1 Experiment with Data from Nonlinear Manifold

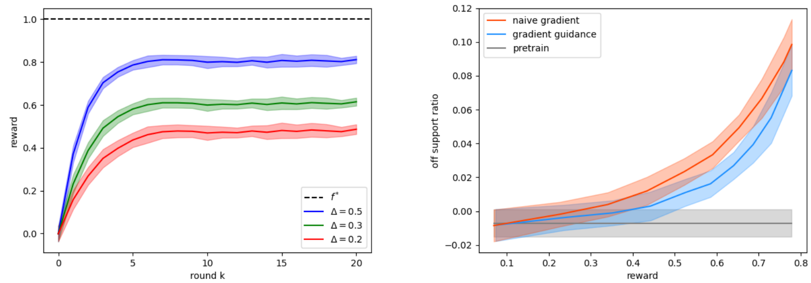

Figure 1: **Nonlinear data structure experiment.** We apply Alg.1 to data uniformly sampled from a unit ball in $\mathbb{R}^{64}$. The objective reward function is $f(x) = \theta^\top x$, where $\|\theta\| = 1$. **Left:** Rewards increase and converge with Alg.1. Higher guidance strength Δ (lower regularization) results in a higher convergent reward. **Right:** For the same reward, gradient guidance achieves a smaller deviation from the unit ball compared to the naive gradient. This indicates that gradient guidance can preserve data structure for nonlinear manifolds.

2 Time Efficiency

	Total runtime (iterations)	Per iteration	No guidance
Simulation	3.8 min (50 iter) | 76 min (1000 iter)	4.6 s	2.6 s
Image	1.3 min (5 iter)	15.8 s	4.9 s

Table 1: Runtime Efficiency of Algorithm 1. Red refers to the total time to converge. No guidance refers to the time for one-time inference of the pre-trained model.

3 Theory for the Failure of Naive Gradient

Lemma* For naive guidance $\mathtt{G}(X_t^{\leftarrow}, t) = b(t)\nabla f(X_t^{\leftarrow})$, suppose $b(t) > b_0 > 0$ for $t > t_0$. For data in subspace under Assumption 1 and reward $f(x) = g^\top x$, $g \perp Span(A)$ with $h(t) = 1 - \exp(-\sqrt{t})$, then the orthogonal component of the generated sample is consistently large:

$$\mathbb{E}[X_{T,\perp}^{\leftarrow}] = Cg, \quad C > \exp(-5/2)b_0.$$

PROOF. Under Assumption 1, the score can be decomposed to terms parallel and orthogonal to Span(A) (Prop 2, Appx D.3) Applying naive guidance, we examine the orthogonal reverse process:

$$\mathrm{d}X_{t,\perp}^{\leftarrow} = \left[\frac{1}{2} - \frac{1}{h(T-t)}\right] X_{t,\perp}^{\leftarrow}\mathrm{d}t + b(t)g\mathrm{d}t + \left(I_D - AA^\top\right)\mathrm{d}\overline{W}_t.$$

Solving this SDE, we get the expectation of the final state following $\mathbb{E}[X_{T,\perp}^{\leftarrow}] = \int_0^T \exp\left(-\int_0^t h^{-1}(s)\mathrm{d}s\right) e^{t/2} b(T-t) g \mathrm{d}t$. For the schedule $h(t) = 1 - \exp(-\sqrt{t})$, we have the coefficient of direction g is larger than $\int_0^T \exp(-t/2 - 2\sqrt{t}) b(T-t)\mathrm{d}t > \int_0^1 \exp(-5/2) b_0 \mathrm{d}t > 0$ where we can assume $T > 1$. Thus, $\mathbb{E}[X_{T,\perp}^{\leftarrow}] \neq 0$. This means the generated sample is going out of the subspace, i.e., naive gradient guidance will violate the latent structure.

4 Expanded Related Work

As suggest by reviewers, we expand our related work on "training-free guidance" and "direct latent optimization". (screen shot provided on the right).

Training-free Guidance. Training free guidance methods [15, 54, 66, 31, 48, 26] utilize the gradient of reward (or loss) function as guidance. [15, 54, 31, 26] is a line of works solving inverse problems on image and [66, 48] aims for guided/conditional image generation. Though not being developed for solving optimization problems, [15, 66] both propose a similar guidance to our guidance $\mathtt{G}_{loss}$: taking gradient on the predicted clean data x_0 with respect to x_t. Differently, our paper provides the first rigorous theoretical study of this gradient-based guidance. In addition, algorithmically, we propose an algorithm that iteratively applies $\mathtt{G}_{loss}$ as a module to the local linearization of optimization objective, which enjoys provable convergence guarantee.

Direct Latent Optimization in Diffusion Models. Besides guidance methods, an alternative training-free route by optimizing the initial value of reverse process [61, 5, 34, 58, 47]. These methods typically backpropagate the gradient of reward directly to the initial latent vector through an ODE solver, utilizing the chain rule. Thus at inference time, the reverse process is unchanged except for being fed with an optimized initialization, different from the guidance method we studied.