

Gradient Guidance for Diffusion Models: An Optimization Perspective

Yingqing Guo[†] Hui Yuan[†] Yukang Yang Minshuo Chen Mengdi Wang^{*}
Princeton University

April 24, 2024

Abstract

Diffusion models have demonstrated empirical successes in various applications and can be adapted to task-specific needs via guidance. This paper introduces a form of gradient guidance for adapting or fine-tuning diffusion models towards user-specified optimization objectives. We study the theoretic aspects of a guided score-based sampling process, linking the gradient-guided diffusion model to first-order optimization. We show that adding gradient guidance to the sampling process of a pre-trained diffusion model is essentially equivalent to solving a regularized optimization problem, where the regularization term acts as a prior determined by the pre-training data. Diffusion models are able to learn data’s latent subspace, however, explicitly adding the gradient of an external objective function to the sample process would *jeopardize* the structure in generated samples. To remedy this issue, we consider a modified form of gradient guidance based on a forward prediction loss, which leverages the pre-trained score function to preserve the latent structure in generated samples. We further consider an iteratively fine-tuned version of gradient-guided diffusion where one can query gradients at newly generated data points and update the score network using new samples. This process mimics a first-order optimization iteration in expectation, for which we proved $\mathcal{O}(1/K)$ convergence rate to the global optimum when the objective function is concave.

1 Introduction

Diffusion models have emerged as a significant advancement in the field of generative artificial intelligence, offering state-of-the-art performance in image generation (Song and Ermon, 2019; Song et al., 2020a; Dhariwal and Nichol, 2021). These models operate by gradually transforming a random noise distribution into a structured output, by using a score function trained from large amounts of data. Such a transforming process is typically modeled as a stochastic differential equation, offering a mathematically grounded approach for sampling. One of the key advantages of diffusion models is their ability to be guided or fine-tuned for specific tasks, which allows them to excel in a wide range of applications (Kong et al., 2020; Ajay et al., 2022; Gruver et al., 2023).

Guidance-based diffusion, a nuanced extension of diffusion models, stands at the forefront of controlled generation in generative AI. This approach involves steering the noise transformation process of a diffusion model towards desired outcomes by incorporating additional “guidance signals”. This guidance can manifest in various forms, such as text prompts, class labels, or even

^{*}† Equal contribution. Emails: {yg6736, huiyuan, yy1325, minshuochen, mengdiw}@princeton.edu.

conditioning on specific attributes and rewards. The core principle behind this technique is to influence the probabilistic pathway of the noise transformation process at each time step, thereby steering the final output towards predefined criteria or objectives. This controlled generation capability opens up opportunities for generative AI in a broad range of tasks, such as in targeted image synthesis, content creation with specific themes, or even in drug design where molecular structures need to meet specifications.

A notable example is the classifier-based diffusion model introduced by Song et al. (2020c); Dhariwal and Nichol (2021), which generates data conditioned on a class label, via guidance signals that are computed from conditional likelihoods from a classifier. Building on this concept, Bansal et al. (2023) extend the classifier-guidance method to a form of “universal guidance”. Such guidance allows the generation process to be influenced by gradient obtained from some external loss function, effectively tailoring the diffusion process to meet specific objectives (Chung et al., 2022a,b; Graikos et al., 2022; Kawar et al., 2022; Lugmayr et al., 2022; Wang et al., 2022). Despite of numerous empirical successes, there remain significant gaps in the theoretical understanding and guarantees associated with guided diffusion models.

Problem and Challenges Suppose we have a pre-trained diffusion model that can generate new samples faithfully from the pre-training data’s distribution and maintain the data’s latent structure. The goal is to adapt this diffusion model to generate new samples that optimize task-specific objectives, while *maintaining the learned structure* in new samples. Compared to classic optimization, the guided diffusion model offers new possibilities to optimize complex design variables such as images, videos, proteins, and genomes (Black et al., 2023; Watson et al., 2023; Liu et al., 2024). Interested readers may refer to recent surveys for a more comprehensive exposure (Yang et al., 2023; Chen et al., 2024; Guo et al., 2023).

To adapt pre-trained diffusion models, existing practical methods largely rely on empirical heuristics and hyperparameter tuning. There remain critical theoretical questions: (i) Why does naively guiding diffusion models using gradient never work in practice? (ii) How to add a guidance signal to improve the target objective without compromising the quality of the generated output? (iii) Can one guarantee the properties of new samples generated by guided diffusion? (iv) What are the limits of adaptability in these guided models? This paper aims to answer these questions from an optimization perspective.

Scope of This Paper We investigate the role of guidance in diffusion models from an optimization perspective. The goal is to generate samples that optimize a given objective function f . Drawing inspiration from gradient-based optimization methods, we construct a guidance signal based on the gradient vector, ∇f . Then we use the gradient signal, in addition to the pre-trained score function, to guide the sampling process towards generating structured output with higher function values. See Figure 1 for illustration our algorithmic framework. Our main results are summarized as follows:

- We focus on structured data. Assume that pre-trained data belongs to a latent low-dimensional subspace, thus the trained score function is capable of discerning and maintaining the latent subspace structure of data during (unguided) generation. To guide the generation, we introduce a gradient-like guidance based on a forward prediction loss (Definition 1). This gradient guidance can be computed based on a pre-trained score network, and it provably preserve any learned low-dimensional structure in the generated output (Lemma 1).

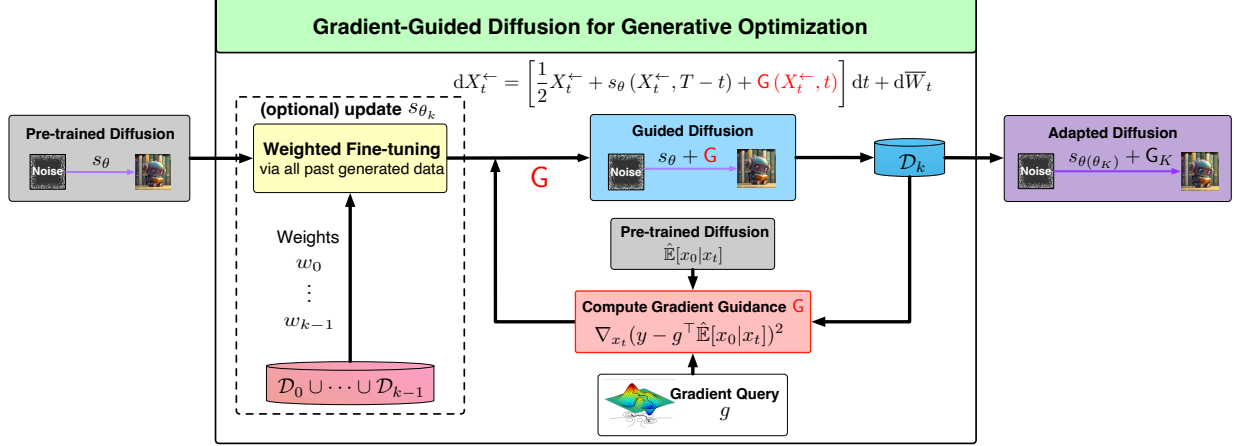


Figure 1: **Gradient-guided diffusion model for generative optimization, with or without adaptive finetuning.** A pre-trained diffusion model is guided with an additional gradient signal from an external objectives function towards generating near-optimal solutions.

- We formalize a mathematical framework of using a gradient-guided diffusion model for generative optimization. While retaining the pre-trained score function, the generation process is iteratively refined and guided using new gradient queries (Algorithm 1; Figure 1 without fine-tuning). Under proper assumptions, we demonstrate that this adapted model generates novel samples whose expectation converges to a solution that is regularized with respect to the original problem (Theorem 2). The regularization ensures that the generated samples remain proximal to the training data. In other words, gradient guidance cannot shift data distribution unboundedly towards higher objective values, revealing a fundamental limit for adapting pre-trained diffusion models.
- Furthermore, we explore an adaptive variant of gradient-guided diffusion, where both the score function and gradient guidance are iteratively fine-tuned using self-generated samples (Algorithm 2; Figure 1 with fine-tuning). Although slightly increasing the computational demand, we provide evidence that this approach generates new samples whose expectation converges to global optima, within the latent subspace, at a rate of $\mathcal{O}(1/K)$ (Theorems 3 and 4), where K denotes the number of iterations and gradient evaluations, matching classical convergence theory of convex optimization.

Our findings suggest that this novel gradient guidance not only preserves the latent subspace structure of the data but also ensures fast convergence towards the optimal solution. Numerical experiments with a pre-trained U-network score function are provided in Section 7 to support these theoretical findings.

2 Related Work

Our study is motivated by recent empirical progress of guidance-based diffusion model fine-tuning for steering sample generation towards specific needs (Dhariwal and Nichol, 2021; Bansal et al., 2023). Upon modeling the specific needs as a reward function, the relevant methods can be summarized into two categories in the sequel.

Guided Generation and Fine-tuning Given an auxiliary reward function judging the sample property of interests, existing research explored diverse mechanisms to guide generation from diffusion models. There are mainly two types of methods. The first type of method incorporates an additive so-called “**guidance**” term into the score function of pre-trained diffusion models at inference time. For example, classifier guidance (Song et al., 2020a; Dhariwal and Nichol, 2021) defines the guidance term as the gradient of an externally trained classifier on noise corrupted data. Classifier-free guidance (Ho and Salimans, 2022) simultaneously trains conditional and unconditional diffusion models, circumventing the training of an external classifier. After the training, the score functions of the conditional and unconditional models are combined together to achieve guided generation. Bansal et al. (2023) draw motivation from classifier guidance and generalize the idea to a “universal guidance” for adapting unconditioned score functions to various external rewards.

The second type of method attempts to directly fine-tune the weight parameters in a pre-trained diffusion model by interacting with the target reward function. For example, Clark et al. (2023) fine-tune diffusion models by directly backpropagating the gradient of the reward function. Recently, a line of works utilizes Reinforcement Learning (RL) techniques for fine-tuning diffusion models (Black et al., 2023; Fan et al., 2023). They formulate the sample generation process of diffusion models as a finite-horizon Markov chain. The score function can be viewed as a policy, the generated samples are the state of the Markov chain, and the target reward function defines the terminal reward. In this way, fine-tuning diffusion models is equivalent to policy optimization and allows the use of policy gradient methods.

Sampling and Statistical Theory of Diffusion Model In contrast to the fruitful empirical advances, the theory of diffusion models is still limited. To the best of our knowledge, a theoretical understanding of fine-tuning diffusion models is absent. Existing results mainly focus on the sampling ability and statistical properties of unconditional diffusion models. In particular, for sampling ability, a line of works shows that the distribution generated by a diffusion model is close to the data distribution, as long as the score function is accurately estimated (De Bortoli et al., 2021; Albergo et al., 2023; Block et al., 2020; Lee et al., 2022a; Chen et al., 2022; Lee et al., 2022b,a; Chen et al., 2022; Lee et al., 2022b). The accuracy of the estimated score function is measured in terms of an L_∞ or L_2 -norm distance. More recently, Chen et al. (2023c,b); Benton et al. (2023) develop refined and tighter analyses using Taylor expansions of the discretized backward process and localization method. It is worth mentioning that the analysis in Chen et al. (2023c,b); Benton et al. (2023) extends to broader sample generation processes such as deterministic ones based on probabilistic ODEs. Going beyond distributions in Euclidean spaces, De Bortoli (2022) analyzes diffusion models for sampling distribution supported on a low-dimensional manifold. Moreover, Montanari and Wu (2023) consider sampling from symmetric spiked models, and El Alaoui et al. (2023) study sampling from Gibbs distributions using diffusion processes.

Turning towards the statistical theory of diffusion models, Song et al. (2020b) and Liu et al. (2022) provide asymptotic analyses, assuming a parametric form of the score function. Unfortunately, asymptotic analysis does not lead to concrete sample complexities. Later, concurrent works, Oko et al. (2023) and Chen et al. (2023a), establish sample complexity bounds of diffusion models for estimating nonparametric data distributions. In high dimensions, their results highlight a curse of dimensionality issue without further assumptions, which also appears in Wibisono et al. (2024) considering kernel methods. More interestingly, these works demonstrate that diffusion models can

circumvent the curse of dimensionality issue if the data has low-dimensional structures. In the same spirit, [Mei and Wu \(2023\)](#) investigate learning high-dimensional graphical models using diffusion models, without the curse of dimensionality. For conditional diffusion models, [Yuan et al. \(2023\)](#); [Fu et al. \(2024\)](#) establish sample complexity bounds for learning generic conditional distributions. We refer readers to [Chen et al. \(2024\)](#) for an overview of contemporary theoretical progress.

Novelty of This Paper Despite the existing theoretical underpinnings of diffusion models, our paper provides the first rigorous study of adapting and fine-tuning diffusion models using gradient guidance from an optimization perspective. Specifically, we first understand why naive gradient guidance does not lead to meaningful optimization performance. Built upon the insights gained from the analysis, we propose gradient guidance that is proven to preserve generated data structures and simultaneously achieve strong optimization guarantees on adapted variants of diffusion models.

We are aware of two recent works ([Uehara et al., 2024](#); [Marion et al., 2024](#)) studying adapting the output distribution of diffusion models to a target reward function. In particular, they define a reward function with respect to the output distribution of the diffusion model. Given a pre-trained diffusion model, for instance, [Uehara et al. \(2024\)](#) utilizes a KL-divergence regularizer penalizing the deviation to the pre-trained model for preventing overfitting in fine-tuning. Through some explicit computation, [Uehara et al. \(2024\)](#); [Marion et al. \(2024\)](#) identify the proper guidance term to adapt the pre-trained model. Yet a sophisticated estimation procedure is needed to find the guidance term in [Uehara et al. \(2024\)](#); our gradient guidance enjoys much simplicity and efficacy, as we demonstrated in both theory and experiments.

3 Preliminaries: Diffusion Models

Score-based diffusion models capture the distribution of pre-training data by learning a sequence of transformations to generate new samples from noise ([Song et al., 2020c](#)). A forward stochastic process progressively adds noise to data, whose sample trajectories are used to train the score function. To generate new samples, a backward denoising process starts from sampling pure noise and gradually transforms the noise guided by the learned score function.

Forward Process The forward process of diffusion models initializes with $X_0 \in \mathbb{R}^D$, a random variable drawn from the pre-training data \mathcal{D} . It introduces noise to via an Ornstein-Uhlenbeck process, i.e.,

$$dX_t = -\frac{1}{2}q(t)X_t dt + \sqrt{q(t)}dW_t \quad \text{for } q(t) > 0, \quad (1)$$

where $(W_t)_{t \geq 0}$ is a standard Wiener process, and $q(t)$ is a non-decreasing weighting function. X_t represents the noise-corrupted data distribution at time t . Given $X_0 = x_0$, the conditional distribution $X_t|X_0 = x_0$ is Gaussian, i.e., $\mathcal{N}(\alpha(t)x_0, h(t)I_D)$ with $\alpha(t) = \exp(-\int_0^t \frac{1}{2}q(s)ds)$ and $h(t) = 1 - \alpha^2(t)$. In practice, the forward process will terminate at a large time T so that the marginal distribution of X_T is close to $\mathcal{N}(0, I_D)$. In our analysis, we take $q(t) \equiv 1$ without loss of generality, where $\alpha(t) := \exp(-t/2)$ and $h(t) := 1 - \exp(-t)$.

Backward Process If reversing the time of the forward process, we can reconstruct the original distribution of the data from pure noise. With $(\bar{W}_t)_{t \geq 0}$ being another independent Wiener process,

the backward SDE below (Anderson, 1982) reverses the time in the forward SDE (1),

$$dX_t^{\leftarrow} = \left[\frac{1}{2} X_t^{\leftarrow} + \underbrace{\nabla \log p_{T-t}(X_t^{\leftarrow})}_{\text{score}} \right] dt + d\bar{W}_t. \quad (2)$$

Here $p_t(\cdot)$ denotes the marginal density of X_t in the forward process. In the forward SDE (2), the *score function* $\nabla \log p_t(\cdot)$ plays a crucial role, but it has to be estimated from data.

Score Matching To learn the unknown score function $\nabla \log p_t(\cdot)$, it is common to train a score network $s_\theta(x, t)$ using samples generated by the forward process. Let \mathcal{D} denote the data for training. Then the score network is learned by minimizing the following loss:

$$\min_{s \in \mathcal{S}} \int_0^T \mathbb{E}_{x_0 \in \mathcal{D}} \mathbb{E}_{x_t | x_0} \left[\|\nabla_{x_t} \log \phi_t(x_t | x_0) - s(x_t, t)\|^2 \right] dt, \quad (3)$$

where \mathcal{S} is a given function class, $\mathbb{E}_{\mathcal{D}}$ denotes the empirical expectation over training data \mathcal{D} and $\mathbb{E}_{x_t | x_0}$ denotes condition expectation over the forward process, $\phi_t(x_t | x_0)$ is the Gaussian transition kernel, i.e., $\frac{1}{(2\pi h(t))^{D/2}} \exp(-\frac{\|x_t - \alpha(t)x_0\|^2}{2h(t)})$.

Generation and Guided Generation Given a pre-trained score function s_θ , one generates new samples by simulating the backward process (2) with the true score replaced by s_θ . Further, one can add additional guidance to the backward SDE to steer its output distribution towards specific properties of interest. Module 1 formalizes the generation process and guided generation process using a pre-trained diffusion model.

Module 1 Guided_BackwardSample(s_θ, \mathbf{G})

- 1: **Input:** Score s_θ , guidance \mathbf{G} default to be zero for unguided generation.
- 2: **Hyper-parameter:** T .
- 3: Initialized at $X_t^{\leftarrow} \sim \mathcal{N}(0, I)$, simulate the following SDE till time T :

$$dX_t^{\leftarrow} = \left[\frac{1}{2} X_t^{\leftarrow} + s_\theta(X_t^{\leftarrow}, T-t) + \mathbf{G}(X_t^{\leftarrow}, T-t) \right] dt + d\bar{W}_t.$$

- 4: **Output:** Sample X_T^{\leftarrow} .
-

Conditional Generation Suppose the goal is to generate X with a desired property $Y = y$ from the distribution $P(X|Y = y)$. To this end, one needs the **conditional score function** $\nabla_{x_t} \log p_t(x_t | y)$, as a replacement of the unconditioned score $\nabla_{x_t} \log p_t(x_t)$. The Bayes rule gives

$$\nabla_{x_t} \log p_t(x_t | y) = \underbrace{\nabla \log p_t(x_t)}_{\text{est. by } s_\theta(x_t, t)} + \underbrace{\nabla_{x_t} \log p_t(y | x_t)}_{\text{to be est. by guidance}}. \quad (4)$$

When a pre-trained score network $s_\theta(x_t, t) \approx \nabla \log p_t(x_t)$, the remaining task is to estimate $\nabla_{x_t} \log p_t(y | x_t)$ and add it as a “guidance” \mathbf{G} to the backward process (Module 1).

Classifier and Classifier-Free Guidance Classifier guidance (Song et al., 2020c; Dhariwal and Nichol, 2021) is a approach for sampling from $P(X|Y = y)$ when Y is a discrete label. This method estimates $\nabla_{x_t} \log p_t(y | x_t)$ by training auxiliary classifiers, denoted as $\hat{p}(y | x_t, t)$, and then computing the gradient of the classifier logits as the guidance, i.e., $\mathbf{G}(x_t, t) = \nabla_{x_t} \log \hat{p}(y | x_t, t)$. An alternative is the classifier-free guidance method (Ho and Salimans, 2022), which jointly trains a conditional and an unconditional diffusion model, and combine the two score estimates via a form of guidance to generate samples.

Notations For a random variable X , P_x represents its distribution, and $p(x)$ denotes its density function. For X, Y jointly distributed, $P(X | Y = y)$ denotes the conditional distribution, and $p(x | y)$ denotes its density function. We use the notation $\mathbb{E}[x | y]$ for the conditional expectation. Let \mathcal{D} be the pre-training data, and let $\mathbb{E}_{\mathcal{D}}$ be the empirical expectation over \mathcal{D} . Let $\bar{\mu}$ and $\bar{\Sigma}$ denote the data’s empirical mean and covariance matrix, i.e., $\bar{\mu} := \mathbb{E}_{x \in \mathcal{D}}[x]$ and $\bar{\Sigma} := \mathbb{E}_{x \in \mathcal{D}}[(x - \bar{\mu})(x - \bar{\mu})^\top]$. For a matrix A , we denote by $\text{Span}(A)$ the subspace spanned by its column vectors. For a square matrix A , we denote by A^{-1} its inverse or Moore–Penrose inverse. For any differentiable function $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$, $\nabla f \in \mathbb{R}^{m \times n}$ denotes Jacobian matrix, i.e., $(\nabla f)_{ij} = \frac{\partial f_i(x)}{\partial x_j}$.

4 A Primer on Gradient Guidance

Suppose we have a pre-trained diffusion model where the score network $s_\theta(x_t, t)$ provides a good approximation to the true score function $\log p(x_t)$. Then this diffusion model is viewed as an implicit density estimator of the pre-training data’s distribution. Its backward process (2) generates samples from this estimated distribution (Oko et al., 2023; Chen et al., 2023a).

Now suppose we want to generate novel samples with desired properties that can be measured by a differentiable function f . We will refer to f as a reward or objective function later on, and it is often user-specified. Motivated by the gradient methodology in optimization, a natural, intuitive way for adding guidance is to steer the generated samples towards the steepest ascent direction of f (Bansal et al., 2023; Clark et al., 2023). This motivates the following guided backward process (Module 1):

$$dX_t^{\leftarrow} = \left[\frac{1}{2} X_t^{\leftarrow} + s_\theta(X_t^{\leftarrow}, T - t) + \mathbf{G}(X_t^{\leftarrow}, t) \right] dt + d\bar{W}_t.$$

Here the guidance term \mathbf{G} is what we focus on and wish to design. Specifically, we want to construct this guidance term \mathbf{G} based on the gradient ∇f of a general objective f .

4.1 Subspace Data and Score Decomposition

Real-world data often has rich intrinsic structures. These structures can be induced by local regularities, global symmetries, and repetitive patterns (Tenenbaum et al., 2000; Roweis and Saul, 2000) and are often low-dimensional (Pope et al., 2021). The power of diffusion models is to model the latent distribution and generate novel samples that preserve important characteristics of real-world data. If we blindly improve f at the cost of losing these characteristics, the quality of new samples would degrade dramatically. This quality degradation, also known as “reward over-optimization”, is a common challenge for adapting diffusion models towards an external reward (Yuan et al., 2023; Uehara et al., 2024).

We aim to design gradient guidance to improve objective function while mitigating the risk of over-optimization. To this end, we focus on data that admits a low-dimensional latent subspace. Let us make the following assumption.

Assumption 1 (Subspace Data). *Data $X \in \mathbb{R}^D$ can be represented as $X = AU$, where $A \in \mathbb{R}^{D \times d}$ is an unknown matrix and the latent variable $U \in \mathbb{R}^d$ follows some distribution P_u with a density p_u . Here $d \ll D$. We assume the empirical covariance of U is full rank.*

Under Assumption 1, the score function $\nabla \log p_t(x)$ decomposes to two orthogonal parts: an on-support component belonging to the subspace; and an orthogonal component. We recall this key result in Proposition 1.

Proposition 1 (Score Decomposition for Subspace Data (Chen et al. (2023a) Lem. 1, Thm. 3)). *Under Assumption 1, the score function $\nabla \log p_t(x)$ decomposes as*

$$\nabla \log p_t(x) = \underbrace{A \nabla \log p_t^{\text{LD}}(A^\top x)}_{\mathbf{s}_{\parallel}(A^\top x, t): \text{on-support score}} - \underbrace{\frac{1}{h(t)} (I_D - AA^\top)}_{\mathbf{s}_{\perp}(x, t): \text{ortho. score}} x. \quad (5)$$

where $p_t^{\text{LD}}(u') = \int \phi_t(u'|u) p_u(u) du$ with $\phi_t(\cdot|u)$ being the density of $\mathcal{N}(\alpha(t)u, h(t)I_d)$ for the same $\alpha(t)$ and $h(t)$ in the forward process (1).

According to Chen et al. (2023a), pre-training a score function on subspace data takes advantage of the decomposition given by (5) and learns the latent subspace. When the pre-trained score network is used to generate new samples, the backward sampling process also decomposes into two orthogonal processes due to (5). Analysis of this backward process proves that the generated output would remain proximal to the latent subspace. This explains why diffusion models can learn and preserve data’s underlying characteristics. We refer interested readers to Chen et al. (2023a) for more discussions.

In the rest of this section, we investigate the principles for designing a guidance based on the gradient of f that ensures generated samples (i) improve the value of f , and at the same time, (ii) adhere to the subspace structure, i.e. generated samples being close to the subspace spanned by A .

4.2 Naive Gradient Does’t Work as Guidance

Motivated by the gradient optimization methodology, a natural, intuitive way for adding guidance is to steer the generated samples towards the steepest ascent direction of f (Bansal et al., 2023; Clark et al., 2023). Therefore, a tempting simple choice of the guidance \mathbf{G} is the steepest ascent direction, which we refer to as *naive gradient guidance* i.e.,

$$\mathbf{G}(X_t^{\leftarrow}, t) \propto \nabla f(X_t^{\leftarrow}). \quad (6)$$

This naive choice of guidance signal \mathbf{G} would steer the movement of the original backward process towards the direction that increases f .

However, the naive gradient guidance (6) is never adopted in practice; existing methods have to resort to more sophisticated forms of guidance or more computationally demanding fine-tuning methods; for example Bansal et al. (2023); Uehara et al. (2024); Marion et al. (2024). Introducing gradient information indiscriminately into the backward SDE has the risk of potentially leading the

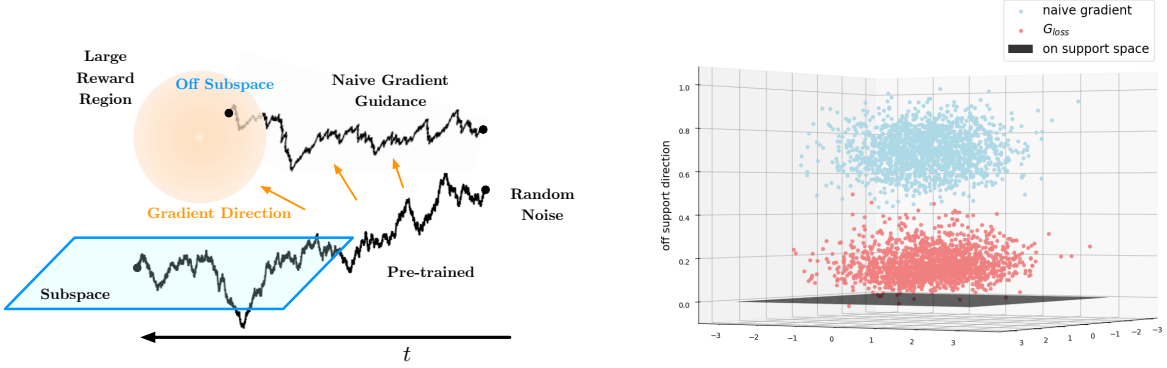


Figure 2: **Directly adding the gradient of the objective function to the backward sampling process sabotages the subspace structure.** Left: When gradients are pointing out of the data subspace, adding them directly to the backward SDE will make samples go off the subspace. Right: Numerical experiments show that naive gradients lead to substantially larger off-subspace error compared to our gradient guidance G_{loss} (Definition 1); see Section 7 for experiment details.

stochastic denoising process to divergence, and compromising the data structures learned during pre-training.

Let us suppose the data distribution is supported on a low-dimensional subspace as in Assumption 1. We explain why naive gradients do not work as guidance. With subspace data, the score function would steer the distribution towards concentrating onto the latent subspace, due to its special decomposition form given by Proposition 1. We refer interested readers to Chen et al. (2023a) for detailed analysis of this phenomenon.

However, naive gradient vectors can be pointing towards any direction, not limited to within the latent subspace. Thus directly adding gradient guidance to the backward process could jeopardize the decomposition form of the score function, and it would also jeopardize the latent structure in the generated output. See Figure 2 for illustration and experiment results. The failure of naive gradient motivates us to seek robust alternatives.

4.3 Motivating Gradient Guidance from Conditional Score Function

We want to study how to add guidance to the sampling process utilizing the gradient of f . To motivate our design of guidance, we start with the most elementary Gaussian probabilistic model. Later we will drop this assumption and consider general data distributions and general f .

Assumption 2 (Gaussian model). *Let data follow a Gaussian distribution, i.e., $X \sim \mathcal{N}(\mu, \Sigma)$, and let $f(x) = g^\top x$ be a linear function for some $g \in \mathbb{R}^D$. Let $Y = f(X) + \epsilon$ with independent, identically distributed noise $\epsilon \sim \mathcal{N}(0, \sigma^2)$ for some $\sigma > 0$.*

To generate samples from $P(X|Y = y)$, we need to train a diffusion model with a conditional score function. By the Bayes' rule, the conditional score function takes the form of a sum given by

$$\nabla_{x_t} \log p_t(x_t | y) = \underbrace{\nabla \log p_t(x_t)}_{\text{est. by } s_\theta(x_t, t)} + \underbrace{\nabla_{x_t} \log p_t(y | x_t)}_{\text{to be est. by guidance}}. \quad (\text{recall (4)})$$

Now if we already have a pre-trained score function s_θ , the remaining task is to estimate the second

term $\log p_t(y | x_t)$. Under the Gaussian assumption, we derive the following closed-form conditional score. The proof is provided in Appendix B.1.

Lemma 1 (Conditional score gives a gradient-like guidance). *Under Assumption 2, we have*

$$\nabla_{x_t} \log p_t(y | x_t) = \beta(t) \left[y - g^\top \mathbb{E}[x_0 | x_t] \right] \cdot (\alpha^2(t)\Sigma + h(t)I_D)^{-1} \Sigma g, \quad (7)$$

where $\mathbb{E}[x_0 | x_t]$ denotes the conditional expectation of x_0 given x_t in the forward process (1), $\alpha(t) = e^{-t/2}$, $h(t) = 1 - e^{-t}$ as in (1), and $\beta(t) = \alpha(t)/(\sigma^2 + g^\top \Sigma^{-1} (I_D + \alpha^2(t)/h(t) \cdot \Sigma)^{-1} g)$.

Observe that, when $\Sigma = I$, (7) suggests the following form of guidance that is aligned with the naive gradient, i.e., the steepest ascent direction:

$$\mathcal{G}(x_t, t) \propto \left[y - g^\top \mathbb{E}[x_0 | x_t] \right] \cdot g.$$

However, even for Gaussian distributions, as long as $\Sigma \neq I$, the term of (7) is no longer proportional to g but becomes a pre-conditioned version of the gradient.

Another observation is that this guidance scales with a residual term $y - g^\top \mathbb{E}[x_0 | x_t]$. In particular, the residual term $y - g^\top \mathbb{E}[x_0 | x_t]$ tunes the *strength of guidance*. Recall $\mathbb{E}[x_0 | x_t]$ denotes the posterior expectation of clean data x_0 given x_t in the forward process. Thus, in a backward view, $\mathbb{E}[x_0 | x_t]$ coincides with the expected sample to be generated conditioned on x_t . In this sense, the quantity $y - g^\top \mathbb{E}[x_0 | x_t]$ measures a **look-ahead gap** between the expected reward of generated samples and the target value. A larger absolute value of the residual means stronger guidance in the backward generation process.




Figure 3: **Plot of $\beta(t), \alpha(t), h(t)$ for $t \in [0, 10]$ when $\Sigma = I$.**

We plot the theoretical choice of $\beta(t)$ and $\alpha(t), h(t)$ to t in Figure 3. In practice, the choice of $\alpha(t), h(t)$ can vary and they are determined by the forward process used for pre-training; and $\beta(t)$ can be treated as a tuning parameter to adjust the strength of guidance.

4.4 Construct Gradient Guidance to Preserve Latent Subspace

When the data distribution is supported on a latent subspace, directly adding gradient guidance to the backward sampling process could jeopardize the data’s latent structure. We saw that this would lead to over-optimization, as illustrated in Figure 2.

To remedy such an issue, we propose the following modification to the gradient guidance. This modified gradient guidance takes advantage of a given score function.

Definition 1 (Gradient Guidance of Look-Ahead Loss). *Given a gradient vector g , define the gradient guidance of look-ahead loss as*

$$\mathcal{G}_{loss}(x_t, t) := -\beta(t) \cdot \nabla_{x_t} \left(y - g^\top \mathbb{E}[x_0 | x_t] \right)^2, \quad (8)$$

where $\beta(t) > 0, y \in \mathbb{R}$ are tuning parameters, and $\mathbb{E}[x_0 | x_t]$ is the conditional expectation of x_0 given x_t in the forward process (1), i.e., $dX_t = -\frac{1}{2}q(t)X_t dt + \sqrt{q(t)} dW_t$.

The formula of (8) generalizes the intuition of a conditional score to work with any data distribution and objective function. The **look-ahead loss** $(y - g^\top \mathbb{E}[x_0|x_t])^2$ resembles the proximal term commonly used in first-order proximal optimization methods. It is worth noting that \mathbf{G}_{loss} coincides with the forward universal guidance $\nabla_{x_t} \ell(y, f(\hat{\mathbb{E}}[x_0|x_t]))$ proposed by Bansal et al. (2023) ((8) in their paper) when ℓ is the square loss and $f = g^\top x$.

When the pre-training data distribution is Gaussian, the gradient guidance (8) is equivalent to Lemma 1 equation (7). This equivalence is a side-product from the proof of Lemma 1 and we provide a sketch here (see details in Appendix B.1). Given the probabilistic model Assumption 2, $\nabla_{x_t} \log p_t(y|x_t)$, the quantity to be estimated by guidance, is the score of a Gaussian distribution $\mathcal{N}(m_y(x_t), \sigma_y^2(x_t))$, with $m_y(x_t)$ and $\sigma_y^2(x_t)$ being mean and variance of the conditional distribution $Y | X_t = x_t$ respectively, i.e.,

$$\nabla_{x_t} \log p_t(y | x_t) = -\nabla_{x_t} \left[\frac{1}{2} \left(\frac{y - m_y(x_t)}{\sigma_y(x_t)} \right)^2 \right] - \nabla_{x_t} \log \sigma_y(x_t), \quad (9)$$

with $m_y(x_t) = g^\top \mathbb{E}[x_0 | x_t]$ and $\sigma_y(x_t)$ not depending on x_t . Thus we see \mathbf{G}_{loss} is equivalent to (7).

A key advantage of \mathbf{G}_{loss} is that it enables preserving the subspace structure, for **any** data distribution under Assumption 1. This result is formally stated in the following theorem, we provide a proof sketch here and the full proof is in Appendix B.2.

Theorem 1 (Faithfulness of \mathbf{G}_{loss} to the Low-Dimensional Subspace of Data). *Under Assumption 1, it holds for any data distribution and $g \in \mathbb{R}^D$ that*

$$\mathbf{G}_{loss}(x_t, t) \in \text{Span}(A). \quad (10)$$

Proof Sketch We have

$$\nabla_{x_t} \left(y - g^\top \mathbb{E}[x_0|x_t] \right)^2 \propto \nabla_{x_t} \mathbb{E}[x_0|x_t]^\top g.$$

Note here that $\nabla_{x_t} \mathbb{E}[x_0|x_t]$ is the Jacobian matrix of $\mathbb{E}[x_0|x_t]$, which is a mapping from \mathbb{R}^D to \mathbb{R}^D . We will show that the Jacobian $\nabla_{x_t} \mathbb{E}[x_0|x_t]$ maps any vector $g \in \mathbb{R}^D$ to $\text{Span}(A)$.

To see this, we utilize the score decomposition result of Proposition 1 which is

$$\nabla \log p_t(x_t) = A \nabla \log p_t^{\text{LD}}(A^\top x_t) - \frac{1}{h(t)} \left(I_D - AA^\top \right) x_t. \quad (\text{recall (5)})$$

Plugging (5) into the equality $\mathbb{E}[x_0|x_t] = \frac{1}{\alpha(t)} (x_t + h(t) \nabla \log p_t(x_t))$ (Tweedie's formula (Efron, 2011)), we have

$$\mathbb{E}[x_0|x_t] = \frac{1}{\alpha(t)} \left(x_t + h(t) \left[A m(A^\top x_t) - \frac{1}{h(t)} x_t \right] \right) = \frac{h(t)}{\alpha(t)} A m(A^\top x_t), \quad (11)$$

here we denote for short $m(u) := \nabla \log p_t^{\text{LD}}(u) + \frac{1}{h(t)} u$. From (11), we see that $\nabla_{x_t} \mathbb{E}[x_0|x_t]^\top$ maps any vector to $\text{Span}(A)$ because $m(\cdot)$ takes $A^\top x_t$ as input in the expression of $\mathbb{E}[x_0|x_t]$. ■

We highlight that the faithfulness of \mathbf{G}_{loss} holds for *arbitrary* data distribution supported on the latent subspace. It takes advantage of the score function's decomposition (5), having the effect of automatically adapting g onto the latent low-dimensional subspace of data.

4.5 Estimation and Implementation of G_{loss}

Theorem 1 asserts that the gradient guidance given by Definition 1 provably preserves the subspace structure of data. However, G_{loss} is not immediately available to compute and it involves the unknown quantity $\mathbb{E}[x_0|x_t]$. Next, we discuss the estimation and computation of G_{loss} based on a pre-trained score function s_θ in practice.

First, we need to estimate the quantity $\mathbb{E}[x_0|x_t]$. It is the conditional expectation of x_0 given x_t in the forward process, thus it depends on the pre-training data distribution. One can construct estimate $\hat{\mathbb{E}}[x_0|x_t]$ based on the pre-trained score network s_θ , by using the Tweedie’s formula (Efron, 2011):

$$\nabla \log p_t(x_t) = -\mathbb{E} \left[\frac{x_t - \alpha(t)x_0}{h(t)} | x_t \right]. \quad (12)$$

Suppose we have a given pre-trained score network that approximates the ground truth, i.e., $s_\theta(x_t, t) \approx \nabla \log p_t(x_t)$. Then a natural estimator of $\hat{\mathbb{E}}[x_0|x_t]$ is given by

$$\hat{\mathbb{E}}[x_0|x_t] := \frac{1}{\alpha(t)} (x_t + h(t)s_\theta(x_t, t)), \quad (13)$$

and we refer to it as the *look-ahead estimator*. The estimator (13) is widely adopted in practice (Song et al., 2020a; Bansal et al., 2023). Here $\alpha(t)$ and $h(t)$ are the noise scheduling used in the forward process (1).

Thus, we have obtained an **implementable version of the gradient guidance** G_{loss} , given by

$$G_{loss}(x_t, t) = -\beta(t) \cdot \nabla_{x_t} \left(y - g^\top \left(\frac{1}{\alpha(t)} (x_t + h(t)s_\theta(x_t, t)) \right) \right)^2, \quad (14)$$

With a slight abuse of notation, we use G_{loss} to refer to this implementable formula (14) in the remainder of this paper.

The gradient guidance (14) has a light-weighted implementation. Suppose the pre-trained score function s_θ is given in the form of a neural network with pre-trained weights. Computing (14) involves calculating the squared loss $\left(y - g^\top \hat{\mathbb{E}}[x_0|x_t] \right)^2$ via a forward pass of the network s_θ and a backward pass utilizing the auto-gradient feature of deep-learning frameworks such as PyTorch and TensorFlow. See Figure 4 for illustration.

Note that the value of y in G_{loss} is a target reward value, inherited from the conditional score analysis under a Gaussian model. In practice, we treat y as a tuning parameter. In our theoretical analysis, we will specify the choices of $y, \beta(t)$ and provide guarantees for general optimization beyond the Gaussian model.

So far, we have finally obtained a gradient guidance (14) that is both implementable and faithful to data’s latent subspace. The next step is to apply this gradient guidance and use it to adapt the generation process of a pre-trained diffusion model. Let us find out what one can obtain using gradient-guided diffusion models.

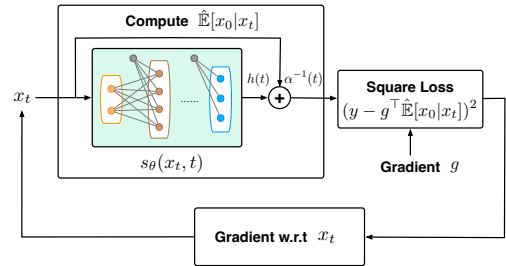


Figure 4: Computation of Gradient Guidance G_{loss} .

5 Gradient-Guided Diffusion Model as Regularized Optimizer

In this section, we study whether gradient guidance steers a pre-trained diffusion model to generate samples of **near-optimal** objective values. We provide a positive answer and our results are twofold: 1) We demonstrate that iteratively applying gradient guidance improves the generated samples towards higher objective values; 2) The pre-trained diffusion model acts as a form of regularization from an optimization perspective.

5.1 Gradient-Guided Generation with A Pre-trained Score

Assume access to a pre-trained score network s_θ and gradient information of the objective function f . Let us present our Algorithm 1 that adapts the pre-trained diffusion model and iteratively updates the gradient guidance (14). The gradient guidance is able to steer the backward sampling process towards generating new samples with higher values of f . See Figure 1 for illustration.

Algorithm 1 takes as input any pre-trained score function $s_\theta(x, t)$ and adapts the backward sampling process with gradient guidance. In each iteration, it evaluates $\nabla f(\cdot)$ at samples generated from the previous iteration (Line 5(i)), and then computes the gradient guidance \mathbf{G}_{loss} using the newly queried gradient (Line 5(ii)). Using the updated gradient guidance, the backward process then generate new samples with improved objective values (Module 1). At the end of iterations, the algorithm outputs an adapted version of the diffusion model, specified by (s_θ, \mathbf{G}_K) , which generates samples with near-optimal objective values.

Algorithm 1 Gradient-Guided Diffusion for Generative Optimization

- 1: **Input:** Pre-trained score network $s_\theta(\cdot, \cdot)$, differentiable objective function f .
 - 2: **Tuning Parameter:** Strength parameters $\beta(t)$, $\{y_k\}_{k=0}^{K-1}$, number of iterations K , batch sizes $\{B_k\}$.
 - 3: **Initialization:** $\mathbf{G}_0 = \text{NULL}$.
 - 4: **for** $k = 0, \dots, K - 1$ **do**
 - 5: **Generate:** Sample $z_{k,i} \sim \text{Guided_BackwardSample}(s_\theta, \mathbf{G}_k)$ using Module 1, for $i \in [B_k]$.
 - 6: **Compute Guidance:**
 - (i) Compute the sample mean $\bar{z}_k := (1/B_k) \sum_{i=1}^{B_k} z_{k,i}$.
 - (ii) Query gradient $g_k = \nabla f(\bar{z}_k)$.
 - (iii) Update gradient guidance $\mathbf{G}_{k+1}(\cdot, \cdot) = \mathbf{G}_{loss}(\cdot, \cdot)$ via (8), using s_θ , gradient vector g_k , and parameters y_k and $\beta(t)$.
 - 7: **end for**
 - 8: **Output:** (s_θ, \mathbf{G}_K) .
-

It is worth highlighting that Algorithm 1 works with *any* pre-trained score network $s_\theta(x_t, t)$. It retains the original score network and only updates the guidance term. The gradient guidance changes the generation process by an additive term to the backward SDE, without having to re-train the score network. Therefore, the algorithm is computationally efficient and easy to implement. We experimented with Algorithm 1 using a pre-trained score network with about 15M parameters; see Section 7 for details. In our experiment, a single run of the backward sampling process (Module 1) takes 4.6s, and Algorithm 1 takes 76min overall. Thus it is a rather light-weighted algorithm to implement and run.

5.2 Gradient-Guided Diffusion Converges to Regularized Optima

We analyze the convergence properties of Algorithm 1 and show that in final iterations, generated samples center around a regularized solution of the optimization objective f . Our theorems allow the pre-training data to have *arbitrary distribution*.

Assumption 3 (Concave smooth objective). *The objective $f : \mathbb{R}^D \rightarrow \mathbb{R}$ is concave and L -smooth with respect to the (semi-)norm $\|\cdot\|_{\bar{\Sigma}^{-1}}$, i.e., $\|\nabla f(x_1) - \nabla f(x_2)\|_{\bar{\Sigma}} \leq L \|x_1 - x_2\|_{\bar{\Sigma}^{-1}}$ for any x_1, x_2 .*

While Algorithm 1 works with any pre-trained score network, we study its optimization properties focusing on the class of linear score functions given by

$$\mathcal{S} = \{s(x, t) = C_t x + b_t : C_t \in \mathbb{R}^{D \times D}, b_t \in \mathbb{R}^D\}. \quad (15)$$

Here (15) is a general linear function class. For comparison, a recent related paper Marion et al. (2024) assumes a more restricted class with $C_t \equiv I$ for when studying parameter optimization in diffusion models. With a linear score function (15), pre-training a diffusion model is essentially the same as using a Gaussian model to estimate the pre-training data distribution and then sampling from this estimated Gaussian. In this case, the guidance \mathbf{G}_{loss} is also linear in x_t , therefore the final output of the guided diffusion model also follows a Gaussian distribution; see (27) in Appendix C.

Recall we aim for an adapted diffusion model (s_θ, \mathbf{G}_K) to generate samples with high values of f . Thus, we focus on the mean of the generated distribution (taking $T \rightarrow \infty$ in the backward sampling process of (s_θ, \mathbf{G}_K)), denoted by μ_K , and establish its optimization guarantee.

Theorem 2 (Convergence to Regularized Maxima in Mean). *Let Assumption 3 hold, and let the pre-training data \mathcal{D} have arbitrary distribution with covariance matrix $\bar{\Sigma} \succ 0$. Suppose the score function s_θ is pre-trained via minimizing the score matching loss (3) over the linear function class (15). Let Alg. 1 take $s_\theta(\cdot, \cdot)$ and f as the input. For any $\lambda > L$, there exists $\{\beta(t)\}$, $\{y_k\}$, $\{B_k\}$ such that, with probability $\geq 1 - \delta$, the mean of the output distribution μ_K converges to be near x_λ^* , and*

$$f(x_\lambda^*) - f(\mu_K) = \lambda \left(\frac{L}{\lambda}\right)^K \mathcal{O}\left(D \log\left(\frac{K}{\delta}\right)\right), \quad (16)$$

where D is the ambient dimension of data, and x_λ^* is a regularized maximizer of f given by

$$x_\lambda^* = \operatorname{argmax}_{x \in \mathbb{R}^D} \left\{ f(x) - \frac{\lambda}{2} \|x - \bar{\mu}\|_{\bar{\Sigma}^{-1}}^2 \right\}, \quad (17)$$

where $\bar{\mu}, \bar{\Sigma}$ are empirical mean and covariance of pre-training data \mathcal{D} .

Proof Sketch Solving the score matching problem (3) with a linear function class (15) yields a pre-trained score as follows

$$s_\theta(x_t, t) = -(\alpha^2(t)\bar{\Sigma} + h(t)I_D)^{-1}(x_t - \alpha(t)\bar{\mu}).$$

With proper choices of $\beta(t)$, gradient guidance \mathbf{G}_{loss} leads to the following output distribution at the end of round k :

$$\mathcal{N}\left(\bar{\mu} + \frac{y_k - g_k^\top \bar{\mu}}{\sigma^2 + g_k^\top \bar{\Sigma} g_k} \bar{\Sigma} g_k, \bar{\Sigma} - \frac{\bar{\Sigma} g_k g_k^\top \bar{\Sigma}}{\sigma^2 + g_k^\top \bar{\Sigma} g_k}\right).$$

Thus, we obtain the mean of the above distribution, i.e., $\mu_{k+1} = \bar{\mu} + \eta_k \bar{\Sigma} \nabla f(\bar{z}_k)$, where \bar{z}_k is the empirical mean of previous samples, η_k is a stepsize determined by y_k . By a rearrangement, we obtain a recursive formula

$$\mu_{k+1} = \bar{z}_k + \eta_k \bar{\Sigma} [\nabla f(\bar{z}_k) - \eta_k^{-1} \bar{\Sigma}^{-1} (\bar{z}_k - \bar{\mu})]. \quad (18)$$

We observe that (18) resembles a gradient ascent update from $\mu_k \approx \bar{z}_k$ to μ_{k+1} corresponding to a regularized optimization problem (17). In this regularized objective, the original objective $f(x)$ incorporates an additional proximal term with $\lambda := 1/\eta_k$. Therefore we can analyze the convergence of μ_k by following the classical argument for gradient optimization. The full proof is provided in Appendix C.1

Remarks. This view of regularized optimization gives the following insights on gradient-guided diffusion models:

- (i) The regularization term $\frac{\lambda}{2} \|x - \bar{\mu}\|_{\bar{\Sigma}^{-1}}^2$ in (17) is centered at the data’s mean $\bar{\mu}$. It penalizes samples that are far away from the pre-training data. The norm $\|\cdot\|_{\bar{\Sigma}^{-1}}$ suggests the regularization is strong in the direction where the original data distribution has low variance. In other words, it reveals that the pre-trained score function acts as a form of “prior” in the guided generation process. This prior favors samples that are proximal to its pre-training data distribution, even when additional guidance are present.
- (ii) The regularization term cannot be made arbitrarily small. In particular, our theorem requires that $\lambda \geq L$. This demonstrates a limit of adapting diffusion models with guidance. As long as the score function stays unchanged, one cannot extrapolate from the pre-training data unlimitedly by solely adding gradient guidance. As a consequence, we *cannot* simply add gradient guidance to a diffusion model in order to reach the global maxima of any objective function. If the goal is to reach global optima, one has to update the pre-trained score network and refine it with newly collected data, we explore this approach in Section 6.
- (iii) The linear convergence rate (16) is determined jointly by the smoothness of the objective function and strength of the regularization. We also pay a linear factor in the dimension D . In the following subsection, we will show that the gradient guidance \mathbf{G}_{loss} can reduce the dimension dependence from D to d if the data admits a latent low-dimensional subspace.

5.3 Gradient Guidance for Optimization in Latent Spaces

Next we focus on data with latent subspace as in Assumption 1. In the next theorem, we show that the generated distribution of our adapted model would converge, in expectation, to the maxima of a regularized version of f within the subspace $\text{Span}(A)$.

Theorem 3 (Convergence to Regularized Maxima in Latent Subspace in Mean). *Let Assumptions 1 and 3 hold. Suppose we use the score function class (15) for pre-training and computing guidance. Then Alg.1 gives an adapted diffusion model that generates new samples that belong to $\text{Span}(A)$. Further, for any $\lambda > L$, there exists $\beta(t), \{y_k\}$ and batch size B_k , such that with high probability $1 - \delta$, the mean of the output distribution μ_K converges to be near $x_{A,\lambda}^*$, and it holds*

$$f(x_{A,\lambda}^*) - f(\mu_K) = \lambda \left(\frac{L}{\lambda}\right)^K \mathcal{O}\left(d \log\left(\frac{K}{\delta}\right)\right),$$

where $x_{A,\lambda}^*$ is an optimal solution of the regularized objective:

$$x_{A,\lambda}^* = \operatorname{argmax}_{x \in \operatorname{Span}(A)} \left\{ f(x) - \frac{\lambda}{2} \|x - \bar{\mu}\|_{\Sigma^{-1}}^2 \right\}. \quad (19)$$

Recall that the gradient guidance \mathbf{G}_{loss} is faithful to the data’s latent subspace, as proved in Theorem 1. As a result, the gradient-guided backward process maintains this latent subspace structure in the generated output. Therefore, all the generated samples and optimization iterates of Algorithm 1 belong to the latent subspace $\operatorname{Span}(A)$. In other words, the entire optimization process happens in the latent low-dimensional subspace. This facilitates a coherent and more efficient exploration in the solution space. Comparing to Theorem 2, the optimization gap in Theorem 3 is substantially smaller, reduced from $\mathcal{O}(D)$ to $\mathcal{O}(d)$. It means that the optimization process leverages the latent subspace and converges much faster.

Finally we note that Theorem 3 only establishes convergence in mean. The final output distribution of Algorithm 1, with a linear score function, is a Gaussian distribution supported on the latent subspace; see Appendix C equation (27).

6 Gradient-Guided Diffusion with Adaptive Fine-Tuning for Global Optimization

In the previous section, we have seen that adding guidance to a pre-trained diffusion model cannot improve the objective function unlimitedly. The pre-trained score function would act as a form of prior to keep the generated output proximal to the original data’s distribution. This leads to a regularization term in the optimization formulation.

Further, we consider adaptively fine-tuning the pre-trained diffusion model for generating samples to attain the unregularized global optima. The idea is not only to update the guidance in the backward sampling process but also to use generated samples to fine-tune the pre-trained score network. Empirically, fine-tuning diffusion models utilizing self-generated samples has been explored by Black et al. (2023); Clark et al. (2023).

6.1 Adaptive Fine-Tuning Algorithm with Gradient Guidance

We propose an adaptive version of the gradient-guided diffusion, where both the gradient guidance and the score networks are iteratively updated utilizing self-generated samples. The full algorithm is given in Algorithm 2.

We introduce a weighting scheme to fine-tune the score network using a mixture of pre-training data and newly generated samples. In Round k , let $\mathcal{D}_1, \dots, \mathcal{D}_k$ be sample batches generated from the previous rounds. Let $\{w_{k,i}\}_{i=0}^k$ be a set of weights. Conceptually, at Round k , we update the model by minimizing the weighted score matching loss:

$$\min_{s \in \mathcal{S}} \int_0^T \sum_{i=0}^k w_{k,i} \mathbb{E}_{x_0 \in \mathcal{D}_i} \mathbb{E}_{x_t | x_0} \left[\|\nabla_{x_t} \log \phi_t(x_t | x_0) - s(x_t, t)\|_2^2 \right] dt, \quad (20)$$

where $\mathcal{D}_0 := \mathcal{D}$ is the pre-training data. For illustration of this algorithm, please see also Figure 1.

In practice, to update the score network incorporating newly generated data, one does not have to exactly solve (20) by re-training the full model from scratch. Instead, (20) can be viewed as a

guideline that motivates more computationally efficient ways for updating the pre-trained score. It is a common practice to only **fine-tune** the weights of the old model by performing gradient descent over a few batches of newly generated data, which is similar to the spirit of (20).

In our experiment, we implemented Algorithm 2 using a pre-trained U-net score function with 15M parameters and tested its performance on synthetic objectives. We implement the finetuning step by making one single Adam step over the new data. In our experiment, the iterated finetuning process of Algorithm 2 takes 91min overall, only slightly longer than the 76min taken by Algorithm 1. For details on the experiment results, please see Section 7.

Algorithm 2 Gradient-Guided Diffusion with **Adaptive Fine-tuning**

- 1: **Input:** Pre-trained score $s_\theta(\cdot, \cdot)$, differentiable objective function f .
 - 2: **Tuning Parameter:** strength parameter $\beta(t)$, $\{y_k\}_{k=0}^{K-1}$, **weights** $\{\{w_{k,i}\}_{i=0}^k\}_{k=0}^{K-1}$, number of iterations K , batch sizes $\{B_k\}$.
 - 3: **Initialize:** $s_{\theta_0} = s_\theta$, $G_0 = \text{NULL}$.
 - 4: **for** $k = 0, \dots, K-1$ **do**
 - 5: **Generate:** Sample a batch $\mathcal{D}_k = \{z_{k,i}\}_{i=1}^{B_k}$ from **Guided_BackwardSample**(s_{θ_k}, G_k) (Module 1).
 - 6: **Compute Guidance:**
 - (i) Compute sample mean $\bar{z}_k = (1/B_k) \sum_{i=1}^{B_k} z_{k,i}$, and query gradient $g_k = \nabla f(\bar{z}_k)$.
 - (ii) **Update** s_{θ_k} **to** $s_{\theta_{k+1}}$ **by minimizing the re-weighted objective** (20).
 - (iii) Compute $G_{k+1}(\cdot, \cdot) = G_{\text{loss}}(\cdot, \cdot)$ in (8), using $s_{\theta_{k+1}}$ and g_k , with parameter $y_k, \beta(t)$.
 - 7: **end for**
 - 8: **Output:** (s_{θ_K}, G_K) .
-

6.2 Guided Generation Finds Unregularized Global Optima

Finally, we analyze the optimization properties for gradient-guided diffusion model with iterative finetuning. We establish that the process of Algorithm 2 yields a final output distribution whose mean, denoted by μ_K , converges to the global optimum of f .

For simplicity of analysis, we study the following function class

$$\mathcal{S}' = \left\{ s(x, t) = \hat{C}_t x + b_t : b_t \in \mathbb{R}^D \right\}, \quad (21)$$

where \hat{C}_t is set to stay the same as in the pre-trained score and only b_t gets updated during iterative fine-tuning. Marion et al. (2024) studied a similar function class where \hat{C}_t is freezed to be $\hat{C}_t \equiv I$.

Theorem 4 (Convergence to Unregularized Maxima in Latent Subspace in Mean). *Let Assumptions 1 and 3 hold, and assume there exists $M > 0$ such that $\|x_{A,\lambda}^*\| < M$ for all $\lambda \geq 0$. Suppose we use the score function class (15) for pre-training s_θ and the class (21) for finetuning it. Then Algorithm 2 gives an adapted diffusion model that generates new samples belonging to $\text{Span}(A)$. Further, there exists $\{\beta(t)\}, \{y_k\}, \{B_k\}$ and $\{w_{k,i}\}$, such that with probability $1 - \delta$,*

$$f_A^* - f(\mu_K) = \mathcal{O} \left(\frac{dL^2 \log K}{K} \cdot \log \left(\frac{K}{\delta} \right) \right), \quad (22)$$

where $f_A^* = \max\{f(x) | x \in \text{Span}(A)\}$.

The proof idea is similar to the proof of Theorem 2. For simplicity, we analyze the case where only the most recent sample batch \mathcal{D}_k is merged with \mathcal{D}_0 for finetuning the score function. More

specifically, we let $w_{k,i} = 0$ for $0 < i < k$ and $w_{k,0} = 1 - w_{k,k}$. Similar to the proof of Theorem 2, we obtain a recursive update rule given by

$$\mu_{k+1} = \bar{z}_k + \eta_k \bar{\Sigma} \left[\nabla f(\bar{z}_k) - (1 - w_{k,k}) \eta_k^{-1} \cdot \bar{\Sigma}^{-1} (\bar{z}_k - \bar{\mu}) \right], \quad (23)$$

where $\bar{z}_k \approx \mu_k$ is the empirical mean of previous samples. This update rule also closely resembles the gradient ascent iteration for maximizing a regularized objective. A key difference here is that we can control the weights $w_{k,i}$ to reduce the impact of \mathcal{D}_0 and make the regularization term vanish to zero. Thus the mean μ_k eventually converges to the global maxima. For the detailed arguments and proof of convergence, we refer readers to Appendix C.3.

Theorem 4 illustrates the effect of finetuning a diffusion model using self-generated data. For comparison, Theorem 3 showed that without finetuning the diffusion model can only generate new samples proximal to the pre-training distribution. Now if we allow finetuning using self-generated samples, the diffusion model can iteratively refine itself and reach global optima, while preserving the latent subspace structure in its generated output.

Now let us take on an optimization view. The convergence rate suggested by Theorem 4 matches with that of standard convex optimization, in terms of their dependence on K the number of gradient evaluations. Further, if we compare the guided diffusion model with a standard gradient solver, the optimality gap of our algorithm scales with the small intrinsic dimension $\mathcal{O}(d)$, while standard gradient ascent converges much more slowly due to the large ambient dimension $\mathcal{O}(D)$. This comparison highlights the merits of “generative optimization”. More specifically, diffusion models leverage pre-training data to learn their intrinsic characteristics. Therefore, when we add gradient guidance to the pre-trained score function and use it for generation, it means that we are solving an optimization problem in its own intrinsic low-dimensional space. This leads to substantially more efficient exploration and faster convergence. This theoretical insight explains the practical successes of guided diffusion models on complex optimization problems, such as video creation, image synthesis and protein AI, where traditional methods do not work at all.

7 Numerical Experiments

We experiment with our design of the gradient guidance as well as Algorithm 1 and Algorithm 2. Going beyond our theoretical assumptions, we adopt a 15M-parameter U-Net as the score function class for training and fine-tuning our diffusion model.

7.1 Experiment Setup

We set the data’s ambient dimension as $D = 64$ and the linear subspace dimension as $d = 16$. The linear subspace is represented by an orthogonal matrix $A \in \mathbb{R}^{D \times d}$. We randomly generate a matrix A and fix it once generated. After that, we sample a data point X by first randomly sampling a latent variable $U \sim \mathcal{N}(0, I_d)$ and computing $X = AU$. We independently sample a total of 65536 data points as our pre-training data set. The objective functions considered in our experiments are $f_1(x) = 10 - (\theta^\top x - 3)^2$ and $f_2(x) = 5 - 0.5\|x - b\|$. Here, θ and b are randomly generated and fixed afterward. Since our data assumes a low-dimensional subspace representation, it is convenient to decompose θ into $\theta_\perp = (I - AA^\top)\theta$ and $\theta_\parallel = AA^\top\theta$, representing the off-support and on-support components. We refer to $\frac{\|\theta_\perp\|}{\|\theta_\parallel\|}$ as the off/on-support ratio. Analogously, for a generated sample, we can also define its off/on-support ratio. Clearly, a small off/on-support ratio indicates close vicinity to the subspace.

Score Network Pre-training We utilize a version of the U-Net (Ronneberger et al., 2015), with 14.8M trainable parameters. Note that this is a complicated network going beyond the linear score function class considered in our theories. Following the implementation of Denoising Diffusion Probabilistic Models (DDPM, Ho et al. (2020)), we train the U-Net to estimate the score function $\nabla \log p_t$, via minimizing the score matching loss introduced in Eqn. (3). We discretize the backward process to have 200 time steps as in Nichol and Dhariwal (2021), and the U-Net is trained using our generated data set for 20 epochs. We use Adam as the optimizer, set the batch size as 32, and set the learning rate to be 10^{-4} . After the pre-training phase, we confirmed that the data subspace structure is well learned, as the generated samples using the pre-trained diffusion model have an average off/on-support ratio of 0.039.

Implementation of Algorithm 1 In each iteration of Algorithm 1, we need to compute the gradient guidance \mathbf{G}_{loss} . We set the targeted y value at the k -th iteration as $y_k = \delta + g_k^\top z_k$, where δ_k specifies the increment per iteration. The choice on δ_k is instance-dependent and we set it via tuning for near-optimal in different experiments. For comparing naive gradient with gradient guidance in Figure 5, we set $\delta = 0.2$ and 0.9 , respectively for using naive gradient \mathbf{G} and gradient guidance \mathbf{G}_{loss} . In Figure 6, we choose δ to be (a) 0.05, (b) 0.2, (c) 1, and (d) 1, corresponding to each panel. We initialize Algorithm 1 with a batch of 32 samples generated by the pre-trained model. Each sample determines an optimization trajectory. We repeat Algorithm 1 for 5 times with different random seeds and report the error bars.

Implementation of Algorithm 2 Algorithm 2 differs from Algorithm 1 in that it allows additional fine-tuning of the pre-trained score network. We adopt a computationally lightweight fine-tuning strategy: We only perform one Adam optimization step using the re-weighted loss given by Eqn. (20) with a batch of 32 generated samples. We set the learning rate as 10^{-6} . This simple strategy already demonstrates good performances as shown in Figure 7. Other implementation details are kept the same as those of Algorithm 1.

We run all experiments using one NVIDIA A100 GPU. Module 1 takes 4.6 seconds to generate a sample. Algorithm 1 takes 76 minutes, and Algorithm 2 takes 91 minutes.

7.2 Results

We first demonstrate our gradient guidance \mathbf{G}_{loss} preserves the subspace structure learned from the pre-trained model. For comparison, we also tested the naive guidance \mathbf{G} defined following Lemma 1 (with $\Sigma = I$). For a quick reference, we repeat the definition here:

$$\mathbf{G}(x_t, t) := \beta(t) \left(y - g^\top \mathbb{E}[x_0|x_t] \right) g,$$

where $\beta(t) > 0$ and $y \in \mathbb{R}$ are tuning parameters, and $\mathbb{E}[x_0|x_t]$ is the conditional expectation of x_0 given noise corrupted data x_t . For implementation, we replace $\mathbb{E}[x_0|x_t]$ by its look-ahead estimator $\hat{\mathbb{E}}[x_0|x_t]$ based on the Tweedie’s formula.

Comparing \mathbf{G} and \mathbf{G}_{loss} on Preserving Subspace Structure Figure 5(a), (c) verify that the naive gradient \mathbf{G} performs much worse than \mathbf{G}_{loss} in preserving the linear subspace structure. It is consistent with our theoretical finding that the gradient guidance \mathbf{G}_{loss} keeps the generated sample close to the latent subspace, with substantially smaller off-support errors. When allowing adaptive

score fine-tuning in Algorithm 2, Figure 5(b), (d) show that the off-support error increases as the model gets fine-tuned using self-generated data, due to increasing distribution shift. Even in this case, the naive gradient G leads to much more severe off-support errors as compared to G_{loss} .

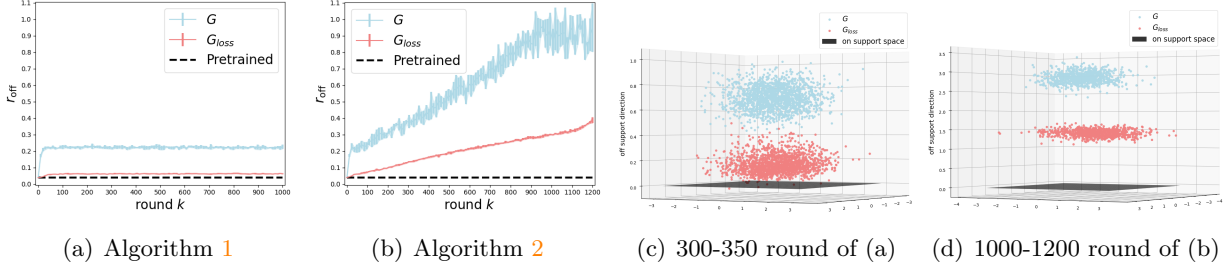


Figure 5: **Comparison between two types of gradient guidance G and G_{loss} .** We plot the off/on support ratio of the generated samples, denoted by $r_{off} = \frac{\|x_{\perp}\|}{\|x_{\parallel}\|}$. The objective function is $f_1(x)$, with θ having an off/on-support ratio of 9.

Algorithm 1 Converges to Regularized Optima We plot the convergence of Algorithm 1 in terms of the objective value in Figure 6. Figure 6 (a),(b) are for the objective function $f_1 = 10 - (\theta^\top x - 3)^2$ as the objective function, while Figure 6(c),(d) are for the objective $f_2 = 5 - 0.5\|x - b\|$. We observe that the algorithm converges to reach some sub-optimal objective value, but there remains a gap to the maximal value. This is consistent with our theory that the pre-trained model essentially acts as a regularization in addition to the objective function. Adding gradient guidance alone cannot reach global maxima. This coincides with our theoretical findings in Theorem 3.

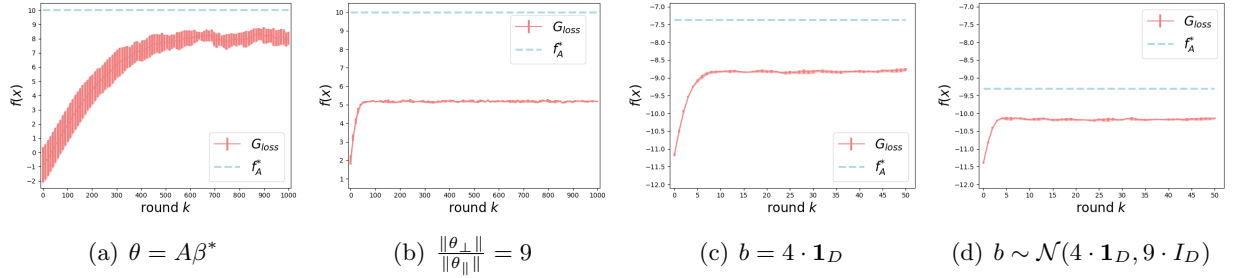


Figure 6: **Convergence of Algorithm 1 under different objectives.** Objectives are $f_1(x)$ for (a) and (b), and $f_2(x)$ for (c) and (d). Parameters θ and b are specified as (a) $\theta = A\beta^*$ with β^* being sampled from the unit ball in \mathbb{R}^d ; (b) the off/on-support ratio of θ being 9 (same as Figure 5); (c) and (d) choosing b as a homogeneous vector or randomly from a Gaussian distribution. All the experiments adopt the gradient guidance G_{loss} .

Algorithm 2 Converges to Global Optima Algorithm 2 converges to the maximal value of the objective function $f_1 = 10 - (\theta^\top x - 3)^2$ as shown in Figure 7(a). In Figure 7(b), we visualize the distribution of generated samples of Algorithm 1 (blue) and 2 (red), respectively, as the iteration evolves. We see that samples from Algorithm 1 mostly stay close to the pre-training data distribution (area described by the dotted contour). In contrast, samples of Algorithm 2 move outside the contour, as the diffusion model gets finetuned using self-generated data.

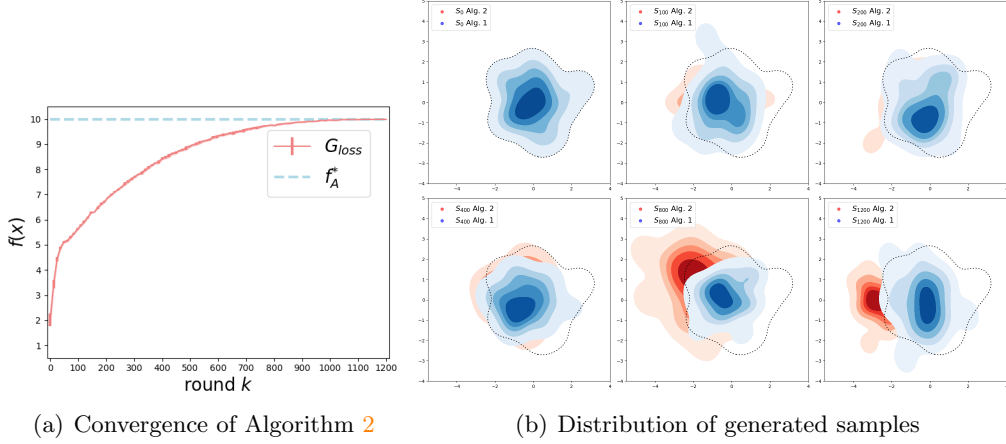


Figure 7: **Convergence of Algorithm 2.** Panel (a) plots the objective values achieved by Algorithm 2 as a function of iterations. Here θ is chosen the same as in Figure 6(b) with off/on-support ratio $\frac{\|\theta_{\perp}\|}{\|\theta_{\parallel}\|} = 9$. Panel (b) visualizes the distribution of the generated samples of Algorithm 2 (red) across the iterations. For comparison, we also visualize the distribution of generated samples of Algorithm 1 (blue).

8 Conclusion

In this paper, we investigate the role and design of gradient guidance for adapting and fine-tuning a pre-trained diffusion model from an optimization perspective. We propose a gradient guidance based on a lookahead loss, as well as two variants of diffusion-based generative optimization algorithm utilizing such guidance. We provide optimization guarantees for adapting/fine-tuning diffusion models towards maximizing any target concave differentiable reward function. Our analysis has also been extended to linear subspace data, where our gradient guidance and adaptive algorithms provably preserve and leverage the latent subspace, thus they achieve faster convergence to near-optimal solutions.

References

- Anurag Ajay, Yilun Du, Abhi Gupta, Joshua Tenenbaum, Tommi Jaakkola, and Pulkit Agrawal. Is conditional generative modeling all you need for decision-making? *arXiv preprint arXiv:2211.15657*, 2022.
- Michael S Albergo, Nicholas M Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unifying framework for flows and diffusions. *arXiv preprint arXiv:2303.08797*, 2023.
- Brian DO Anderson. Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12(3):313–326, 1982.
- 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 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 843–852, 2023.
- Joe Benton, Valentin De Bortoli, Arnaud Doucet, and George Deligiannidis. Linear convergence bounds for diffusion models via stochastic localization. *arXiv preprint arXiv:2308.03686*, 2023.
- Kevin Black, Michael Janner, Yilun Du, Ilya Kostrikov, and Sergey Levine. Training diffusion models with reinforcement learning. *arXiv preprint arXiv:2305.13301*, 2023.
- Adam Block, Youssef Mroueh, and Alexander Rakhlin. Generative modeling with denoising auto-encoders and langevin sampling. *arXiv preprint arXiv:2002.00107*, 2020.
- Sébastien Bubeck et al. Convex optimization: Algorithms and complexity. *Foundations and Trends® in Machine Learning*, 8(3-4):231–357, 2015.
- Minshuo Chen, Kaixuan Huang, Tuo Zhao, and Mengdi Wang. Score approximation, estimation and distribution recovery of diffusion models on low-dimensional data. *arXiv preprint arXiv:2302.07194*, 2023a.
- Minshuo Chen, Song Mei, Jianqing Fan, and Mengdi Wang. An overview of diffusion models: Applications, guided generation, statistical rates and optimization. *arXiv preprint arXiv:2404.07771*, 2024.
- Sitan Chen, Sinho Chewi, Jerry Li, Yuanzhi Li, Adil Salim, and Anru R Zhang. Sampling is as easy as learning the score: theory for diffusion models with minimal data assumptions. *arXiv preprint arXiv:2209.11215*, 2022.
- Sitan Chen, Sinho Chewi, Holden Lee, Yuanzhi Li, Jianfeng Lu, and Adil Salim. The probability flow ode is provably fast. *arXiv preprint arXiv:2305.11798*, 2023b.
- Sitan Chen, Giannis Daras, and Alex Dimakis. Restoration-degradation beyond linear diffusions: A non-asymptotic analysis for ddim-type samplers. In *International Conference on Machine Learning*, pages 4462–4484. PMLR, 2023c.
- Hyungjin Chung, Jeongsol Kim, Michael T Mccann, Marc L Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. *arXiv preprint arXiv:2209.14687*, 2022a.

- Hyungjin Chung, Byeongsu Sim, Dohoon Ryu, and Jong Chul Ye. Improving diffusion models for inverse problems using manifold constraints. *Advances in Neural Information Processing Systems*, 35:25683–25696, 2022b.
- Kevin Clark, Paul Vicol, Kevin Swersky, and David J Fleet. Directly fine-tuning diffusion models on differentiable rewards. *arXiv preprint arXiv:2309.17400*, 2023.
- Valentin De Bortoli. Convergence of denoising diffusion models under the manifold hypothesis. *arXiv preprint arXiv:2208.05314*, 2022.
- Valentin De Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. Diffusion schrödinger bridge with applications to score-based generative modeling. *Advances in Neural Information Processing Systems*, 34:17695–17709, 2021.
- Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021.
- Bradley Efron. Tweedie’s formula and selection bias. *Journal of the American Statistical Association*, 106(496):1602–1614, 2011.
- Ahmed El Alaoui, Andrea Montanari, and Mark Sellke. Sampling from mean-field gibbs measures via diffusion processes. *arXiv preprint arXiv:2310.08912*, 2023.
- Ying Fan, Olivia Watkins, Yuqing Du, Hao Liu, Moonkyung Ryu, Craig Boutilier, Pieter Abbeel, Mohammad Ghavamzadeh, Kangwook Lee, and Kimin Lee. Dpok: Reinforcement learning for fine-tuning text-to-image diffusion models. *arXiv preprint arXiv:2305.16381*, 2023.
- Hengyu Fu, Zhuoran Yang, Mengdi Wang, and Minshuo Chen. Unveil conditional diffusion models with classifier-free guidance: A sharp statistical theory. *arXiv preprint arXiv:2403.11968*, 2024.
- Alexandros Graikos, Nikolay Malkin, Nebojsa Jojic, and Dimitris Samaras. Diffusion models as plug-and-play priors. *Advances in Neural Information Processing Systems*, 35:14715–14728, 2022.
- Nate Gruver, Samuel Stanton, Nathan C Frey, Tim GJ Rudner, Isidro Hotzel, Julien Lafrance-Vanasse, Arvind Rajpal, Kyunghyun Cho, and Andrew Gordon Wilson. Protein design with guided discrete diffusion. *arXiv preprint arXiv:2305.20009*, 2023.
- Zhiye Guo, Jian Liu, Yanli Wang, Mengrui Chen, Duolin Wang, Dong Xu, and Jianlin Cheng. Diffusion models in bioinformatics: A new wave of deep learning revolution in action. *arXiv preprint arXiv:2302.10907*, 2023.
- Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising diffusion restoration models. *Advances in Neural Information Processing Systems*, 35:23593–23606, 2022.
- Zhifeng Kong, Wei Ping, Jiaji Huang, Kexin Zhao, and Bryan Catanzaro. Diffwave: A versatile diffusion model for audio synthesis. *arXiv preprint arXiv:2009.09761*, 2020.

- Holden Lee, Jianfeng Lu, and Yixin Tan. Convergence for score-based generative modeling with polynomial complexity. *arXiv preprint arXiv:2206.06227*, 2022a.
- Holden Lee, Jianfeng Lu, and Yixin Tan. Convergence of score-based generative modeling for general data distributions. *arXiv preprint arXiv:2209.12381*, 2022b.
- Xingchao Liu, Lemeng Wu, Mao Ye, and Qiang Liu. Let us build bridges: Understanding and extending diffusion generative models. *arXiv preprint arXiv:2208.14699*, 2022.
- Yixin Liu, Kai Zhang, Yuan Li, Zhiling Yan, Chujie Gao, Ruoxi Chen, Zhengqing Yuan, Yue Huang, Hanchi Sun, Jianfeng Gao, et al. Sora: A review on background, technology, limitations, and opportunities of large vision models. *arXiv preprint arXiv:2402.17177*, 2024.
- Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11461–11471, 2022.
- Pierre Marion, Anna Korba, Peter Bartlett, Mathieu Blondel, Valentin De Bortoli, Arnaud Doucet, Felipe Llinares-López, Courtney Paquette, and Quentin Berthet. Implicit diffusion: Efficient optimization through stochastic sampling. *arXiv preprint arXiv:2402.05468*, 2024.
- Song Mei and Yuchen Wu. Deep networks as denoising algorithms: Sample-efficient learning of diffusion models in high-dimensional graphical models. *arXiv preprint arXiv:2309.11420*, 2023.
- Andrea Montanari and Yuchen Wu. Posterior sampling from the spiked models via diffusion processes. *arXiv preprint arXiv:2304.11449*, 2023.
- Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *Proceedings of the International Conference on Machine Learning*, pages 8162–8171. PMLR, 2021.
- Kazusato Oko, Shunta Akiyama, and Taiji Suzuki. Diffusion models are minimax optimal distribution estimators. *arXiv preprint arXiv:2303.01861*, 2023.
- Phillip Pope, Chen Zhu, Ahmed Abdelkader, Micah Goldblum, and Tom Goldstein. The intrinsic dimension of images and its impact on learning. *arXiv preprint arXiv:2104.08894*, 2021.
- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015.
- Sam T Roweis and Lawrence K Saul. Nonlinear dimensionality reduction by locally linear embedding. *science*, 290(5500):2323–2326, 2000.
- Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020a.
- Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in Neural Information Processing Systems*, 32, 2019.

- Yang Song, Sahaj Garg, Jiaxin Shi, and Stefano Ermon. Sliced score matching: A scalable approach to density and score estimation. In *Uncertainty in Artificial Intelligence*, pages 574–584. PMLR, 2020b.
- Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020c.
- Joshua B Tenenbaum, Vin de Silva, and John C Langford. A global geometric framework for nonlinear dimensionality reduction. *science*, 290(5500):2319–2323, 2000.
- Masatoshi Uehara, Yulai Zhao, Kevin Black, Ehsan Hajiramezanali, Gabriele Scalia, Nathaniel Lee Diamant, Alex M Tseng, Tommaso Biancalani, and Sergey Levine. Fine-tuning of continuous-time diffusion models as entropy-regularized control. *arXiv preprint arXiv:2402.15194*, 2024.
- Yinhuai Wang, Jiwen Yu, and Jian Zhang. Zero-shot image restoration using denoising diffusion null-space model. *arXiv preprint arXiv:2212.00490*, 2022.
- Joseph L Watson, David Juergens, Nathaniel R Bennett, Brian L Trippe, Jason Yim, Helen E Eisenach, Woody Ahern, Andrew J Borst, Robert J Ragotte, Lukas F Milles, et al. De novo design of protein structure and function with rfdiffusion. *Nature*, 620(7976):1089–1100, 2023.
- Andre Wibisono, Yihong Wu, and Kaylee Yingxi Yang. Optimal score estimation via empirical bayes smoothing. *arXiv preprint arXiv:2402.07747*, 2024.
- Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Wentao Zhang, Bin Cui, and Ming-Hsuan Yang. Diffusion models: A comprehensive survey of methods and applications. *ACM Computing Surveys*, 56(4):1–39, 2023.
- Hui Yuan, Kaixuan Huang, Chengzhuo Ni, Minshuo Chen, and Mengdi Wang. Reward-directed conditional diffusion: Provable distribution estimation and reward improvement. *arXiv preprint arXiv:2307.07055*, 2023.

A Characterization for Output Distribution of Backward Process

In this section, we provide analytical characterizations for the output distribution of the backward process guided by \mathbf{G}_{loss} when the pre-trained score is linear. We first give the result of score matching as follows.

Lemma 2 (Pre-training with Linear Score Functions). *Suppose for pre-training the score network, the class in (20) is*

$$\mathcal{S} = \{s(x, t) = C_t x + b_t : C_t \in \mathbb{R}^{D \times D}, b_t \in \mathbb{R}^D\}, \quad (\text{recall (15)})$$

If we freeze C_t in (15), that is, minimizing the score matching objective (20) over the class $\{s(x, t) = C_t x + b_t : b_t \in \mathbb{R}^D\}$ gives

$$s_\theta(x_t, t) = C_t (x_t - \alpha(t)\bar{x}),$$

where $\bar{x} = \left(\sum_{i=0}^k w_{k,i} \mathbb{E}_{x \in \mathcal{D}_i}[x]\right) / \left(\sum_{i=0}^k w_{k,i}\right)$. Moreover, minimizing the score matching objective (20) over the class (15) yields

$$s_\theta(x_t, t) = -(\alpha^2(t)\bar{\Sigma} + h(t)I_D)^{-1} (x_t - \alpha(t)\bar{x}),$$

where $\bar{\Sigma} = \left(\sum_{i=0}^k w_{k,i} \mathbb{E}_{x \in \mathcal{D}_i}[(x - \bar{x})(x - \bar{x})^\top]\right) / \left(\sum_{i=0}^k w_{k,i}\right)$ are weighted data covariance.

Proof. Using the linear score network class \mathcal{S} with freezing C_t , we cast the score matching loss (20) into

$$\begin{aligned} & \int_0^T \sum_{i=0}^k w_{k,i} \mathbb{E}_{x_0 \in \mathcal{D}_i} \mathbb{E}_{x_t \sim \mathcal{N}(\alpha(t)x_0, h(t)I_D)} \left[\|\nabla_{x_t} \log \phi_t(x_t|x_0) - s(x_t, t)\|^2 \right] dt \\ &= \int_0^T \sum_{i=0}^k w_{k,i} \mathbb{E}_{x_0 \in \mathcal{D}_i} \mathbb{E}_{x_t \sim \mathcal{N}(\alpha(t)x_0, h(t)I_D)} \left[\left\| -\frac{1}{h(t)}(x_t - \alpha(t)x_0) - C_t x_t - b_t \right\|^2 \right] dt \\ &= \int_0^T \sum_{i=0}^k w_{k,i} \mathbb{E}_{x_0 \in \mathcal{D}_i} \mathbb{E}_{x_t \sim \mathcal{N}(\alpha(t)x_0, h(t)I_D)} \left[\left\| \left(C_t + \frac{1}{h(t)} I_D \right) (x_t - \alpha(t)x_0) + (\alpha(t)C_t x_0 + b_t) \right\|^2 \right] dt \\ &\stackrel{(i)}{=} \int_0^T \sum_{i=0}^k w_{k,i} \mathbb{E}_{x_0 \in \mathcal{D}_i} \left[\|b_t + \alpha(t)C_t x_0\|^2 \right] + w \int_0^T \text{trace} \left(h(t) \left(C_t + \frac{1}{h(t)} I_D \right)^\top \left(C_t + \frac{1}{h(t)} I_D \right) \right) dt, \end{aligned}$$

where $w = \sum_{i=0}^k w_{k,i}$, equality (i) follows from computing the expectation over the conditional Gaussian distribution of $x_t|x_0$. We note that b_t should minimize $\sum_{i=0}^k w_{k,i} \mathbb{E}_{x_0 \in \mathcal{D}_i} [\|b_t + \alpha(t)C_t x_0\|^2]$ for any t , which leads to

$$\hat{b}_t = -\alpha(t)C_t \bar{x}.$$

Now, we solve C_t for the second result. Substituting \hat{b}_t into the optimization objective (20) yields:

$$\begin{aligned} & \int_0^T \sum_{i=0}^k w_{k,i} \mathbb{E}_{x_0 \in \mathcal{D}_i} \left[\|\alpha(t)C_t \bar{x} - \alpha(t)C_t x_0\|^2 \right] + w \int_0^T \text{trace} \left(h(t) \left(C_t + \frac{1}{h(t)} I_D \right)^\top \left(C_t + \frac{1}{h(t)} I_D \right) \right) dt \\ &= \int_0^T \alpha^2(t) \sum_{i=0}^k w_{k,i} \mathbb{E}_{x_0 \in \mathcal{D}_i} \left[\|C_t(x_0 - \bar{x})\|^2 \right] + w \int_0^T \text{trace} \left(h(t) \left(C_t^\top C_t + \frac{1}{h(t)} C_t + \frac{1}{h(t)} C_t^\top + \frac{1}{h^2(t)} I_D \right) \right) dt. \end{aligned}$$

Taking the gradient for C_t , we get

$$2\alpha^2(t)C_t \sum_{i=0}^k w_{k,i} \mathbb{E}_{x_0 \in \mathcal{D}_i} \left[(x_0 - \bar{x})(x_0 - \bar{x})^\top \right] + 2wh(t)C_t + 2wI_D.$$

Setting the gradient above to 0, we get the solution for C_t as

$$\hat{C}_t = - \left(\alpha^2(t)w^{-1} \sum_{i=0}^k w_{k,i} \mathbb{E}_{x_0 \in \mathcal{D}_i} \left[(x_0 - \bar{x})(x_0 - \bar{x})^\top \right] + h(t)I_D \right)^{-1}.$$

Therefore, the proof is completed. \square

When $w_{k,0} = 1, w_{k,i} = 0, i \in [k]$, Lemma 2 reduces to the pre-training score matching.

Corollary 1. *Let \mathcal{D} be the pre-training data. Minimizing the score matching objective (3) over the function class (15) gives*

$$s_\theta(x_t, t) = -(\alpha^2(t)\bar{\Sigma} + h(t)I_D)^{-1}(x_t - \alpha(t)\bar{\mu}). \quad (24)$$

The following lemma characterizes the output distribution of the backward process guided by \mathbf{G}_{loss} when the pre-trained score is linear.

Lemma 3. *If the pre-trained score $s_\theta(x_t, t)$ is (24), substituting the score function with $s_\theta(x_t, t) + \mathbf{G}_{loss}(x_t, t)$ in the backward SDE (2) yields, when $T \rightarrow \infty$,*

$$X_T^\leftarrow \stackrel{d}{=} \mathcal{N} \left(\bar{\mu} + \frac{y - g^\top \bar{\mu}}{\sigma^2 + g^\top \bar{\Sigma} g} \bar{\Sigma} g, \bar{\Sigma} - \frac{\bar{\Sigma} g g^\top \bar{\Sigma}}{\sigma^2 + g^\top \bar{\Sigma} g} \right).$$

with $\beta(t)$ assigned as $\beta(t) = \frac{1}{2} \left(\sigma^2 + g^\top \bar{\Sigma}^{-1} (I_D + \alpha^2(t)\bar{\Sigma}/h(t))^{-1} g \right)^{-1}$. Moreover, if pre-training data reside in $\text{Span}(A)$ following Assumption 1, it holds $X_T^\leftarrow \in \text{Span}(A)$.

Proof. Consider $X_0 \stackrel{d}{=} \mathcal{N}(\bar{\mu}, \bar{\Sigma})$, $Y = g^\top X_0 + \epsilon$ where $\epsilon \sim \mathcal{N}(0, \sigma^2)$. Let X_0 be the initialization of the forward process. Similar to the proof in Appendix B.1, we get

$$\begin{bmatrix} X_t \\ Y \end{bmatrix} \stackrel{d}{=} \mathcal{N} \left(\begin{bmatrix} \alpha(t)\bar{\mu} \\ g^\top \bar{\mu} \end{bmatrix}, \begin{bmatrix} \alpha^2(t)\bar{\Sigma} + h(t)I_D & \alpha(t)\bar{\Sigma}g \\ \alpha(t)g^\top \bar{\Sigma} & \sigma^2 + g^\top \bar{\Sigma}g \end{bmatrix} \right).$$

Thus, we get $s_\theta(x_t, t)$ is exactly the score of marginal distribution of X_t , i.e., $\nabla \log p_t(x_t) = s_\theta(x_t, t)$. According to the proof in Appendix B.1, we get $s_\theta(x_t, t) + \mathbf{G}_{loss}(x_t, t) = \nabla \log p_t(x_t | y)$. Thus, the backward SDE turns out to be

$$dX_t^\leftarrow = \left[\frac{1}{2} X_t^\leftarrow + \nabla \log p_{T-t}(X_t^\leftarrow | y) \right] dt + d\bar{W}_t, \quad X_0^\leftarrow \stackrel{d}{=} \mathcal{N}(0, I_D). \quad (25)$$

Since the initial distribution $p_0(x_0 | y)$ of the forward process can also be obtained by (25) where we replace the initial distribution as $p_T(x_T | y)$. According to Girsanov theorem, we get the KL divergence between the terminal distribution p_T^\leftarrow of (25) and $p_0(x_0 | y)$.

$$\text{KL}(p_0 | p_T^\leftarrow) = -\mathbb{E}_{p_0} \left[\log \left(\frac{p_T}{\varphi} \right) \right],$$

where p_0, p_T are short hands for $p_0(x_0 | y)$ and $p_T(x_T | y)$, and $\varphi(\cdot)$ is the density for the standard normal distribution $\mathcal{N}(0, I_D)$. Since in the forward process, $p_T \rightarrow \varphi$ when $T \rightarrow \infty$, we have $\text{KL}(p_0 | p_T^\leftarrow) \rightarrow 0$ when $T \rightarrow \infty$. We complete the first part of the lemma. As for the second part, if data reside in $\text{Span}(A)$ following Assumption 1, we have $\bar{\mu} = A\bar{u}$ and $\bar{\Sigma} = A\bar{\Sigma}_u A^\top$, where $\bar{u} = \mathbb{E}_{x \in \mathcal{D}, u=A^\top x}[u]$, $\bar{\Sigma}_u = \mathbb{E}_{x \in \mathcal{D}, u=A^\top x}[(u - \bar{u})(u - \bar{u})^\top]$. Thus, the covariance matrix of X_T^\leftarrow is

$$\bar{\Sigma} - \frac{\bar{\Sigma} g g^\top \bar{\Sigma}}{\sigma^2 + g^\top \bar{\Sigma} g} = A \left[\bar{\Sigma}_u - \frac{\bar{\Sigma}_u A^\top g g^\top A \bar{\Sigma}_u}{\sigma^2 + g^\top \bar{\Sigma} g} \right] A^\top,$$

and due to X_T^\leftarrow follows Gaussian distribution, we get $X_T^\leftarrow \in \text{Span}(A)$. Thus, the proof is completed. \square

B Omitted Proofs in Section 4

Contents

B.1 Proof of Lemma 1	28
B.2 Proof of Theorem 1	29

Throughout the proof, we omit the subscript x_t of ∇ in score function $\nabla \log p_t(x_t)$ and conditional score function $\nabla \log p_t(x_t|y)$.

B.1 Proof of Lemma 1

Proof. Recall $\{X_t\}_{t \geq 0}$ is the stochastic process from the forward process. $X_0 \stackrel{d}{=} \mathcal{N}(\mu, \Sigma)$, $Y = g^\top X_0 + \epsilon$ where $\epsilon \sim \mathcal{N}(0, \sigma^2)$ is independent with X_0 . Since X_0, X_t and Y are joint Gaussian, we have $Y | X_t$ also follows Gaussian distribution, denoted as $\mathcal{N}(m_y(x_t), \sigma_y^2(x_t))$. Then, the closed form of $\nabla_{x_t} \log p_t(y | x_t)$ can be derived as

$$\nabla_{x_t} \log p_t(y | x_t) = -\nabla_{x_t} \left[\frac{1}{2} \left(\frac{y - m_y(x_t)}{\sigma_y(x_t)} \right)^2 \right] - \nabla_{x_t} \log \sigma_y(x_t).$$

Due to the linearity of Y with regard to X_0 , $m_y(x_t)$ can be computed as

$$m_y(x_t) = \mathbb{E}[y | x_t] = \mathbb{E}[g^\top x_0 + \epsilon | x_t] = \mathbb{E}[g^\top x_0 | x_t] = g^\top \mathbb{E}[x_0 | x_t]. \quad (26)$$

To get the variance $\sigma_y^2(x_t)$, we compute the joint distribution (X_t, Y) . In the forward process, given $X_0 = x_0$, X_t can be written as $\alpha(t)x_0 + Z_t$ for $Z_t \stackrel{d}{=} \mathcal{N}(0, h(t)I_D)$ independent of x_0 . Due to the linear function assumption, we have

$$\begin{bmatrix} X_t \\ Y \end{bmatrix} = \begin{bmatrix} \alpha(t)I_D & 0 & I_D \\ g^\top & 1 & 0 \end{bmatrix} \begin{bmatrix} x_0 \\ \epsilon \\ Z_t \end{bmatrix}.$$

Observing that the joint distribution of (x_0, ϵ, z_t) is Gaussian, we deduce

$$\begin{bmatrix} X_t \\ Y \end{bmatrix} \stackrel{d}{=} \mathcal{N} \left(\begin{bmatrix} \alpha(t)\mu \\ g^\top \mu \end{bmatrix}, \begin{bmatrix} \alpha^2(t)\Sigma + h(t)I_D & \alpha(t)\Sigma g \\ \alpha(t)g^\top \Sigma & \sigma^2 + g^\top \Sigma g \end{bmatrix} \right).$$

Thus, we get $\sigma_y^2(x_t) = \sigma^2 + g^\top \Sigma g - \alpha^2(t) g^\top \Sigma (\alpha^2(t) \Sigma + h(t) I_D)^{-1} \Sigma g$. Together with the derivation of the mean $m_y(x_t)$ (26), we get

$$\begin{aligned} \nabla_{x_t} \log p_t(y | x_t) &= -\frac{1}{2\sigma_y^2(x_t)} \nabla_{x_t} \left[\left(y - g^\top \mathbb{E}[x_0 | x_t] \right)^2 \right] \\ &= \frac{1}{\sigma_y^2(x_t)} \left(y - g^\top \mathbb{E}[x_0 | x_t] \right) \nabla_{x_t} \mathbb{E}[x_0 | x_t] g. \end{aligned}$$

To get $\mathbb{E}[x_0 | x_t]$, we derive the joint distribution of (X_0, X_t) :

$$\begin{bmatrix} X_0 \\ X_t \end{bmatrix} \stackrel{d}{=} \mathcal{N} \left(\begin{bmatrix} \mu \\ \alpha(t)\mu \end{bmatrix}, \begin{bmatrix} \Sigma & \alpha(t)\Sigma \\ \alpha(t)\Sigma & \alpha^2(t)\Sigma + h(t)I_D \end{bmatrix} \right).$$

Thus, we get $\mathbb{E}[x_0 | x_t] = \mu + \alpha(t) (\alpha^2(t)\Sigma + h(t)I_D)^{-1} \Sigma (x_t - \alpha(t)\mu)$. As a consequence, we have

$$\nabla_{x_t} \log p_t(y | x_t) = \frac{1}{\sigma_y^2(x_t)} \left(y - g^\top \mathbb{E}[x_0 | x_t] \right) \alpha(t) (\alpha^2(t)\Sigma + h(t)I_D)^{-1} \Sigma g.$$

Together with the following equality by Woodbury identity, we get the result.

$$\begin{aligned} \sigma_y^{-2}(x_t) &= \left(\sigma^2 + g^\top \Sigma g - \alpha^2(t) g^\top \Sigma (\alpha^2(t)\Sigma + h(t)I_D)^{-1} \Sigma g \right)^{-1} \\ &= \left[\sigma^2 + g^\top \Sigma^{-1} \left(I_D + \frac{\alpha^2(t)}{h(t)} \Sigma \right)^{-1} g \right]^{-1}. \end{aligned}$$

□

B.2 Proof of Theorem 1

Proof. We expand the derivative in \mathbf{G}_{loss} as

$$\mathbf{G}_{loss}(x_t, t) = 2\beta(t)(y - g^\top \mathbb{E}[x_0 | x_t]) (\nabla_{x_t} \mathbb{E}[x_0 | x_t])^\top g.$$

It holds $\mathbb{E}[x_0 | x_t] = \frac{1}{\alpha(t)}(x_t + h(t)\nabla \log p_t(x_t))$. Via the score decomposition under linear subspace data in Chen et al. (2023a, Lemma 1), we have

$$\nabla \log p_t(x_t) = A \nabla \log p_t^{\text{LD}}(A^\top x_t) - \frac{1}{h(t)} \left(I_D - AA^\top \right) x_t,$$

where p_t^{LD} denotes the diffused latent distribution, i.e., $p_t^{\text{LD}}(u') = \int \phi_t(u' | u_0) p_u(u_0) du_0$. Recall that ϕ_t is the Gaussian transition kernel of the forward process and p_u is the density of latent variable u_0 in Assumption 1.

To ease the derivation, we denote $m(u) = \nabla \log p_t^{\text{LD}}(u) + \frac{1}{h(t)} u$. It then holds that

$$\begin{aligned} \mathbb{E}[x_0 | x_t] &= \frac{1}{\alpha(t)} \left(x_t + h(t) \left[A m(A^\top x_t) - \frac{1}{h(t)} x_t \right] \right) \\ &= \frac{h(t)}{\alpha(t)} A m(A^\top x_t). \end{aligned}$$

As a consequence, we can verify that

$$\nabla_{x_t} \mathbb{E}[x_0|x_t] = \frac{h(t)}{\alpha(t)} A \left[\nabla m(A^\top x_t) \right] A^\top,$$

where $\nabla m(A^\top x_t) \in \mathbb{R}^{d \times d}$ is the Jacobian matrix of m at $A^\top x_t$. Plugging the last display into \mathbf{G}_{loss} , we conclude that

$$\begin{aligned} \mathbf{G}_{loss}(x_t, t) &= 2\beta(t)(y - g^\top \mathbb{E}[x_0|x_t]) (\nabla_{x_t} \mathbb{E}[x_0|x_t])^\top g \\ &= \ell_t \cdot g \end{aligned}$$

for $\ell_t = 2\beta(t)(y - g^\top \mathbb{E}[x_0|x_t])$ and $g' = (\nabla_{x_t} \mathbb{E}[x_0|x_t])^\top g = \frac{h(t)}{\alpha(t)} A \left[\nabla m(A^\top x_t) \right] A^\top g \in \text{Span}(A)$. The proof is complete. \square

C Omitted Proofs in Sections 5 and 6

Contents

C.1 Proof of Theorem 2	30
C.2 Proof of Theorem 3	30
C.3 Proof of Theorem 4	33
C.4 Auxiliary Lemma	34

C.1 Proof of Theorem 2

Proof. The proof is a special case of Theorem 3 in Appendix C.2, via setting the representation matrix $A = I_D$. \square

C.2 Proof of Theorem 3

Proof of Theorem 3. Define a filtration $\{\mathcal{H}_k\}_{k=0}^{K-1}$ with \mathcal{H}_k be the information accumulated after k rounds of Alg.1.

$$\begin{aligned} \mathcal{H}_0 &:= \sigma(\bar{\mu}), \\ \mathcal{H}_k &:= \sigma \left(\mathcal{H}_{k-1}, \sigma \left(z_{k-1,1}, \dots, z_{k-1,B_{k-1}} \right) \right), \quad k \in [K]. \end{aligned}$$

Define the expectation of samples generated at k -th round as

$$\mu_k := \mathbb{E}[z_{k,i} \mid \mathcal{H}_{k-1}], \quad k \in [K-1].$$

Applying Corollary 1, we get the pre-trained score as

$$s_\theta(x_t, t) = - \left(\alpha^2(t) \bar{\Sigma} + h(t) I_D \right)^{-1} (x_t - \alpha(t) \bar{\mu}),$$

If we set y_k as follows

$$y_k = \eta \cdot \left(\sigma^2 + g_k^\top \bar{\Sigma} g_k \right) + g_k^\top \bar{\mu},$$

where $\eta = 1/\lambda$. And we choose $\beta(t)$ at k -round as $\beta(t) = \frac{1}{2} \left(\sigma^2 + g_{k-1}^\top \bar{\Sigma}^{-1} (I_D + \alpha^2(t) \bar{\Sigma}/h(t))^{-1} g_{k-1} \right)^{-1}$. Then, Lemma 3 provides the generated distribution in k -th round:

$$\mathcal{N} \left(\bar{\mu} + \eta \bar{\Sigma} g_{k-1}, \bar{\Sigma} - \frac{\bar{\Sigma} g_{k-1} g_{k-1}^\top \bar{\Sigma}}{\sigma^2 + g_{k-1}^\top \bar{\Sigma} g_{k-1}} \right). \quad (27)$$

Define the empirical covariance matrix of the latent variable U as $\bar{\Sigma}_u = \mathbb{E}_{x \in \mathcal{D}, u=A^\top x}[(u - \bar{u})(u - \bar{u})^\top]$ where $\bar{u} = \mathbb{E}_{\mathcal{D}}[u]$. Then in the subspace setting, the empirical mean and covariance of data X can be written as $\bar{\mu} = AA^\top \bar{\mu}$ and $\bar{\Sigma} = A\bar{\Sigma}_u A^\top$ respectively. The mean of the sample $z_{k,i}$ follows

$$\mu_k = \mathbb{E}[z_{k,i} \mid \mathcal{H}_{k-1}] = AA^\top \bar{\mu} + \eta \cdot A\bar{\Sigma}_u A^\top g_{k-1},$$

where $g_{k-1} = \nabla f(\bar{z}_{k-1})$ and $\bar{z}_{k-1} = (1/B) \sum_i^B z_{k-1,i}$. We rearrange the update rule to show a gradient ascent formula as follows

$$\begin{aligned} \mu_k &= AA^\top \mu_{k-1} - AA^\top (\mu_{k-1} - \bar{\mu}) + \eta \cdot A\bar{\Sigma}_u A^\top \nabla f(\mu_{k-1}) + \eta \cdot A\bar{\Sigma}_u A^\top (g_{k-1} - \nabla f(\mu_{k-1})) \\ &= AA^\top \mu_{k-1} - A\bar{\Sigma}_u A^\top A\bar{\Sigma}_u^{-1} A^\top (\mu_{k-1} - \bar{\mu}) + \eta \cdot A\bar{\Sigma}_u A^\top \nabla f(\mu_{k-1}) + \eta \cdot A\bar{\Sigma}_u A^\top (g_{k-1} - \nabla f(\mu_{k-1})) \\ &= AA^\top \mu_{k-1} + \eta \cdot A\bar{\Sigma}_u A^\top \left[\nabla f(\mu_{k-1}) - \lambda A\bar{\Sigma}_u^{-1} A^\top (\mu_{k-1} - \bar{\mu}) \right] + \eta \cdot A\bar{\Sigma}_u A^\top (g_{k-1} - \nabla f(\mu_{k-1})). \end{aligned}$$

where $\lambda = 1/\eta$. Define $h(x) := f(x) - \lambda/2 \|x - \bar{\mu}\|_{\bar{\Sigma}^{-1}}^2$, we have

$$\mu_k = AA^\top \mu_{k-1} + \eta \cdot \bar{\Sigma} \nabla h(\mu_{k-1}) + \eta \cdot \bar{\Sigma} (g_{k-1} - \nabla f(\mu_{k-1})).$$

Recall the notation for the optimum: $x_{A,\lambda}^* = \operatorname{argmax}_{x=Au} h(x)$. We consider the distance of μ_k to $x_{A,\lambda}^*$ under the semi-norm $\|\cdot\|_{\bar{\Sigma}^{-1}}$.

$$\begin{aligned} \|\mu_k - x_{A,\lambda}^*\|_{\bar{\Sigma}^{-1}} &= \|\mu_{k-1} - x_{A,\lambda}^* + \eta \bar{\Sigma} \nabla h(\mu_{k-1}) + \eta \bar{\Sigma} (g_{k-1} - \nabla f(\mu_{k-1}))\|_{\bar{\Sigma}^{-1}} \\ &\leq \underbrace{\|\mu_{k-1} - x_{A,\lambda}^* + \eta \bar{\Sigma} \nabla h(\mu_{k-1})\|_{\bar{\Sigma}^{-1}}}_{:=I_1} + \underbrace{\|\eta \bar{\Sigma} (g_{k-1} - \nabla f(\mu_{k-1}))\|_{\bar{\Sigma}^{-1}}}_{:=I_2}. \end{aligned} \quad (28)$$

We bound the second term I_2 first. According to f is L -smooth with respect to $\|\cdot\|_{\bar{\Sigma}^{-1}}$, we have

$$I_2 = \eta \|g_{k-1} - \nabla f(\mu_{k-1})\|_{\bar{\Sigma}} \leq \eta L \|\bar{z}_{k-1} - \mu_{k-1}\|_{\bar{\Sigma}^{-1}},$$

Lemma 3 shows the distribution of $z_{k-1,i}$. Therefore, according to concentration inequality for Gaussian distribution, with the probability at least $1 - \delta/K$, it holds

$$\|\bar{z}_{k-1} - \mu_{k-1}\|_{\bar{\Sigma}^{-1}}^2 \leq 2 \log \left(\frac{2K}{\delta} \right) \cdot \frac{\operatorname{trace}(\mathbb{V}(z_{k-1,i}) \cdot \bar{\Sigma}^{-1})}{B_{k-1}}.$$

We have $\operatorname{trace}(\mathbb{V}(z_{k-1,i}) \cdot \bar{\Sigma}^{-1}) \leq \operatorname{trace}(\bar{\Sigma} \cdot \bar{\Sigma}^{-1}) = d$. Therefore, I_2 is bounded by

$$I_2 \leq M_0 / \sqrt{B_{k-1}},$$

where $M_0 := \eta L \sqrt{2 \log \left(\frac{2K}{\delta} \right) \cdot d}$. Next, we consider the first term in (28). Since $x_{A,\lambda}^*$ is the optimum of h within $\operatorname{Span}(A)$, the gradient $\nabla h(x_{A,\lambda}^*)$ is in the orthogonal subspace, i.e., $A^\top \nabla h(x_{A,\lambda}^*) = 0$, thus $\bar{\Sigma} \nabla h(x_{A,\lambda}^*) = 0$. The first term in (28) can be written as

$$\begin{aligned} I_1^2 &= \|(\mu_{k-1} - x_{A,\lambda}^*) + \eta \bar{\Sigma} (\nabla h(\mu_{k-1}) - \nabla h(x_{A,\lambda}^*))\|_{\bar{\Sigma}^{-1}}^2 \\ &= \|\mu_{k-1} - x_{A,\lambda}^*\|_{\bar{\Sigma}^{-1}}^2 + \eta^2 \|\nabla h(\mu_{k-1}) - \nabla h(x_{A,\lambda}^*)\|_{\bar{\Sigma}}^2 \\ &\quad + 2 \langle \mu_{k-1} - x_{A,\lambda}^*, \eta (\nabla h(\mu_{k-1}) - \nabla h(x_{A,\lambda}^*)) \rangle. \end{aligned}$$

Recall h is f adding a $\|\cdot\|_{\bar{\Sigma}^{-1}}$ regularized term. We get h is $(L + \lambda)$ -smooth with respect to semi norm $\|\cdot\|_{\bar{\Sigma}^{-1}}$ which is derived from f L -smooth. Also, h is λ -strongly concave with respect to semi norm $\|\cdot\|_{\bar{\Sigma}^{-1}}$ since f is concave. According to Lemma 4, we derive

$$\langle \mu_{k-1} - x_{A,\lambda}^*, \nabla h(\mu_{k-1}) - \nabla h(x_{A,\lambda}^*) \rangle \leq -\frac{\lambda(L + \lambda)}{L + 2\lambda} \|\mu_{k-1} - x_{A,\lambda}^*\|_{\bar{\Sigma}^{-1}}^2 - \frac{1}{L + 2\lambda} \|\nabla h(\mu_{k-1}) - \nabla h(x_{A,\lambda}^*)\|_{\bar{\Sigma}}^2.$$

Plugin the formula of I_1 , we get

$$I_1^2 \leq \left(1 - \frac{2\eta\lambda(L + \lambda)}{L + 2\lambda}\right) \|\mu_{k-1} - x_{A,\lambda}^*\|_{\bar{\Sigma}^{-1}}^2 + \left(\eta^2 - \frac{2\eta}{L + 2\lambda}\right) \|\nabla h(\mu_{k-1}) - \nabla h(x_{A,\lambda}^*)\|_{\bar{\Sigma}}^2.$$

Since $\eta = 1/\lambda$, it holds $\eta^2 - \frac{2\eta}{L + 2\lambda} > 0$. Due to h $(L + \lambda)$ -smoothness, we get

$$\begin{aligned} I_1^2 &\leq \left(1 - \frac{2\eta\lambda(L + \lambda)}{L + 2\lambda}\right) \|\mu_{k-1} - x_{A,\lambda}^*\|_{\bar{\Sigma}^{-1}}^2 + \left(\eta^2 - \frac{2\eta}{L + 2\lambda}\right) (L + \lambda)^2 \|\mu_{k-1} - x_{A,\lambda}^*\|_{\bar{\Sigma}^{-1}}^2 \\ &= (1 - \eta(L + \lambda))^2 \|\mu_{k-1} - x_{A,\lambda}^*\|_{\bar{\Sigma}^{-1}}^2, \end{aligned}$$

thus, we get the bound of I_1

$$I_1 \leq \zeta \|\mu_{k-1} - x_{A,\lambda}^*\|_{\bar{\Sigma}^{-1}},$$

where $\zeta := |1 - \eta(L + \lambda)|$. Combing the upper bound of I_1 and I_2 , we get with probability at least $1 - \delta/K$, for $1 < k \leq K$,

$$\|\mu_k - x_{A,\lambda}^*\|_{\bar{\Sigma}^{-1}} \leq \zeta \|\mu_{k-1} - x_{A,\lambda}^*\|_{\bar{\Sigma}^{-1}} + \frac{M_0}{\sqrt{B_{k-1}}}.$$

As for $k = 1$, by similar derivation, we can obtain $\|\mu_1 - x_{A,\lambda}^*\|_{\bar{\Sigma}^{-1}} \leq \zeta \|z_0 - x_{A,\lambda}^*\|_{\bar{\Sigma}^{-1}}$. By induction, we get with probability at least $1 - ((K - 1)/K)\delta$,

$$\|\mu_K - x_{A,\lambda}^*\|_{\bar{\Sigma}^{-1}} \leq \zeta^K \|z_0 - x_{A,\lambda}^*\|_{\bar{\Sigma}^{-1}} + M_0 \sum_{k=1}^{K-1} \frac{\zeta^{K-k-1}}{\sqrt{B_k}}.$$

Choose $B_k \geq \zeta^{-4k}(1 - \zeta)^{-2}$ for all $k \in [K - 1]$, then we can get

$$\|\mu_K - x_{A,\lambda}^*\|_{\bar{\Sigma}^{-1}} \leq \zeta^K \left(\|z_0 - x_{A,\lambda}^*\|_{\bar{\Sigma}^{-1}}^2 + M_1 \cdot \sqrt{d} \right) \quad (29)$$

where $M_1 := \eta L \sqrt{2 \log \left(\frac{2K}{\delta} \right)}$. Since h is $(L + \lambda)$ -smooth with respect to $\|\cdot\|_{\bar{\Sigma}^{-1}}$, it holds

$$|h(\mu_K) - h(x_{A,\lambda}^*) - \langle \nabla h(x_{A,\lambda}^*), \mu_K - x_{A,\lambda}^* \rangle| \leq \frac{L + \lambda}{2} \|\mu_K - x_{A,\lambda}^*\|_{\bar{\Sigma}^{-1}}^2.$$

Considering that $\nabla h(x_{A,\lambda}^*) \perp \text{Span}(A)$ yields $\langle \nabla h(x_{A,\lambda}^*), \mu_K - x_{A,\lambda}^* \rangle = 0$, we obtain the following by rearranging the equation above

$$\begin{aligned} f(x_{A,\lambda}^*) - f(\mu_K) &\leq \frac{\lambda}{2} \left(\|x_{A,\lambda}^* - \bar{\mu}\|_{\bar{\Sigma}^{-1}}^2 - \|\mu_K - \bar{\mu}\|_{\bar{\Sigma}^{-1}}^2 \right) + \frac{L + \lambda}{2} \|\mu_K - x_{A,\lambda}^*\|_{\bar{\Sigma}^{-1}}^2 \\ &\leq \left[\lambda \|\bar{\mu} - x_{A,\lambda}^*\|_{\bar{\Sigma}^{-1}} \|\mu_K - x_{A,\lambda}^*\|_{\bar{\Sigma}^{-1}} + (L + \lambda) \|\mu_K - x_{A,\lambda}^*\|_{\bar{\Sigma}^{-1}}^2 \right]. \end{aligned} \quad (30)$$

Substitute (29) into above upper bound, with $z_0 = \bar{\mu}$ we have

$$f(x_{A,\lambda}^*) - f(\mu_K) \lesssim \zeta^K \cdot (L + \lambda) \left[\|\bar{\mu} - x_{A,\lambda}^*\|_{\bar{\Sigma}^{-1}}^2 + M_1^2 d \right].$$

Since $\|\bar{\mu} - x_{A,\lambda}^*\|_{\bar{\Sigma}^{-1}}^2 = \|A^\top (\bar{\mu} - x_{A,\lambda}^*)\|_{\bar{\Sigma}_u^{-1}}^2$ is the distance within $\text{Span}(A)$, i.e., $\mathcal{O}(d)$. Recall $\eta = 1/\lambda$, $\zeta = |1 - \eta(L + \lambda)| = L/\lambda$, $\lambda > L$, and $M_1 = \eta L \sqrt{2 \log \left(\frac{2K}{\delta} \right)}$. Therefore, we get the final result:

$$f(x_{A,\lambda}^*) - f(\mu_K) \lesssim \lambda \left(\frac{L}{\lambda} \right)^K d \log \left(\frac{K}{\delta} \right), \quad \text{w.p.} 1 - \delta.$$

□

C.3 Proof of Theorem 4

Proof. Define

$$\begin{aligned} \mathcal{H}_0 &:= \sigma(\bar{\mu}), \\ \mathcal{H}_k &:= \sigma \left(\mathcal{H}_{k-1}, \sigma(z_{k-1,1}, \dots, z_{k-1,B_{k-1}}) \right), \quad k \in [K-1], \\ \mu_k &:= \mathbb{E}[z_{k,i} \mid \mathcal{H}_{k-1}], \quad k \in [K]. \end{aligned}$$

According to Lemma 2, with freezing C_t in class (15), the pre-trained score in Round k is $s_{\theta_{k+1}}(x_t, t) = -(\alpha^2(t)\bar{\Sigma} + h(t)I_D)(x_t - \alpha(t)\bar{x}_k)$ where $\bar{x}_k = \sum_{j=0}^k w_{k,j} \bar{z}_j$ and $\bar{z}_j = \mathbb{E}_{x \in \mathcal{D}_j}[x]$. By choosing y_k , and weights $w_{k,j}$ as

$$\begin{aligned} y_k &= \eta_k \cdot \left(\sigma^2 + g_k^\top A \bar{\Sigma}_u A^\top g_k \right) + g_k^\top A A^\top \bar{x}_k, \\ w_{k,0} &= 1 - w_k \\ w_{k,j} &= 0, \quad 1 \leq j < k, \\ w_{k,k} &= w_k, \end{aligned}$$

where $\eta_k > 0$, $0 < w_k < 1$ will be specified later. And we choose $\beta(t)$ at Round k as $\beta(t) = \frac{1}{2} \left(\sigma^2 + g_{k-1}^\top \bar{\Sigma}^{-1} (I_D + \alpha^2(t)\bar{\Sigma}/h(t))^{-1} g_{k-1} \right)^{-1}$. Lemma 3 gives the mean of distribution of $z_{k+1,i}$ as

$$\mu_{k+1} = \bar{x}_k + \eta_k \bar{\Sigma} g_k, \quad (31)$$

and the output distribution

$$\mathcal{N} \left(\bar{x}_{K-1} + \eta_{K-1} \bar{\Sigma} g_{K-1}, \bar{\Sigma} - \frac{\bar{\Sigma} g_{K-1} g_{K-1}^\top \bar{\Sigma}}{\sigma^2 + g_{K-1}^\top \bar{\Sigma} g_{K-1}} \right). \quad (32)$$

Applying Lemma 3 yields $z_{k,i} \in \text{Span}(A)$, thus, $\bar{x}_k = A A^\top \bar{x}_k$ and $\bar{\Sigma} = A \bar{\Sigma}_u A^\top$, we get the update rule reduced to

$$\begin{aligned} \mu_{k+1} &= A A^\top ((1 - w_k) \bar{\mu} + w_k \bar{z}_k) + \eta_{k-1} A \bar{\Sigma}_u A^\top g_{k-1} \\ &= A A^\top \bar{z}_k + \eta_k A \bar{\Sigma}_u A^\top \left(\nabla f(\bar{z}_k) - \eta_k^{-1} (1 - w_k) A \bar{\Sigma}_u^{-1} A^\top (\bar{z}_k - \bar{\mu}) \right). \end{aligned}$$

We set $w_k = 1 - \eta_k \lambda$ and set $\eta_k = \eta$, where $\lambda, \eta > 0$ will be specified later. Therefore, we have

$$\mu_{k+1} = A A^\top \bar{z}_k + \eta A \bar{\Sigma}_u A^\top \nabla h_\lambda(\bar{z}_k), \quad (33)$$

where $h_\lambda(x) := f(x) - (\lambda/2) \|x - \bar{\mu}\|_{\bar{\Sigma}^{-1}}^2$. Define $x_{A,\lambda}^* = \operatorname{argmax}_{x=Au} h_\lambda(x)$. With some similar steps in proof in Appendix C.2, by choosing $B_k \geq \zeta^{-4k}(1 - \zeta)^{-2}$, together with $z_0 = \bar{\mu}$, we get

$$\|\mu_K - x_{A,\lambda}^*\|_{\bar{\Sigma}^{-1}} \lesssim \zeta^K \left(\|\bar{\mu} - x_{A,\lambda}^*\|_{\bar{\Sigma}^{-1}} + M_1 \cdot \sqrt{d} \right), \quad \text{w.p. } 1 - \delta,$$

with $\eta = \frac{2}{L+2\lambda}$, $\zeta = |1 - \eta(L + \lambda)|$ and $M_1 = 2L\sqrt{(1 + \eta^2) \log\left(\frac{2K}{\delta}\right)}$. Also, we can get (30) as in proof in Appendix C.2. We restate it here:

$$f(x_{A,\lambda}^*) - f(\tilde{z}_K) \leq \frac{\lambda}{2} \left(\|x_{A,\lambda}^* - \bar{\mu}\|_{\bar{\Sigma}^{-1}}^2 - \|\mu_K - \bar{\mu}\|_{\bar{\Sigma}^{-1}}^2 \right) + \frac{L + \lambda}{2} \|\mu_K - x_{A,\lambda}^*\|_{\bar{\Sigma}^{-1}}^2. \quad (34)$$

Since f is concave,

$$f(x_A^*) - f(x_{A,\lambda}^*) \leq \langle \nabla f(x_{A,\lambda}^*), x_A^* - x_{A,\lambda}^* \rangle = \lambda \langle \bar{\Sigma}^{-1}(x_{A,\lambda}^* - \bar{\mu}), x_A^* - x_{A,\lambda}^* \rangle.$$

Adding (34), it holds

$$f(x_{A,\lambda}^*) - f(\mu_K) \leq \frac{\lambda}{2} \left(\|x_{A,\lambda}^* - \bar{\mu}\|_{\bar{\Sigma}^{-1}}^2 - \|x_{A,\lambda}^* - x_A^*\|_{\bar{\Sigma}^{-1}}^2 - \|\mu_K - \bar{\mu}\|_{\bar{\Sigma}^{-1}}^2 \right) + \frac{L + \lambda}{2} \|\mu_K - x_{A,\lambda}^*\|_{\bar{\Sigma}^{-1}}^2.$$

Due to (33), we have, it holds w.p. $1 - \delta$,

$$f(x_{A,\lambda}^*) - f(\mu_K) \lesssim \left[\lambda \|x_A^* - \bar{\mu}\|_{\bar{\Sigma}^{-1}}^2 + (L + \lambda) \zeta^K \left(\|\bar{\mu} - x_{A,\lambda}^*\|_{\bar{\Sigma}^{-1}}^2 + M_1 d \right) \right].$$

We choose $\lambda = L \log K / (4K)$ and get

$$f(x_{A,\lambda}^*) - f(\mu_K) \lesssim \frac{L \log K}{K} \cdot \left[\|x_A^* - \bar{\mu}\|_{\bar{\Sigma}^{-1}}^2 + \|\bar{\mu} - x_{A,\lambda}^*\|_{\bar{\Sigma}^{-1}}^2 + M_1 d \right], \quad \text{w.p. } 1 - \delta.$$

With assuming $\|x_{A,\lambda}^*\|$ is bounded, we derive

$$f(x_{A,\lambda}^*) - f(\mu_K) = \mathcal{O} \left(\frac{dL^2 \log K}{K} \cdot \log \left(\frac{K}{\delta} \right) \right), \quad \text{w.p. } 1 - \delta.$$

C.4 Auxiliary Lemma

The following is a standard result in convex optimization utilized in previous proofs.

Lemma 4. *Let f be α -strongly concave and β -smooth with respect to the (semi) norm $\|\cdot\|_{\Sigma^{-1}}$, for all x and y , it holds*

$$-\langle \nabla f(x) - \nabla f(y), x - y \rangle \geq \frac{\alpha\beta}{\alpha + \beta} \|x - y\|_{\Sigma^{-1}}^2 + \frac{1}{\alpha + \beta} \|\nabla f(x) - \nabla f(y)\|_{\Sigma}^2. \quad (35)$$

Proof. See Bubeck et al. (2015, Lemma 3.11) for a proof. □

□