A_*)n})$ regret bound for a linear bandit, where Γ is the maximum ratio of regret to information gain and $H(A_*)$ is the entropy of the distribution of the optimal action under the prior. This bound holds for any prior and says that a lower entropy $H(A_*)$, which corresponds to more informative priors, yields a lower regret.

We also believe that our ideas can be extended beyond GLMs. The key idea in Section 4.4 is to use the Laplace approximation of the likelihood. This approximation can be computed exactly in GLMs. More generally though, it is a good approximation whenever the likelihood can be approximated well by a single Gaussian distribution. By the central limit theorem, under appropriate assumptions, this is expected for any observation model when the number of observations is large.

References

- [1] Yasin Abbasi-Yadkori, David Pal, and Csaba Szepesvari. Improved algorithms for linear stochastic bandits. In *Advances in Neural Information Processing Systems 24*, pages 2312–2320, 2011.
- [2] Marc Abeille and Alessandro Lazaric. Linear Thompson sampling revisited. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, 2017.
- [3] Shipra Agrawal and Navin Goyal. Thompson sampling for contextual bandits with linear payoffs. In *Proceedings of the 30th International Conference on Machine Learning*, pages 127–135, 2013.
- [4] Imad Aouali. Linear diffusion models meet contextual bandits with large action spaces. In *NeurIPS 2023 Workshop on Foundation Models for Decision Making Workshop*, 2023.
- [5] Imad Aouali, Branislav Kveton, and Sumeet Katariya. Mixed-effect Thompson sampling. In *Proceedings of the 26th International Conference on Artificial Intelligence and Statistics*, 2023.

- [6] Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47:235–256, 2002.
- [7] Arpit Bansal, Hong-Min Chu, Avi Schwarzschild, Soumyadip Sengupta, Micah Goldblum, Jonas Geiping, and Tom Goldstein. Universal guidance for diffusion models. In *Proceedings of the 12th International Conference on Learning Representations*, 2024.
- [8] Fan Bao, Chongxuan Li, Jiacheng Sun, Jun Zhu, and Bo Zhang. Estimating the optimal covariance with imperfect mean in diffusion probabilistic models. In *Proceedings of the 39th International Conference on Machine Learning*, 2022.
- [9] Soumya Basu, Branislav Kveton, Manzil Zaheer, and Csaba Szepesvari. No regrets for learning the prior in bandits. In *Advances in Neural Information Processing Systems 34*, 2021.
- [10] Christopher Bishop. *Pattern Recognition and Machine Learning*. Springer, New York, NY, 2006.
- [11] Olivier Chapelle and Lihong Li. An empirical evaluation of Thompson sampling. In *Advances in Neural Information Processing Systems 24*, pages 2249–2257, 2011.
- [12] Hyungjin Chung, Jeongsol Kim, Michael Thompson Mccann, Marc Louis Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. In *Proceedings of the 11th International Conference on Learning Representations*, 2023.
- [13] Varsha Dani, Thomas Hayes, and Sham Kakade. Stochastic linear optimization under bandit feedback. In *Proceedings of the 21st Annual Conference on Learning Theory*, pages 355–366, 2008.
- [14] Mark Davenport and Justin Romberg. An overview of low-rank matrix recovery from incomplete observations. *IEEE Journal of Selected Topics in Signal Processing*, 10(4):608–622, 2016.
- [15] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. In *Advances in Neural Information Processing Systems 34*, 2021.
- [16] Arnaud Doucet, Nando de Freitas, and Neil Gordon. *Sequential Monte Carlo Methods in Practice*. Springer, New York, NY, 2001.
- [17] Sarah Filippi, Olivier Cappe, Aurelien Garivier, and Csaba Szepesvari. Parametric bandits: The generalized linear case. In *Advances in Neural Information Processing Systems 23*, pages 586–594, 2010.
- [18] Alexandros Graikos, Nikolay Malkin, Nebojsa Jojic, and Dimitris Samaras. Diffusion models as plug-and-play priors. In *Advances in Neural Information Processing Systems 35*, 2022.
- [19] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems 33*, 2020.
- [20] Joey Hong, Branislav Kveton, Sumeet Katariya, Manzil Zaheer, and Mohammad Ghavamzadeh. Deep hierarchy in bandits. In *Proceedings of the 39th International Conference on Machine Learning*, 2022.
- [21] Joey Hong, Branislav Kveton, Manzil Zaheer, and Mohammad Ghavamzadeh. Hierarchical Bayesian bandits. In *Proceedings of the 25th International Conference on Artificial Intelligence and Statistics*, 2022.
- [22] Joey Hong, Branislav Kveton, Manzil Zaheer, Mohammad Ghavamzadeh, and Craig Boutilier. Thompson sampling with a mixture prior. In *Proceedings of the 25th International Conference on Artificial Intelligence and Statistics*, 2022.
- [23] Yu-Guan Hsieh, Shiva Kasiviswanathan, Branislav Kveton, and Patrick Blobaum. Thompson sampling with diffusion generative prior. In *Proceedings of the 40th International Conference on Machine Learning*, 2023.
- [24] Kwang-Sung Jun, Aniruddha Bhargava, Robert Nowak, and Rebecca Willett. Scalable generalized linear bandits: Online computation and hashing. In *Advances in Neural Information Processing Systems 30*, pages 98–108, 2017.
- [25] Jaya Kawale, Hung Bui, Branislav Kveton, Long Tran-Thanh, and Sanjay Chawla. Efficient Thompson sampling for online matrix-factorization recommendation. In *Advances in Neural Information Processing Systems 28*, pages 1297–1305, 2015.

- [26] Branislav Kveton, Manzil Zaheer, Csaba Szepesvari, Lihong Li, Mohammad Ghavamzadeh, and Craig Boutilier. Randomized exploration in generalized linear bandits. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics*, 2020.
- [27] Tze Leung Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6(1):4–22, 1985.
- [28] Shyong Lam and Jon Herlocker. MovieLens Dataset. <http://grouplens.org/datasets/movielens/>, 2016.
- [29] John Langford and Tong Zhang. The epoch-greedy algorithm for contextual multi-armed bandits. In *Advances in Neural Information Processing Systems 20*, pages 817–824, 2008.
- [30] Tor Lattimore and Csaba Szepesvari. *Bandit Algorithms*. Cambridge University Press, 2019.
- [31] Yann LeCun, Corinna Cortes, and Christopher Burges. MNIST Handwritten Digit Database. <http://yann.lecun.com/exdb/mnist>, 2010.
- [32] Lihong Li, Wei Chu, John Langford, and Robert Schapire. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th International Conference on World Wide Web*, 2010.
- [33] Lihong Li, Yu Lu, and Dengyong Zhou. Provably optimal algorithms for generalized linear contextual bandits. In *Proceedings of the 34th International Conference on Machine Learning*, pages 2071–2080, 2017.
- [34] Lisha Li, Kevin Jamieson, Giulia DeSalvo, Afshin Rostamizadeh, and Ameet Talwalkar. Hyperband: A novel bandit-based approach to hyperparameter optimization. *Journal of Machine Learning Research*, 18(185):1–52, 2018.
- [35] Shuai Li, Alexandros Karatzoglou, and Claudio Gentile. Collaborative filtering bandits. In *Proceedings of the 39th Annual International ACM SIGIR Conference*, 2016.
- [36] Xiuyuan Lu and Benjamin Van Roy. Information-theoretic confidence bounds for reinforcement learning. In *Advances in Neural Information Processing Systems 32*, 2019.
- [37] P. McCullagh and J. A. Nelder. *Generalized Linear Models*. Chapman & Hall, 1989.
- [38] Xiangming Meng and Yoshiyuki Kabashima. Diffusion model based posterior sampling for noisy linear inverse problems. *CoRR*, abs/2211.12343, 2023. URL <https://arxiv.org/abs/2211.12343>.
- [39] Friedrich Pukelsheim. *Optimal Design of Experiments*. Society for Industrial and Applied Mathematics, 2006.
- [40] Carlos Riquelme, George Tucker, and Jasper Snoek. Deep Bayesian bandits showdown: An empirical comparison of Bayesian deep networks for Thompson sampling. In *Proceedings of the 6th International Conference on Learning Representations*, 2018.
- [41] Litu Rout, Negin Raof, Giannis Daras, Constantine Caramanis, Alex Dimakis, and Sanjay Shakkottai. Solving linear inverse problems provably via posterior sampling with latent diffusion models. In *Advances in Neural Information Processing Systems 36*, 2023.
- [42] Daniel Russo and Benjamin Van Roy. Learning to optimize via posterior sampling. *Mathematics of Operations Research*, 39(4):1221–1243, 2014.
- [43] Daniel Russo and Benjamin Van Roy. An information-theoretic analysis of Thompson sampling. *Journal of Machine Learning Research*, 17(68):1–30, 2016.
- [44] Daniel Russo, Benjamin Van Roy, Abbas Kazerouni, Ian Osband, and Zheng Wen. A tutorial on Thompson sampling. *Foundations and Trends in Machine Learning*, 11(1):1–96, 2018.
- [45] Tim Sainburg, Leland McInnes, and Timothy Gentner. Parametric UMAP embeddings for representation and semisupervised learning. *Neural Computation*, 33(11):2881–2907, 2021.
- [46] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *Proceedings of the 32nd International Conference on Machine Learning*, pages 2256–2265, 2015.
- [47] Bowen Song, Soo Min Kwon, Zecheng Zhang, Xinyu Hu, Qing Qu, and Liyue Shen. Solving inverse problems with latent diffusion models via hard data consistency. In *Proceedings of the 12th International Conference on Learning Representations*, 2024.

- [48] Jiaming Song, Arash Vahdat, Morteza Mardani, and Jan Kautz. Pseudoinverse-guided diffusion models for inverse problems. In *Proceedings of the 11th International Conference on Learning Representations*, 2023.
- [49] Yang Song, Jascha Sohl-Dickstein, Diederik Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *Proceedings of the 9th International Conference on Learning Representations*, 2021.
- [50] William R. Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3-4):285–294, 1933.
- [51] Yinhuai Wang, Jiwen Yu, and Jian Zhang. Zero-shot image restoration using denoising diffusion null-space model. In *Proceedings of the 11th International Conference on Learning Representations*, 2023.
- [52] R. Wolke and H. Schwetlick. Iteratively reweighted least squares: Algorithms, convergence analysis, and numerical comparisons. *SIAM Journal on Scientific and Statistical Computing*, 9(5):907–921, 1988.
- [53] Luhuan Wu, Brian Trippe, Christian Naesseth, David Blei, and John Cunningham. Practical and asymptotically exact conditional sampling in diffusion models. In *Advances in Neural Information Processing Systems 36*, 2023.
- [54] Xiaoxue Zhao, Weinan Zhang, and Jun Wang. Interactive collaborative filtering. In *Proceedings of the 22nd ACM International Conference on Information and Knowledge Management*, pages 1411–1420, 2013.

A Proofs and Supporting Lemmas

This section contains proofs of our main claims and supporting lemmas.

A.1 Proof of Lemma 1

All derivations are based on basic rules of probability and the chain structure in Figure 1, and are exact. From Figure 1, the joint probability distribution conditioned on $H = h$ factors as

$$p(s_{0:T} | h) = p(s_T | h) \prod_{t=1}^T p(s_{t-1} | s_{t:T}, h) = p(s_T | h) \prod_{t=1}^T p(s_{t-1} | s_t, h).$$

We use that $p(s_{t-1} | s_{t:T}, h) = p(s_{t-1} | s_t, h)$ in the last equality. We consider two cases.

Derivation of $p(s_{t-1} | s_t, h)$. By Bayes' rule, we get

$$p(s_{t-1} | s_t, h) = \frac{p(h | s_{t-1}, s_t) p(s_{t-1} | s_t)}{p(h | s_t)} \propto p(h | s_{t-1}) p(s_{t-1} | s_t).$$

In the last step, we use that $p(h | s_t)$ is a constant, since s_t and h are fixed, and that $p(h | s_{t-1}, s_t) = p(h | s_{t-1})$. Note that the last term $p(s_{t-1} | s_t)$ is the conditional prior distribution. When $t > 1$, the first term can be expressed as

$$\begin{aligned} p(h | s_{t-1}) &= \int_{s_0} p(h, s_0 | s_{t-1}) ds_0 = \int_{s_0} p(h | s_0, s_{t-1}) p(s_0 | s_{t-1}) ds_0 \\ &= \int_{s_0} p(h | s_0) p(s_0 | s_{t-1}) ds_0. \end{aligned}$$

In the last equality, we use that our graphical model is a chain (Figure 1), and thus $p(h | s_0, s_{t-1}) = p(h | s_0)$. Finally, we chain all identities and get that

$$p(s_{t-1} | s_t, h) \propto \int_{s_0} p(h | s_0) p(s_0 | s_{t-1}) ds_0 p(s_{t-1} | s_t). \quad (11)$$

Derivation of $p(s_T | h)$. By Bayes' rule, we get

$$p(s_T | h) = \frac{p(h | s_T) p(s_T)}{p(h)} \propto p(h | s_T) p(s_T).$$

In the last step, we use that $p(h)$ is a constant, since h is fixed. The first term can be rewritten as

$$\begin{aligned} p(h | s_T) &= \int_{s_0} p(h, s_0 | s_T) ds_0 = \int_{s_0} p(h | s_0, s_T) p(s_0 | s_T) ds_0 \\ &= \int_{s_0} p(h | s_0) p(s_0 | s_T) ds_0. \end{aligned}$$

Finally, we chain all identities and get that

$$p(s_T | h) \propto \int_{s_0} p(h | s_0) p(s_0 | s_T) ds_0 p(s_T). \quad (12)$$

This completes the derivations.

A.2 Proof of Theorem 2

This proof has two parts.

Derivation of $p(s_{t-1} | s_t, h)$. From (6) and Assumption 1, it follows that

$$\begin{aligned} \int_{s_0} p(h | s_0) p(s_0 | s_{t-1}) ds_0 &\propto p(h | s_{t-1}/\sqrt{\bar{\alpha}_{t-1}}) \propto \mathcal{N}(s_{t-1}/\sqrt{\bar{\alpha}_{t-1}}; \bar{\theta}, \bar{\Sigma}) \\ &\propto \mathcal{N}(s_{t-1}; \sqrt{\bar{\alpha}_{t-1}}\bar{\theta}, \bar{\alpha}_{t-1}\bar{\Sigma}). \end{aligned}$$

The last step treats $\bar{\alpha}_{t-1}$ and $\bar{\Sigma}$ as constants, because the forward process and evidence h are fixed. Now we apply Lemma 6 to distributions

$$p(s_{t-1} | s_t) = \mathcal{N}(s_{t-1}; \mu_t(s_t), \Sigma_t), \quad \mathcal{N}(s_{t-1}; \sqrt{\bar{\alpha}_{t-1}}\bar{\theta}, \bar{\alpha}_{t-1}\bar{\Sigma}),$$

and get that

$$p(s_{t-1} | s_t, h) \propto \mathcal{N}(s_{t-1}; \hat{\mu}_t(s_t, h), \hat{\Sigma}_t(h)),$$

where $\hat{\mu}_t(s_t, h)$ and $\hat{\Sigma}_t(h)$ are defined in the claim. This is a product of two Gaussians: the prior with mean $\mu_t(s_t)$ and covariance Σ_t , and the evidence with mean $\sqrt{\bar{\alpha}_{t-1}}\bar{\theta}$ and covariance $\bar{\alpha}_{t-1}\bar{\Sigma}$.

Derivation of $p(s_T | h)$. Analogously to the derivation of $p(s_{t-1} | s_t, h)$, we establish that

$$\int_{s_0} p(h | s_0) p(s_0 | s_T) ds_0 \propto \mathcal{N}(s_T; \sqrt{\bar{\alpha}_T}\bar{\theta}, \bar{\alpha}_T\bar{\Sigma}).$$

Then we apply Lemma 6 to distributions

$$p(s_T) = \mathcal{N}(s_T; \mathbf{0}_d, I_d), \quad \mathcal{N}(s_T; \sqrt{\bar{\alpha}_T}\bar{\theta}, \bar{\alpha}_T\bar{\Sigma}),$$

and get that

$$p(s_T | h) \propto \mathcal{N}(s_T; \hat{\mu}_{T+1}(h), \hat{\Sigma}_{T+1}(h)),$$

where $\hat{\mu}_{T+1}(h)$ and $\hat{\Sigma}_{T+1}(h)$ are defined in the claim. This is a product of two Gaussians: the prior with mean $\mathbf{0}_d$ and covariance I_d , and the evidence with mean $\sqrt{\bar{\alpha}_T}\bar{\theta}$ and covariance $\bar{\alpha}_T\bar{\Sigma}$.

A.3 Proof of Theorem 3

We start with the triangle inequality

$$\|\tilde{\theta} - \theta_*\|_2 = \|\tilde{\theta} - \bar{\theta} + \bar{\theta} - \theta_*\|_2 \leq \|\tilde{\theta} - \bar{\theta}\|_2 + \|\bar{\theta} - \theta_*\|_2,$$

where we introduce $\bar{\theta}$ from Section 2.1. Now we bound each term on the right-hand side.

Upper bound on $\|\tilde{\theta} - \bar{\theta}\|_2$. This part of the proof is based on analyzing the asymptotic behavior of the conditional densities in Theorem 2.

As a first step, note that $S_T \sim \mathcal{N}(\hat{\mu}_{T+1}(h), \hat{\Sigma}_{T+1}(h))$, where

$$\hat{\mu}_{T+1}(h) = \hat{\Sigma}_{T+1}(h)(I_d \mathbf{0}_d + \bar{\Sigma}^{-1}\bar{\theta}/\sqrt{\bar{\alpha}_T}), \quad \hat{\Sigma}_{T+1}(h) = (I_d + \bar{\Sigma}^{-1}/\bar{\alpha}_T)^{-1}.$$

Since $\lambda_d(\bar{\Sigma}^{-1}) \rightarrow \infty$, we get

$$\hat{\Sigma}_{T+1}(h) \rightarrow \bar{\alpha}_T\bar{\Sigma}, \quad \hat{\mu}_{T+1}(h) \rightarrow \sqrt{\bar{\alpha}_T}\bar{\theta}.$$

Moreover, $\lambda_d(\bar{\Sigma}^{-1}) \rightarrow \infty$ implies $\lambda_1(\bar{\Sigma}) \rightarrow 0$, and thus $\lim_{N \rightarrow \infty} \|S_T - \sqrt{\bar{\alpha}_T}\bar{\theta}\|_2 = 0$.

The same argument can be applied inductively to later stages of the reverse process. Specifically, for any $t \in [T]$, $S_{t-1} \sim \mathcal{N}(\hat{\mu}_t(S_t, h), \hat{\Sigma}_t(h))$, where

$$\hat{\mu}_t(S_t, h) = \hat{\Sigma}_t(h)(\Sigma_t^{-1}\mu_t(S_t) + \bar{\Sigma}^{-1}\bar{\theta}/\sqrt{\bar{\alpha}_{t-1}}), \quad \hat{\Sigma}_t(h) = (\Sigma_t^{-1} + \bar{\Sigma}^{-1}/\bar{\alpha}_{t-1})^{-1}.$$

Since $\lambda_d(\bar{\Sigma}^{-1}) \rightarrow \infty$ and $S_t \rightarrow \sqrt{\bar{\alpha}_t}\bar{\theta}$ by induction, we get

$$\hat{\Sigma}_t(h) \rightarrow \bar{\alpha}_{t-1}\bar{\Sigma}, \quad \hat{\mu}_t(S_t, h) \rightarrow \sqrt{\bar{\alpha}_{t-1}}\bar{\theta}.$$

Moreover, $\lambda_d(\bar{\Sigma}^{-1}) \rightarrow \infty$ implies $\lambda_1(\bar{\Sigma}) \rightarrow 0$, and thus $\lim_{N \rightarrow \infty} \|S_{t-1} - \sqrt{\bar{\alpha}_{t-1}}\bar{\theta}\|_2 = 0$ for any $t \in [T]$. In the last stage, $t = 1$, $\bar{\alpha}_0 = 1$, and $S_0 = \tilde{\theta}$. Therefore,

$$\lim_{N \rightarrow \infty} \|S_{t-1} - \sqrt{\bar{\alpha}_{t-1}}\bar{\theta}\|_2 = \lim_{N \rightarrow \infty} \|\tilde{\theta} - \bar{\theta}\|_2 = 0.$$

Upper bound on $\|\bar{\theta} - \theta_*\|_2$. This part of the proof uses the definition of $\bar{\theta}$ in Section 2.1 and that $\varepsilon_\ell \sim \mathcal{N}(0, \sigma^2)$ is independent noise. By definition,

$$\bar{\theta} - \theta_* = \sigma^{-2}\bar{\Sigma} \sum_{\ell=1}^N \phi_\ell y_\ell - \theta_* = \sigma^{-2}\bar{\Sigma} \sum_{\ell=1}^N \phi_\ell (\phi_\ell^\top \theta_* + \varepsilon_\ell) - \theta_* = \sigma^{-2}\bar{\Sigma} \sum_{\ell=1}^N \phi_\ell \varepsilon_\ell.$$

Since ε_ℓ is independent zero-mean Gaussian noise with variance σ^2 , $\bar{\theta} - \theta_*$ is a Gaussian random variable with mean $\mathbf{0}_d$ and covariance

$$\text{cov} \left[\sigma^{-2} \bar{\Sigma} \sum_{\ell=1}^N \phi_\ell \varepsilon_\ell \right] = \sigma^{-4} \bar{\Sigma} \left(\sum_{\ell=1}^N \phi_\ell \text{var} [\varepsilon_\ell] \phi_\ell^\top \right) \bar{\Sigma} = \bar{\Sigma} \frac{\sum_{\ell=1}^N \phi_\ell \phi_\ell^\top}{\sigma^2} \bar{\Sigma} = \bar{\Sigma}.$$

Since $\lambda_d(\bar{\Sigma}^{-1}) \rightarrow \infty$ implies $\lambda_1(\bar{\Sigma}) \rightarrow 0$, we get

$$\lim_{N \rightarrow \infty} \|\bar{\theta} - \theta_*\|_2 = 0.$$

This completes the proof.

A.4 Proof of Theorem 4

This proof has two parts.

Derivation of $p(s_{t-1} | s_t, h)$. From (6), we have

$$\int_{s_0} p(h | s_0) p(s_0 | s_{t-1}) ds_0 \propto p(h | s_{t-1} / \sqrt{\bar{\alpha}_{t-1}}).$$

Since $p(s_{t-1} | s_t)$ is a Gaussian, we have

$$p(s_{t-1} | s_t) = \mathcal{N}(s_{t-1}; \mu_t(s_t), \Sigma_t) \propto \mathcal{N}(\gamma s_{t-1}; \gamma \mu_t(s_t), \gamma^2 \Sigma_t)$$

for $\gamma = 1/\sqrt{\bar{\alpha}_{t-1}}$. Then by the Laplace approximation,

$$p(h | \gamma s_{t-1}) \mathcal{N}(\gamma s_{t-1}; \gamma \mu_t(s_t), \gamma^2 \Sigma_t) \propto \mathcal{N}(\gamma s_{t-1}; \dot{\theta}_t, \dot{\Sigma}_t) \propto \mathcal{N}(s_{t-1}; \dot{\theta}_t / \gamma, \dot{\Sigma}_t / \gamma^2),$$

where $\dot{\theta}_t, \dot{\Sigma}_t \leftarrow \text{IRLS}(\gamma \mu_t(s_t), \gamma^2 \Sigma_t, h)$.

Derivation of $p(s_T | h)$. Analogously to the derivation of $p(s_{t-1} | s_t, h)$, we establish that

$$\int_{s_0} p(h | s_0) p(s_0 | s_T) ds_0 \propto p(h | s_T / \sqrt{\bar{\alpha}_T}).$$

Then by the Laplace approximation for $\gamma = 1/\sqrt{\bar{\alpha}_T}$, we get

$$p(h | \gamma s_T) \mathcal{N}(s_T; \mathbf{0}_d, I_d) \propto \mathcal{N}(s_T; \dot{\theta}_{T+1} / \gamma, \dot{\Sigma}_{T+1} / \gamma^2),$$

where $\dot{\theta}_{T+1}, \dot{\Sigma}_{T+1} \leftarrow \text{IRLS}(\mathbf{0}_d, \gamma^2 I_d, h)$.

A.5 Supporting Lemmas

We state and prove our supplementary lemmas next.

Lemma 5. Let $p(x) = \mathcal{N}(x; \mu_1, \Sigma_1)$ and $q(x) = \mathcal{N}(x; \mu_2, \Sigma_2)$, where $\mu_1, \mu_2 \in \mathbb{R}^d$ and $\Sigma_1, \Sigma_2 \in \mathbb{R}^{d \times d}$. Then

$$d(p, q) = \frac{1}{2} \left((\mu_2 - \mu_1)^\top \Sigma_2^{-1} (\mu_2 - \mu_1) + \text{tr}(\Sigma_2^{-1} \Sigma_1) - \log \frac{\det(\Sigma_1)}{\det(\Sigma_2)} - d \right).$$

Moreover, when $\Sigma_1 = \Sigma_2$,

$$d(p, q) = \frac{1}{2} (\mu_2 - \mu_1)^\top \Sigma_2^{-1} (\mu_2 - \mu_1).$$

Proof. The proof follows from the definitions of KL divergence and multivariate Gaussians. \square

Lemma 6. Fix $\mu_1 \in \mathbb{R}^d$, $\Sigma_1 \succeq 0$, $\mu_2 \in \mathbb{R}^d$, and $\Sigma_2 \succeq 0$. Then

$$\mathcal{N}(x; \mu_1, \Sigma_1) \mathcal{N}(x; \mu_2, \Sigma_2) \propto \mathcal{N}(x; \mu, \Sigma),$$

where

$$\mu = \Sigma(\Sigma_1^{-1} \mu_1 + \Sigma_2^{-1} \mu_2), \quad \Sigma = (\Sigma_1^{-1} + \Sigma_2^{-1})^{-1}.$$

Proof. This is a classic result, which is proved as

$$\begin{aligned}
\mathcal{N}(x; \mu_1, \Sigma_1) \mathcal{N}(x; \mu_2, \Sigma_2) &\propto \exp \left[-\frac{1}{2} \left((x - \mu_1)^\top \Sigma_1^{-1} (x - \mu_1) + (x - \mu_2)^\top \Sigma_2^{-1} (x - \mu_2) \right) \right] \\
&\propto \exp \left[-\frac{1}{2} \left(x^\top \Sigma_1^{-1} x - 2x^\top \Sigma_1^{-1} \mu_1 + x^\top \Sigma_2^{-1} x - 2x^\top \Sigma_2^{-1} \mu_2 \right) \right] \\
&= \exp \left[-\frac{1}{2} \left(x^\top \Sigma^{-1} x - 2x^\top \Sigma^{-1} \Sigma (\Sigma_1^{-1} \mu_1 + \Sigma_2^{-1} \mu_2) \right) \right] \\
&\propto \exp \left[-\frac{1}{2} (x - \mu)^\top \Sigma^{-1} (x - \mu) \right] \propto \mathcal{N}(x; \mu, \Sigma).
\end{aligned}$$

The neglected factors depend on constants $\mu_1, \mu_2, \Sigma_1,$ and Σ_2 . This completes the proof. \square

B Learning the Reverse Process

One property of our model is that $q(s_T) \approx \mathcal{N}(s_T; \mathbf{0}_d, I_d)$ when T is sufficiently large [19]. Since s_T has the same distribution in the reverse process p , p can be learned from the forward process q by simply reversing it. This is done as follows. Using the definition of the forward process in (4), Ho et al. [19] showed that

$$q(s_{t-1} | s_t, s_0) = \mathcal{N}(s_{t-1}; \tilde{\mu}_t(s_t, s_0), \tilde{\beta}_t I_d) \quad (13)$$

holds for any s_0 and s_t , where

$$\tilde{\mu}_t(s_t, s_0) = \frac{\sqrt{\bar{\alpha}_{t-1}} \beta_t}{1 - \bar{\alpha}_t} s_0 + \frac{\sqrt{\bar{\alpha}_t} (1 - \bar{\alpha}_{t-1})}{1 - \bar{\alpha}_t} s_t, \quad \tilde{\beta}_t = \frac{1 - \bar{\alpha}_{t-1}}{1 - \bar{\alpha}_t} \beta_t, \quad \bar{\alpha}_t = \prod_{\ell=1}^t \alpha_\ell. \quad (14)$$

Therefore, the latent variable in stage $t - 1$, s_{t-1} , is easy to sample when s_t and s_0 are known. To estimate s_0 , which is unknown when sampling from the reverse process, we use the forward process again. In particular, (4) implies that $s_t = \sqrt{\bar{\alpha}_t} s_0 + \sqrt{1 - \bar{\alpha}_t} \varepsilon_t$, where $\varepsilon_t \sim \mathcal{N}(\mathbf{0}_d, I_d)$ is a standard Gaussian noise. This identity can be rearranged as

$$s_0 = \frac{1}{\sqrt{\bar{\alpha}_t}} (s_t - \sqrt{1 - \bar{\alpha}_t} \varepsilon_t).$$

To obtain ε_t , which is unknown when sampling from p , we learn to regress it from s_t [19].

The regressor is learned as follows. Let $\varepsilon_t(\cdot; \psi)$ be a regressor of ε_t parameterized by ψ and $\mathcal{D} = \{s_0\}$ be a dataset of training examples. We sample s_0 uniformly at random from \mathcal{D} and then solve

$$\psi_t = \arg \min_{\psi} \mathbb{E}_q [\|\varepsilon_t - \varepsilon_t(S_t; \psi)\|_2^2] \quad (15)$$

per stage. The expectation is approximated by sampled s_0 . Note that we slightly depart from Ho et al. [19]. Since each regressor has its own parameters, the original optimization problem over T stages decomposes into T subproblems.

C Additional Experiments

This section contains four additional experiments.

C.1 Additional Synthetic Problems

In Section 6.2, we show results for three hand-selected problems out of six. We report results on the other three problems in Figure 4. We observe the same trends as in Section 6.2.

C.2 MNIST Experiment

The next experiment is on the MNIST dataset [31]. We start with learning an MLP-based multi-way classifier for digits and extract their $d = 8$ dimensional embeddings. These are used as features in

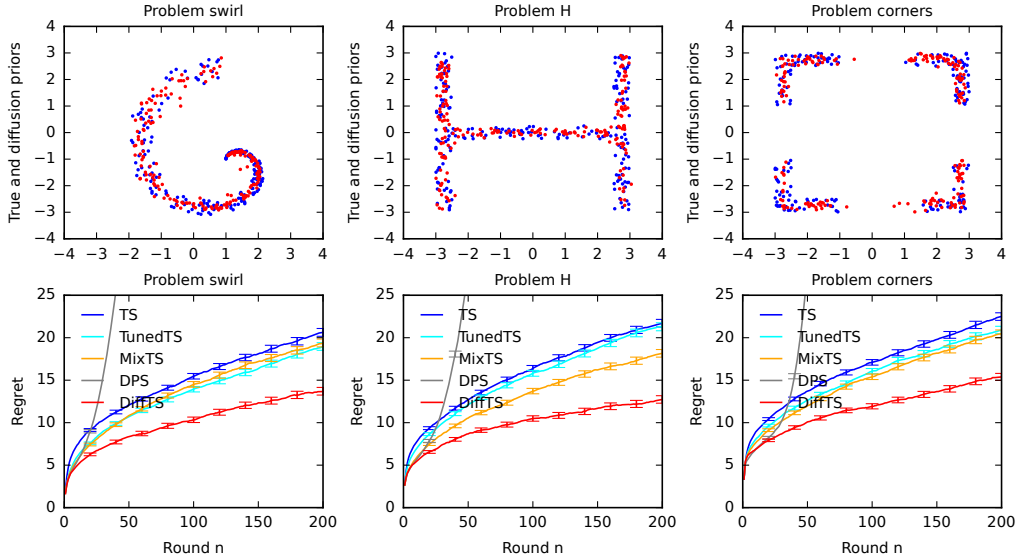


Figure 4: Evaluation of **DiffTS** on another three synthetic problems. The first row shows samples from the true (blue) and diffusion model (red) priors. The second row shows the regret of **DiffTS** and the baselines as a function of round n .




Figure 5: Evaluation of **DiffTS** on the MNIST dataset: (a) shows samples from the true (blue) and diffusion model (red) priors, (b) shows regret in the linear bandit, and (c) shows regret in the logistic bandit.

our experiment. We generate a distribution over model parameters θ_* as follows: (1) we choose a random positive label, assign it reward 1, and assign reward -1 to all other labels; (2) we subsample a random dataset of size 20, with 50% positive and 50% negative labels; (3) we train a linear model, which gives us a single θ_* . We repeat this 10 000 times and get a distribution over θ_* .

We consider both linear and logistic bandits. In both, the model parameter θ_* is initially sampled from the prior. In each round, $K = 10$ random actions are chosen randomly from all digits. In the linear bandit, the mean reward for a digit with embedding x is $x^\top \theta_*$ and the reward noise is $\sigma = 1$. In the logistic bandit, the mean reward is $g(x^\top \theta_*)$, where g is a sigmoid.

Our MNIST results are reported in Figure 5. We observe again that **DiffTS** has a lower regret than all baselines, because the learned prior captures the underlying distribution of θ_* well. We note that both the prior and diffusion prior distributions exhibit a strong cluster structure (Figure 5a), where each cluster represents one label.

C.3 Ablation Studies

We conduct three ablation studies on the cross problem in Figure 2.

In all experiments, the number of samples for training diffusion priors was 10 000. In Figure 6a, we vary it from 100 to 10 000. We observe that the regret decreases as the number of samples increases,




Figure 6: An ablation study of **DiffTS** on the cross problem: (a) we vary the number of samples for training the diffusion prior and report regret, (b) we vary the number of diffusion stages T and report regret, and (c) we vary the number of diffusion stages T and report computation time.

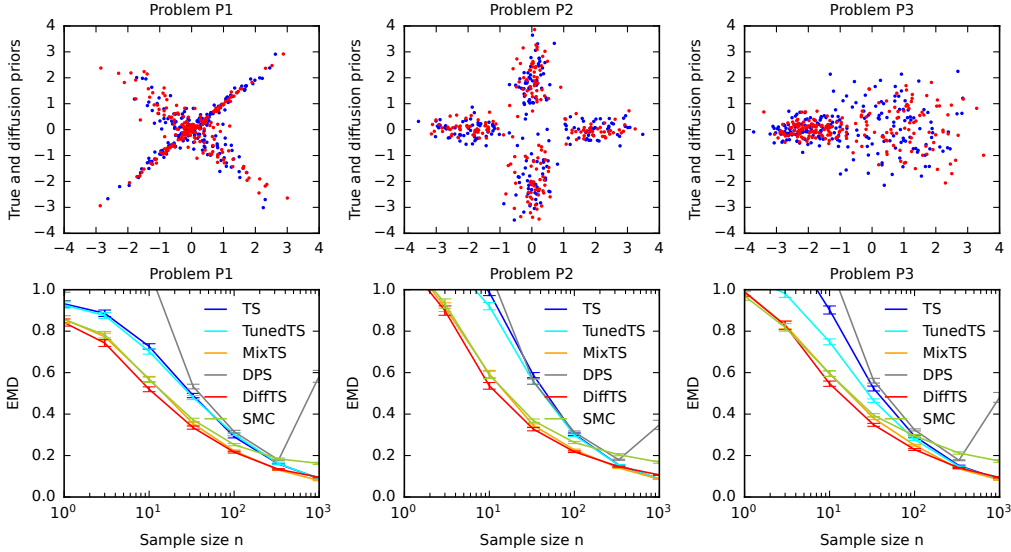


Figure 7: Evaluation on Gaussian mixture variants of the synthetic problems in Figure 2. The first row shows samples from the true (blue) and diffusion model (red) priors. The second row shows the earth mover’s distance of **DiffTS** and baseline posteriors from the true posterior as a function of sample size n .

due to learning a better prior approximation. The trend stabilizes around 3 000 training samples. We conclude that the quality of the learned prior approximation has a major impact on regret.

In all experiments, the number of diffusion stages was $T = 100$. In Figure 6b, we vary it from 1 to 300 and observe its impact on regret. While the regret at $T = 1$ is high, it decreases quickly as T increases. It stabilizes around $T = 100$, which we used in our experiments. In Figure 6c, we vary T from 1 to 300 and observe its effect on the computation time of posterior sampling. The time is linear in T , as suggested in Section 4.2. The main contributor is the neural network regressor.

C.4 Non-Bandit Evaluation

We use Gaussian mixture variants of the synthetic problems in Figure 2 for our non-bandit evaluation. The action in round k is chosen uniformly at random (not adaptively). Since the priors are Gaussian mixtures, the true posterior distribution can be computed in a closed form using **MixTS** and we can measure the distance of posterior approximations from it. We use the *earth mover’s distance (EMD)* between posterior samples from the true posterior and its approximation. We also considered the KL divergence. However, we could not apply it because the posteriors of **DiffTS** and **DPS** do not have analytical forms.

Algorithm 4 DPS of Chung et al. [12].

- 1: **Input:** Model parameters $\tilde{\sigma}_t$ and ζ_t
 - 2: Initial sample $S_T \sim \mathcal{N}(\mathbf{0}_d, I_d)$
 - 3: **for** stage $t = T, \dots, 1$ **do**
 - 4: $\hat{S} \leftarrow -\frac{\varepsilon_t(S_t; \psi_t)}{1 - \bar{\alpha}_t}$
 - 5: $\hat{S}_0 \leftarrow \frac{1}{\sqrt{\bar{\alpha}_t}}(S_t + (1 - \bar{\alpha}_t)\hat{S})$
 - 6: $Z \sim \mathcal{N}(\mathbf{0}_d, I_d)$
 - 7: $S_{t-1} \leftarrow \frac{\sqrt{\bar{\alpha}_{t-1}\beta_t}}{1 - \bar{\alpha}_t}\hat{S}_0 + \frac{\sqrt{\bar{\alpha}_t(1 - \bar{\alpha}_{t-1})}}{1 - \bar{\alpha}_t}S_t + \tilde{\sigma}_t Z - \zeta_t \nabla \sum_{\ell=1}^N (y_\ell - \phi_\ell^\top \hat{S}_0)^2$
 - 8: **Output:** Posterior sample S_0
-

We evaluate all methods from Figure 2. In addition, we implement a *sequential Monte Carlo (SMC)* sampler [16]. The initial particles are sampled uniformly at random from the prior. At each round, the particles are perturbed by a Gaussian noise. The standard deviation of the noise is initialized as a fraction of the observation noise and decays over time, as the posterior concentrates. The particles are weighted according to the likelihood of the observation in the round. Finally, we use normalized likelihood weights to resample the particles. The number of particles is 3 000 and we tune SMC to get good posterior approximations. The computational cost of SMC is comparable to DiffTS.

Our results are reported in Figure 7. We observe that DiffTS approximations are comparable to MixTS, which has an exact posterior in this setting. The second best performing method is SMC. Its approximations worsen as the sample size n increases. DPS approximations also get worse as n increases, which caused instability in Figure 2.

D Implementation of Chung et al. [12]

In our experiments, we compare to diffusion posterior sampling (DPS) with a Gaussian observation noise (Algorithm 1 in Chung et al. [12]). Our implementation is presented in Algorithm 4. The score is $\hat{S} = -\varepsilon_t(S_t; \psi_t)/(1 - \bar{\alpha}_t)$, where $\varepsilon_t(S_t; \psi_t)$ is a regression estimate of the forward process noise ε_t in Appendix B. We set $\tilde{\sigma}_t = \sqrt{\hat{\beta}_t}$, which is the same amount of noise as in our reverse process (Section 3). The term

$$\nabla \sum_{\ell=1}^N (y_\ell - \phi_\ell^\top \hat{S}_0)^2$$

is the gradient of the negative log-likelihood with respect to S_t .

As discussed in Appendices C.2 and D.1 of Chung et al. [12], ζ_t in DPS needs to be tuned for good performance. This is because $\nabla \sum_{\ell=1}^N (y_\ell - \phi_\ell^\top \hat{S}_0)^2$ grows with the number of observations, which causes instability. We also observed this in our experiments (Section 6.2). To make DPS work well, we follow Chung et al. [12] and set

$$\zeta_t = \frac{1}{\sqrt{\sum_{\ell=1}^N (y_\ell - \phi_\ell^\top \hat{S}_0)^2}}.$$

While this significantly improves the performance of DPS, it does not prevent failures. The fundamental problem is that gradient-based optimization is sensitive to the step size, especially when the optimized function is steep. Note that LaplaceDPS does not have any such hyper-parameter.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope?

Answer: [Yes]

Justification: The abstract and introduction clearly state all contributions. The introduction also points to where those contributions are made.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: The increase in computational cost is discussed in Section 4.2 and shown empirically in Section 6.2. We also conduct an ablation study in Appendix C.3, where we show how the regret of DiffTS scales with the number of samples used for pre-training the prior and the number of stages in the diffusion model prior.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: The main claims are stated and discussed in Section 4. Their proofs are in Appendix A.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We also include code to reproduce the synthetic results in Figures 2 and 4.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We include code to reproduce the synthetic results in Figures 2 and 4.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: The experiments are described to a sufficient level to be reproducible. To make sure, we include code to reproduce the synthetic results in Figures 2 and 4.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: All plots in the paper have error bars.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [No]

Justification: Our experiments are not large scale.

9. **Code Of Ethics**

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

Answer: [Yes]

Justification: We checked the link and comply.

10. **Broader Impacts**

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: This work is algorithmic and not tied to a particular application that would have immediate negative impact.

11. **Safeguards**

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: This paper does not pose such a risk.

12. **Licenses for existing assets**

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: All used assets are stated and cited.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: This paper does not release new assets.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: No crowdsourcing or research with human subjects.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: No crowdsourcing or research with human subjects.