
On the Guidance of Flow Matching

Ruiqi Feng¹ Chenglei Yu^{* 1} Wenhao Deng^{* 1} Peiyan Hu² Tailin Wu¹

Abstract

Flow matching has shown state-of-the-art performance in various generative tasks, ranging from image generation to decision-making, where generation under energy guidance (abbreviated as guidance in the following) is pivotal. However, the guidance of flow matching is more general than and thus substantially different from that of its predecessor, diffusion models. Therefore, the challenge in guidance for general flow matching remains largely underexplored. In this paper, we propose the first framework of general guidance for flow matching. From this framework, we derive a family of guidance techniques that can be applied to general flow matching. These include a new training-free asymptotically exact guidance, novel training losses for training-based guidance, and two classes of approximate guidance that cover classical gradient guidance methods as special cases. We theoretically investigate these different methods to give a practical guideline for choosing suitable methods in different scenarios. Experiments on synthetic datasets, image inverse problems, and offline reinforcement learning demonstrate the effectiveness of our proposed guidance methods and verify the correctness of our flow matching guidance framework. Code to reproduce the experiments can be found at https://github.com/AI4Science-WestlakeU/flow_guidance.

1. Introduction

Flow matching has emerged as a prominent class of generative models. It features the ability to use a vector field

^{*}Equal contribution ¹Department of Artificial Intelligence, Westlake University, Hangzhou, China ²Academy of Mathematics and Systems Science, Chinese Academy of Sciences, Beijing, China. Correspondence to: Tailin Wu <wutailin@westlake.edu.cn>.

to transform samples from a source distribution into samples following a target distribution, thus realizing generative modeling (Lipman et al., 2023). The probability distribution the samples follow during the flow is called the probability path. By designing the probability path in a large design space, flow matching has shown improved generative modeling fidelity as well as higher sampling efficiency in a variety of generative modeling tasks including image generation (Lipman et al., 2023), decision-making (Zheng et al., 2023), audio generation, and molecular structure design (Gat et al., 2024; Chen & Lipman, 2024; Ben-Hamu et al., 2024). Flow matching substantially extends diffusion models (Ho et al., 2020; Song et al., 2021). Most diffusion models leverage the score matching process (Song & Ermon, 2019; Song et al., 2020; 2021), inherently limiting them to using the Gaussian distribution as the source distribution to construct a special probability path. Meanwhile, flow matching can learn the mapping between any source distribution and target distributions (Lipman et al., 2023; 2024; Chen & Lipman, 2024; Gat et al., 2024).

Guiding flow matching models refers to steering the generated samples toward desired properties, e.g., sampling from a distribution weighted with some objective function (Lu et al., 2023) or conditioned on class labels (Song et al., 2021)¹. It is vital in many generative modeling applications (Song et al., 2023b; Zheng et al., 2023), but contrary to well-studied guidance in diffusion models (Song et al., 2023b; Chung et al., 2023; Dhariwal & Nichol, 2021; Song et al., 2023a; Zheng et al., 2024; Lu et al., 2023; Dou & Song, 2024; Trippe et al., 2023), the guidance of flow matching remains less investigated. Most existing guidance methods only apply to a subset of flow matching that assumes the source distribution to be Gaussian and the probability path to have a certain simple form (Lipman et al., 2024; Zheng et al., 2023; Pogle et al., 2024; Zhang et al., 2025; Kollovich et al., 2025). In these cases, it is allowed to simplify the guidance of flow matching to be essentially the same as diffusion model guidance, but flow matching’s power of generating more flexible probability paths than diffusion models (Tong et al., 2024; Chen & Lipman, 2024; Gat et al., 2024) is restricted. There have been other controlled generation methods for flow matching, with a notable stream following the paradigm of optimizing some objective functions via

¹Mathematically, they are essentially equivalent (Section 3.1).

differentiating through the sampling process (Ben-Hamu et al., 2024; Liu et al., 2023b; Wang et al., 2025). However, their goal differs from our guidance of weighting the generated distribution. Therefore, the guidance for flow-matching models remains unrevealed in the rest of the ample design space.

To fill this gap, in this work, we start from a similar assumption as diffusion guidance and propose a general framework of flow matching guidance. From the perspective of this framework, we propose *Monte Carlo-based training-free asymptotically exact guidance* for flow matching. We also propose different training losses for exact *training-based guidance*, one of which covers existing losses as special cases (Lu et al., 2023; Zhang et al., 2025). For *approximate guidance methods*, we can theoretically derive from our framework many famous guidance methods that have appeared in the literature, including DPS (Chung et al., 2023), Π GDM (Song et al., 2023a), LGD (Song et al., 2023b), as well as their flow-matching extensions that are theoretically justified for general flow matching models. We demonstrate the effectiveness of our proposed method in both synthetic datasets and decision-making (offline RL) benchmarks. Furthermore, more extensive experiments are conducted on image inverse problems to provide an empirical comparison of different types of guidance methods for a guideline of choosing different methods.

We summarize our contributions as follows: **(1)** We propose a theoretically justified unified framework to construct guidance for general flow matching, *i.e.*, with arbitrary source distribution, coupling, and conditional paths. **(2)** The framework inspires us to propose a family of new guidance methods, including Monte Carlo sampling-based asymptotically exact guidance and training-based exact guidance for flow matching. **(3)** The framework can exactly cover multiple classical guidance methods in flow matching and diffusion models. Contrary to previous derivations relying on the flow to have a Gaussian source distribution, our derivation provides theoretical justification of these methods for general flow matching. **(4)** Empirical comparisons between guidance methods are conducted in different tasks, providing insights into choosing appropriate guidance methods for different generative modeling tasks.

2. Background

Let \mathbb{R}^d denote the data space where the data samples $x_t \in \mathbb{R}^d$. Here the subscript $t \in [0, 1]$ denotes inference time such that $p_1(x_1)$ is a target distribution we want to generate, and $p_0(x_0)$ is a base distribution that is easy to sample from. Flow-based generative models (Lipman et al., 2023; 2024) define a vector field $v_t(x_t) : [0, 1] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ that generates a probability path $p_t(x_t) : [0, 1] \times \mathbb{R}^d \rightarrow \mathbb{R}_{>0}$ connecting the tractable base distribution $p_0(x_0)$ and

the target distribution $p_1(x_1)$. By first sampling x_0 from $p_0(x_0)$ and then solving the ordinary differential equation (ODE) $\frac{d}{dt}x_t = v_t(x_t)$, one can generate clean samples $x_1 := x_t|_{t=1}$ that follow the target distribution $p_1(x_1)$.

An efficient way to learn the vector field $v_t(\cdot)$ by a model $v_\theta(\cdot, t)$ is to use flow matching (Lipman et al., 2023; 2024; Tong et al., 2024). It works by first finding a conditional vector field $v_{t|z}(x_t|z)$ that generates a *conditional* probability path $p_t(x_t|z)$, where z denotes sample pairs (x_0, x_1) ². The pairs (couplings) follow the probability distribution of $p(z) = \pi(x_0, x_1)$ ³. We use (x_0, x_1) and z interchangeably throughout the paper.

It has been proved that the *marginal* vector field

$$v_t(x_t) := \int v_{t|z}(x_t|z)p(z|x_t)dz,$$

where $p(z|x_t) = \frac{p_t(x_t|z)p(z)}{p(x_t)}$, will generate the *marginal* probability path $p_t(x_t) = \int p_t(x_t|z)p(z)dz$ (Lipman et al., 2023). Thus, one only needs to fit the *marginal* vector field $v_t(x_t) = \int v_{t|z}(x_t|z)p_t(z|x_t)dz$ using a model $v_\theta(x_t, t)$.

It is intuitive to construct the loss

$$\mathbb{E}_{t \sim \mathcal{U}(0,1), x_t \sim p(x_t)} \left[\left\| v_\theta(x_t, t) - \underbrace{v_t(x_t)}_{\text{intractable}} \right\|_2^2 \right],$$

which is, unfortunately, intractable. To cope with this problem, an equivalent conditional flow matching loss has been proposed (Lipman et al., 2023; Tong et al., 2024):

$$\mathbb{E}_{t \sim \mathcal{U}(0,1), z \sim p(z), x_t \sim p(x_t|z)} \left[\left\| v_\theta(x_t, t) - v_{t|z}(x_t|z) \right\|_2^2 \right],$$

which is tractable and can be used to train v_θ .

3. Guidance Vector Fields

This section is organized as follows: Section 3.1 proposes the general flow matching guidance framework. In Section 3.2, we derive a new guidance method g^{MC} based on Monte Carlo estimation, which is asymptotically exact. In Section 3.3, we derive a guidance g_t^{local} proportional to the gradient of the energy function J by approximating g_t with Taylor expansion. Then, we introduce the affine path assumption (Assumption 3.2) to obtain a tractable $g^{\text{cov}} \approx g_t^{\text{local}}$, and show that under the stronger uncoupled affine Gaussian path assumption (Assumption 3.3), g^{local} covers classical diffusion guidance methods. Section 3.4 introduces an alternative approximation that $p(z|x_t)$ is a Gaussian. Under this approximation and the affine path assumption (Assumption 3.2), we derive a derivative-free guidance $g^{\text{sim-MC}}$, and a specialized guidance of $g^{\text{sim-A}}$ for inverse problems. Finally,

²The notation z can also represent x_1 alone (Lipman et al., 2024). Our analysis is in the general setting where $z = (x_0, x_1)$, but for ease of interpretation, one may consider z as simply x_1 .

³We use π to denote the probability density of data couplings.

Section 3.5 provides different training losses for training-based exact guidance g_ϕ . Figure 1 visualizes the above outline and the relations between these guidance methods.

3.1. General Guidance Vector Fields

Guided generative modeling aims to generate samples according to specific requirements, and energy-guided sampling is one of the primary approaches. Appendix A.1 shows that energy guidance is fundamentally equivalent to conditional sampling or posterior sampling. Accordingly, we use the notation of energy guidance throughout, with the theoretical results directly extending to these tasks. Given an energy function $J : \mathbb{R}^d \rightarrow \mathbb{R}$ and a pre-trained generative model for $p(x)$, energy-guided samples follow $x \sim p'(x) = \frac{1}{Z} p(x) e^{-J(x)}$, where $Z = \int p(x) e^{-J(x)} dx$ is the normalizing constant. Samples with lower energy values J are more likely to be generated. Thus, the problem of guidance for flow-matching models becomes:

How can we alter the original vector field (VF) v_t that generates $p(x)$ such that the new VF $v_t'^4$ can generate samples from the new distribution $p'(x) = \frac{1}{Z} p(x) e^{-J(x)}$?

A natural choice that is commonly used in diffusion models (Dhariwal & Nichol, 2021) is to add a guidance VF $g_t(x_t) : [0, 1] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ to the original VF $v_t(x_t)$, such that the new VF $v_t'(x_t) = g_t(x_t) + v_t(x_t)$ is a VF formed with the same flow matching hyperparameters but generates the new probability path arriving at the new distribution $p'(x)$. Therefore,

$$g_t(x_t) = \int v_t'|_z(x_t|z) p'(z|x_t) dz - \int v_t|_z(x_t|z) p(z|x_t) dz,$$

where $p(z|x_t) = \frac{p_t(x_t|z)p(z)}{p(x_t)}$ and $p'(z|x_t) = \frac{p_t'(x_t|z)p'(z)}{p'(x_t)}$. Recall that new VF $v_t'(x_t)$ has the conditional probability path and conditional VF as that of the original VF, i.e., $v_t'(x_t|z) = v_t(x_t|z)$ and $p_t'(x_t|z) = p_t(x_t|z)$. Then, we have the following theorem (proof in Appendix A.2).

Theorem 3.1. *Adding the guidance VF $g_t(x_t)$ to the original VF $v_t(x_t)$ will form VF $v_t'(x_t)$ that generates $p_t'(x_t) = \int p_t(x_t|z) p'(z) dz$, as long as $g_t(x_t)$ follows:*

$$g_t(x_t) = \int \left(\mathcal{P} \frac{e^{-J(x_1)}}{Z_t(x_t)} - 1 \right) v_t|_z(x_t|z) p(z|x_t) dz, \quad (1)$$

$$\text{where } Z_t(x_t) = \int \mathcal{P} e^{-J(x_1)} p(z|x_t) dz, \quad (2)$$

$\mathcal{P} = \frac{\pi'(x_0|x_1)}{\pi(x_0|x_1)}$ is the reverse coupling ratio, where $\pi'(x_0|x_1)$ is the reverse data coupling for the new VF, i.e., the distribution of x_0 given x_1 sampled from the target distribution.

⁴The prime symbol ' denotes probability distributions and vector fields corresponding to the new distribution $p'(x)$.

In this paper, we consider the case where \mathcal{P} is or can be approximated as 1. \mathcal{P} is exactly 1 when the coupling is independent $\pi(x_0, x_1) = p(x_0)p(x_1)$, which results in $\pi'(x_0|x_1) = \pi(x_0|x_1) = p(x_0)$. $\mathcal{P} = 1$ is also reasonable in many practical flow matching methods with dependent couplings, such as mini-batch optimal transport (OT) conditional flow matching (Tong et al., 2024), which we elaborate in Appendix A.3. However, there may be cases where the VF of dependent couplings is significantly different from that of independent couplings. In these cases, the impact of \mathcal{P} can be non-negligible, which we leave for future work. Later, we will show this approximation allows us to cover many existing guidance techniques for flow matching or diffusion models.

Uncoupled affine Gaussian path flow matching, where $\pi(x_0, x_1) = p(x_0)p(x_1)$, $v_t|_1(x_t|x_1) = \alpha_t x_t + \beta_t x_1$, and $p(x_0) = \mathcal{N}(x_0; \mu, \Sigma)$, is known to be equivalent to diffusion models, with little difference in the noise schedule (Zheng et al., 2023; Ma et al., 2024). In this case, our general guidance for flow matching Eq. (1) can be reduced to a commonly used guidance term in diffusion models: $g_t(x_t) \propto \nabla_{x_t} \log Z_t(x_t)$ (proof in Appendix A.4). Thus, most existing works that only consider uncoupled affine Gaussian path flow matching essentially apply the same guidance techniques for diffusion models (Dhariwal & Nichol, 2021; Song et al., 2023b; Chung et al., 2023).

Next, we will explore more challenging scenarios of flow matching guidance, which are substantially different from diffusion guidance: either the coupling is dependent, or the source distribution is non-Gaussian.

3.2. Monte Carlo Estimation, g_t^{MC}

To start with, we discuss the Monte Carlo (MC) method to estimate the guidance $g_t(x_t)$ in Eq. (1), which is asymptotically exact while being training-free. The MC estimation of the integrals in Eq. (1) and (2) requires sampling from intractable $p(z|x_t)$, but they can be converted into:

$$\begin{aligned} g_t^{\text{MC}}(x_t) &\triangleq \mathbb{E}_{x_1, x_0 \sim p(z)} \left[\left(\frac{e^{-J(x_1)}}{Z_t} - 1 \right) v_t|_z(x_t|z) \frac{p_t(x_t|z)}{p_t(x_t)} \right], \quad (3) \\ Z_t(x_t) &= \mathbb{E}_{x_1, x_0 \sim p(z)} \left[e^{-J(x_1)} \frac{p_t(x_t|z)}{p_t(x_t)} \right], \end{aligned}$$

where only $p_t(x_t)$ remains intractable. We can use the same MC samples to self-normalize the distribution $p_t(x_t|z)$, using $p_t(x_t) = \mathbb{E}_{x_1, x_0 \sim p(z)} [p_t(x_t|z)]$. In the expressions above, $p_t(x_t|z)$ is usually designed to be a simple, known distribution (Tong et al., 2024) and $p(z)$ can be sampled either using the learned generative model or from the training distribution if accessible. The above method can be understood as executing importance sampling to convert the expectation under the intractable distribution $\mathbb{E}_{z \sim p(z|x_t)}[\cdot]$ to

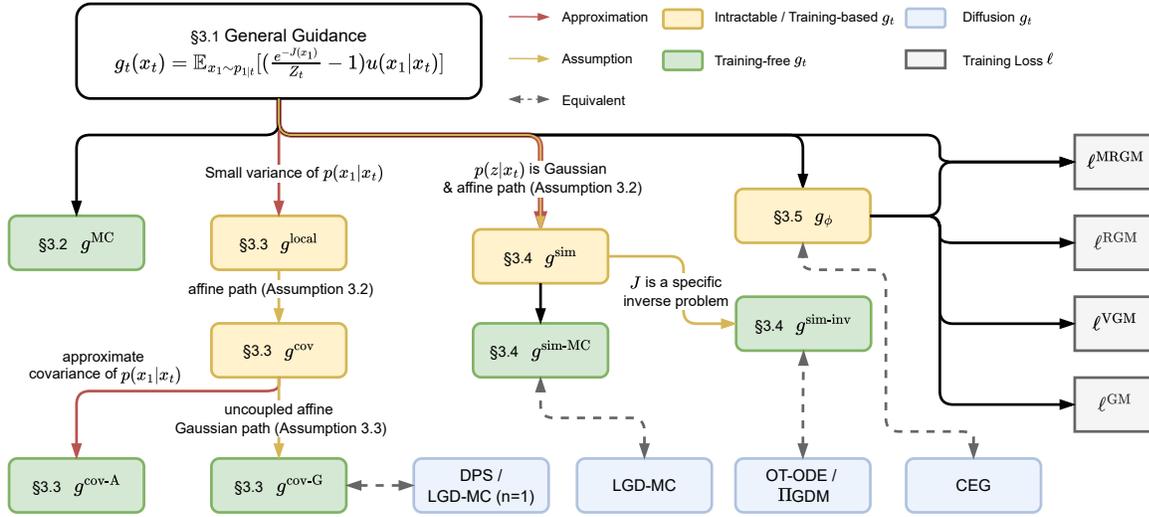


Figure 1: Overview of guidance methods in the paper. We start with a unified guidance expression and derive different guidance methods, including training-free and training-based methods, and cover many classical diffusion guidances.

that under a tractable distribution $\mathbb{E}_{z \sim p(z)} \left[\frac{p(x_t | z)}{p(x_t)} (\cdot) \right]$. The pseudocode for computing $g_t^{\text{MC}}(x_t)$ can be found in Algorithm 1, and a simplified version of g^{MC} under the assumption of independent coupling is provided in Appendix A.8.

It should be noted that this method is unbiased and applicable to *any* source distribution. This enables the guidance for flow matching with different source distributions such as uniform (Chen & Lipman, 2024), Gaussian process (Andrae et al., 2025), mixture of Gaussian (Hiranaka et al., 2025; Papamakarios et al., 2017), and others (Mathieu & Nickel, 2020; Stimper et al., 2022). This method can also be applied to flow matching with dependent couplings of x_1, x_0 , e.g., optimal transport couplings (Tong et al., 2024) and rectified flow (Liu et al., 2023a).

To our knowledge, g_t^{MC} is the *first* to provide an asymptotically exact training-free estimation of the guidance VF for flow matching whose source distribution is non-Gaussian. However, due to the high variance of MC given a limited number of samples, g^{MC} is more suitable for tasks where the energy function J varies gently and the data is relatively low-dimensional.

Since the efficiency of g^{MC} is restricted by the high variance in the MC estimation, it has an unsatisfactory performance in high-dimensional generation tasks like image inverse problems. To improve the scalability of g^{MC} , there are many techniques that can be readily applied. For example, we can adopt importance sampling in Eq. (3) to reduce the variance

of the MC estimation:

$$g_t^{\text{MC-IS}}(x_t) \triangleq \mathbb{E}_{x_1, x_0 \sim \tilde{p}(z)} \left[\frac{p(z)}{\tilde{p}(z)} \left(\frac{e^{-J(x_1)}}{Z_t} - 1 \right) \frac{p_t(x_t | z)}{p_t(x_t)} v_t(x_t | z) \right],$$

$$Z_t(x_t) = \mathbb{E}_{x_1, x_0 \sim \tilde{p}(z)} \left[\frac{p(z)}{\tilde{p}(z)} e^{-J(x_1)} \frac{p_t(x_t | z)}{p_t(x_t)} \right],$$

where $\tilde{p}(z)$ is an alternative distribution that can be sampled from. To enhance the scalability of $g^{\text{MC-IS}}$, we need to select a \tilde{p} such that $\frac{e^{-J(x_1)}}{\tilde{p}(z)}$ has lower variance, i.e., when $e^{-J(x_1)}$ is large, $\tilde{p}(z)$ is also large, and vice versa. This can be achieved by sampling using another guided VF, such as those to be proposed subsequently, which generates samples close to (but not exactly) $\frac{1}{Z} p(x_1) e^{-J(x_1)}$. The probability density ratio $\frac{p(z)}{\tilde{p}(z)}$ can be estimated using, for example, the Hutchinson trace estimator to preserve scalability (Song et al., 2021). It should be noted that this estimation is still unbiased.⁵

3.3. Localized $p(x_1 | x)$

Many practical guidance methods rely on approximations (Song et al., 2023b; Pokle et al., 2024), so contrary to the unbiased MC estimation of g_t , we investigate approximate (and thus biased) guidance methods in this subsection. We start from an intuitive assumption that the probability mass of $p(x_1 | x_t)$ is centered around its mean. Following this,

⁵We leave empirical investigation into $g^{\text{MC-IS}}$ to future work.

it is natural to approximate the integrals in Eq. (1) with the Taylor expansion that captures local behaviors around the mean. Thus, Eq. (1) can be simplified and becomes tractable.

First, we approximate the normalizing constant $Z_t(x_t)$ in Eq. (2) as (Appendix A.9):

$$Z_t(x_t) = \int p(z|x_t) e^{-J(x_1)} dz \approx e^{-J(\hat{x}_1)} \quad (4)$$

where $\hat{x}_1 := \mathbb{E}_{x_0, x_1 \sim p(z|x_t)}[x_1]$. Likewise, Eq. (1) can be approximated as $g_t(x_t) \approx g_t^{\text{local}}$, and:

$$g_t^{\text{local}}(x_t) \triangleq -\mathbb{E}_{z \sim p(z|x_t)} [(x_1 - \hat{x}_1) v_{1|t}(x_t|z)] \nabla_{\hat{x}_1} J(\hat{x}_1),$$

where \hat{x}_1 is defined as above. The approximation error $\|\delta g\|_2^2$ can be bounded: $\|\delta g\|_2^2 \leq |\lambda_h \sigma_1 d / e^{-J(\hat{x}_1)}|^2 (C_1 + C_2)$ (Appendix A.9), where λ_h is the maximum eigenvalue of the Hessian of $e^{-J(x)}$, σ_1 is the L_2 norm of the covariance matrix of the distribution $p(x_1|x_t)$, d is the dimensionality of the data, and C_1, C_2 are constants dependent on the variance of the original conditional VF and the guided VF.

The error bound gives insights on the approximation quality of g_t^{local} (detailed discussion in Appendix A.9):

- The error is small when J is smooth, in which case the Hessian of $e^{-J(x)}$ will approach zero. This corresponds to the mild guidance, where approximation-based g^{local} works well.
- The error is small when σ_1 is small, *i.e.* the covariance matrix Σ_{11} has a small Frobenius norm. This is the case when the flow time $t \rightarrow 1$ (and $\sigma_t = 0$), where x_t predicts x_1 well.

The gradient in g_t^{local} is a natural outcome of the Taylor expansion near the mean of $p(x_1|x_t)$. g_t^{local} is not only applicable to more general flow matching but also originates differently from diffusion guidance (Dhariwal & Nichol, 2021) where the gradient naturally emerges from the score function $\nabla_{x_t} \log p(x_t)$. Moreover, the error bound we provide here is more practical compared to those previously proposed for diffusion guidance, *e.g.*, the Jensen gap in diffusion posterior sampling (DPS) (Chung et al., 2023), which only bounds the error in $\mathbb{E}[e^{-J(x_1)}]$, but it is $\nabla_{x_t} \log Z_t$ that is the guidance VF, whose error is not bounded.

In order to obtain \hat{x}_1 , we need the following assumption:

Assumption 3.2. The affine path assumption. We assume the conditional probability path to be affine, *i.e.*, $x_t = \alpha_t x_1 + \beta_t x_0 + \sigma_t \varepsilon$, where ε is a random noise and $\sigma_t, \dot{\sigma}_t$ are both sufficiently small.⁶

⁶This choice is a widely used one (Lipman et al., 2023; Tong et al., 2024). With a small random noise $\sigma_t \varepsilon$, x_t under conditional VF flows is almost exactly from x_0 to x_1 .

Note that this assumption does not prevent the samples from having dependent coupling: $\pi(x_0|x_1) \neq p(x_0)$. Under Assumption 3.2, we can use the x_1 -parameterization (Lipman et al., 2024) to express \hat{x}_1 with the VF v_t that is learned by the model v_θ (Appendix A.10):

$$\hat{x}_1 \approx -\frac{\dot{\beta}_t}{\dot{\alpha}_t \beta_t - \dot{\beta}_t \alpha_t} x_t + \frac{\beta_t}{\dot{\alpha}_t \beta_t - \dot{\beta}_t \alpha_t} v_t. \quad (5)$$

With the commonly chosen schedule $\alpha_t = t, \beta_t = 1 - t$ (Lipman et al., 2023; Tong et al., 2024), $\hat{x}_1 = x_t + (1 - t)v_t$ coincides with the 1-step generated x_1 under the Euler scheme. Also under this affine path assumption, g_t^{local} can be expressed with the covariance matrix of $p(x_1|x_t)$, $\Sigma_{1|t}$, and the gradient $\nabla_{\hat{x}_1} J(\hat{x}_1)$ (proof in Appendix A.11):

$$g_t^{\text{cov}} \triangleq -\underbrace{\frac{\dot{\alpha}_t \beta_t - \dot{\beta}_t \alpha_t}{\beta_t}}_{\text{schedule}} \Sigma_{1|t} \nabla_{\hat{x}_1} J(\hat{x}_1). \quad (6)$$

Intuitively, the guidance is the gradient of estimated J preconditioned with the covariance matrix of $p(z|x_t)$. The preconditioning squeezes the guidance vector into the $p(z|x_t)$ manifold. Next, we discuss different ways to obtain this covariance term, resulting in $g^{\text{cov-A}}, g^{\text{cov-L}}$, and $g^{\text{cov-G}}$.

The simplest way is to approximate the covariance matrix with a manually set schedule $\lambda'_t I$, resulting in $g_t^{\text{cov-A}}$. This allows us to tune the guidance’s schedule, a common practice in diffusion model guidance (Song et al., 2023b;a).

$$g_t^{\text{cov-A}} \triangleq -\lambda_t^{\text{cov-A}} \nabla_{\hat{x}_1} J(\hat{x}_1),$$

where the original schedule in g_t^{cov} is already included in the hyperparameter $\lambda_t^{\text{cov-A}}$. Since $p(x_1|x_t)$ is localized when $t \rightarrow 1$, a general guideline is to set $\lambda_t^{\text{cov-A}}$ decaying. $g_t^{\text{cov-A}}$ only assumes affine path (Assumption 3.2), and is thus theoretically justified for mini-batch optimal transport conditional flow matching (Tong et al., 2024) for which few existing guidance methods have been proposed. Besides, $g_t^{\text{cov-A}}$ is efficient as its computation involves no extra number of function evaluations (NFE) than unguided sampling.

Alternatively, we can use Proposition 3.5 to train a model to fit the actual covariance matrix of $p(x_1|x_t)$ and acquire $g_t^{\text{cov-L}}$. Note that this covariance matrix is determined by the original distribution and is agnostic to the energy function J . The original flow matching essentially learns $\mathbb{E}_{x_1 \sim p(x_1|x_t)}[x_1]$, but to achieve better approximate guided generation, more detailed information of the distribution $p(x_1|x_t)$, its covariance, also needs to be learned.

The important special case where flow matching is a diffusion model with a new schedule is formalized with:

Assumption 3.3. The uncoupled affine Gaussian path assumption. In addition to Assumption 3.2, the source distribution $p(x_0)$ is a standard Gaussian and the coupling is independent $\pi(x_0, x_1) = p(x_0)p(x_1)$.

We will demonstrate that under Assumption 3.3, g_t^{cov} simplifies into $g_t^{\text{cov-G}}$ (Eq. (7)), which covers classical guidance methods in diffusion models. Specifically, using the second-order Tweedie’s formula, we can express the covariance matrix $\Sigma_{1|t}$ in Eq. (6) in terms of the Hessian of the log-probability, $\nabla_{x_t} \nabla_{x_t} \log p(x_t)$ (Rozet et al., 2024; Boys et al., 2023; Ye et al., 2024). Then, under Assumption 3.3, $\nabla_{x_t} \log p(x_t)$ depends affinely on the VF v_t in Eq. (5) which depends affinely on \hat{x}_1 . Therefore, the covariance matrix $\Sigma_{1|t}$ can be expressed with the derivative of the Jacobian matrix $\frac{\partial \hat{x}_1}{\partial x_t}$. We refer to this relationship as the Jacobian trick (proof in Appendix A.12):

Proposition 3.4. The Jacobian trick. *Under Assumption 3.3, the inverse covariance matrix of $p(x_1|x_t)$, $\Sigma_{1|t}$, depends affinely on the Jacobian of the VF $\frac{\partial v_t}{\partial x_t}$, and is proportional to the Jacobian $\frac{\partial \hat{x}_1}{\partial x_t}$:*

$$\Sigma_{1|t} = \frac{\beta_t^2}{\alpha_t(\dot{\alpha}_t\beta_t - \dot{\beta}_t\alpha_t)}(-\dot{\beta}_t + \beta_t \frac{\partial v_t}{\partial x_t}) = \frac{\beta_t^2}{\alpha_t} \frac{\partial \hat{x}_1}{\partial x_t}.$$

Inserting back to Eq. (6), we have:

$$g_t^{\text{cov-G}} \triangleq \lambda_t^{\text{cov-G}} \nabla_{x_t} J(\hat{x}_1), \quad (7)$$

where $\lambda_t^{\text{cov-G}} = -\beta_t(\dot{\alpha}_t\beta_t - \dot{\beta}_t\alpha_t)/\alpha_t$ is the schedule.

$g_t^{\text{cov-G}}$ covers classical diffusion model guidance including loss-guided diffusion⁷ (LGD) (Song et al., 2023b) and diffusion posterior sampling (DPS) (Chung et al., 2023). In diffusion models, the guidance can be expressed as $\nabla_{x_t} \log \mathbb{E}_{z \sim p(z|x_t)}[e^{-J(x_1)}]$, and DPS approximates this by neglecting the gap in the Jensen inequality (Chung et al., 2023) to move the expectation into $e^{J(\cdot)}$, and LGD approximates the expectation with a point estimation. Both methods arrive at $-\nabla_{x_t} J(\hat{x}_1)$, which is covered by our Eq. (7).

3.4. Gaussian Approximation of $p(z|x_t)$

Instead of approximating guidance by expanding it near its mean, we can alternatively approximate $p(z|x_t)$ with a Gaussian distribution $\tilde{p}(z|x_t) = \mathcal{N}(z; \mu_t(x_t), \Sigma_t(x_t))$, and get $g_t \approx g_t^{\text{sim}}$. To minimize the KL divergence between the approximate distribution and the original distribution, we need to choose $\mu_t(x_t)$ and $\Sigma_t(x_t)$ to that of the actual distribution $p(z|x_t)$ (Rozet et al., 2024; Boys et al., 2023). Under Assumption 3.2, \hat{x}_1 can be estimated using Eq. (5), \hat{x}_0 can be similarly estimated (Appendix A.10), and $\Sigma_t(x_t)$ can be either set as a hyperparameter as in $g_t^{\text{cov-A}}$ or computed as in $g_t^{\text{cov-G}}$ in uncoupled affine Gaussian path flow matching.

Now that if J has an unknown expression, we can use Monte Carlo (MC) sampling to estimate $g_t^{\text{sim}} \approx g_t^{\text{sim-MC}}$, since we

⁷We refer to the simplest version of LGD-MC($n = 1$) as LGD here. LGD-MC with $n > 1$ will be covered in Section 3.4.

can sample from $\tilde{p}(z|x_t)$ which is a Gaussian distribution:

$$g_t^{\text{sim-MC}} \triangleq \sum_i^N \left(\frac{e^{-J(x_1^i)}}{\tilde{Z}_t} - 1 \right) v_{t|z}(x_t|z^i), \quad (8)$$

where $\tilde{Z}_t := \frac{1}{N} \sum_i^N e^{-J(x_1^i)}$ is an estimated normalizing constant. This shares the spirit of LGD-MC (Song et al., 2023b), which uses MC to estimate Z_t and then computes diffusion guidance $\nabla_{x_t} \log Z_t$.

In some guided generation tasks like inverse problems where J is known analytically, we can derive analytical expressions of g_t^{sim} . For example, when the measurement involves a known degradation operator H applied to x and then adding Gaussian noise with scale σ_y to Hx , we have $J(x_1) \propto \frac{1}{2\sigma_y^2} \|y - Hx\|_2^2$. Thus we can derive g_t^{sim} under the affine path assumption (Assumption 3.2), and propose a practical approximation $g_t^{\text{sim-inv-A}}$ (Appendix A.13):

$$g_t^{\text{sim-inv-A}} \triangleq -\lambda_t \left(\frac{\sigma_y^2}{r_t^2} + H^T H \right)^{-1} H^T (y - H\hat{x}_1), \quad (9)$$

where λ_t and r_t are both hyperparameters. $g_t^{\text{sim-inv-A}}$ extends IIGDM (Song et al., 2023a) to affine flow matching but requires a further approximation $\frac{\partial \hat{x}_1}{\partial x_t} \approx I$, which is accurate when $t \rightarrow 1$.

In uncoupled affine Gaussian path (Assumption 3.3), g_t^{sim} in our framework can exactly cover IIGDM and OT-ODE (Song et al., 2023a; Pokle et al., 2024) (Appendix A.13). Note that our g_t^{sim} is theoretically justified for dependent couplings, such as optimal transport conditional flow matching (OT-CFM) (Tong et al., 2024).

3.5. Training-based Guidance g_ϕ

Previously, we have discussed training-free guidance methods. In this subsection, we discuss how to train a neural network g_ϕ to fit guidance g_t . To construct a tractable training loss, we extend the conditional loss in flow matching to arbitrary conditional variables in the following proposition (proof in Appendix A.5).

Proposition 3.5. *Any marginal variable $f(x_t, t) := \mathbb{E}_{z \sim p_t(z|x_t)}[f_{t|z}(x_t, z, t)]$, $z = (x_0, x_1)$ has an intractable marginal loss:*

$$\mathcal{L}_t = \mathbb{E}_{x_t \sim p(x_t)} \left[\left\| f_\theta(x_t, t) - \mathbb{E}_{z \sim p_t(z|x_t)}[f_{t|z}(x_t, z, t)] \right\|_2^2 \right],$$

whose gradient w.r.t. θ is identical to that of the tractable conditional loss:

$$\mathcal{L}_{t|z} = \mathbb{E}_{x_t, z \sim p(x_t, z)} \left[\left\| f_\theta(x_t, t) - f_{t|z}(x_t, z, t) \right\|_2^2 \right].$$

Guidance Matching. Based on proposition 3.5, We can train g_ϕ by first learning a surrogate model $Z_{\phi_Z}(x_t, t) \approx Z_t$, and then train $g_\phi \approx g_t$ with the following guidance

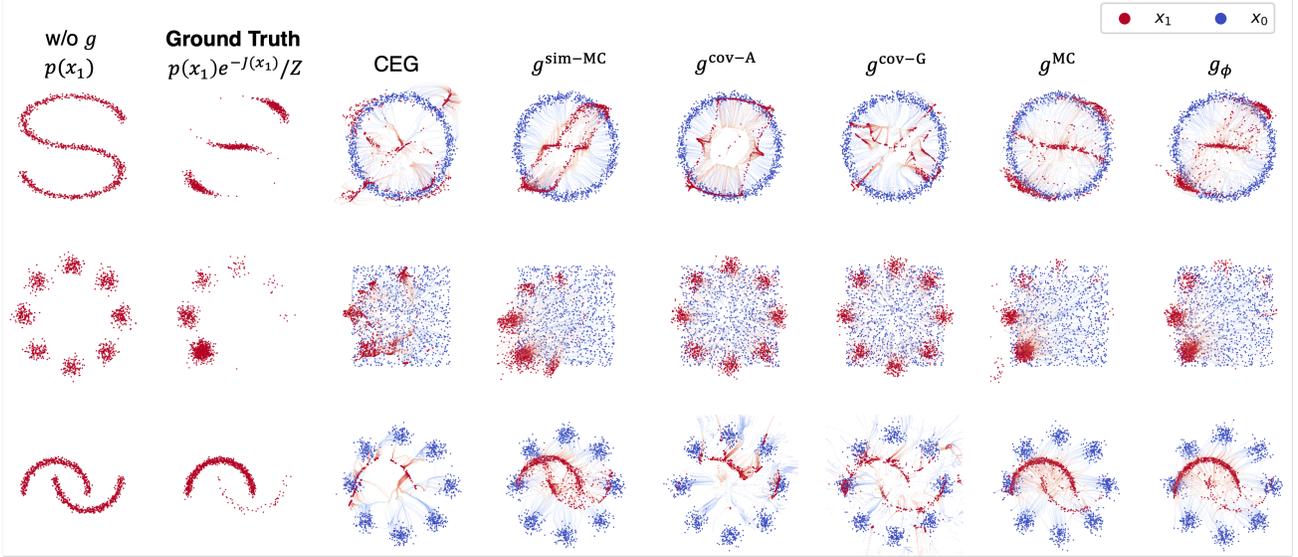


Figure 2: Results of the synthetic dataset with different source (blue) and target (red) distributions. We visualize the start/end points and the flow trajectories. g^{MC} and g_{ϕ} yield the best guidance across different settings while diffusion guidance fails.

matching (GM) losses.

$$\mathcal{L}_{\phi_Z, \phi} = \mathbb{E}_{t \sim \mathcal{U}(0,1), z \sim p(z), x_t \sim p(x_t|z)} [\ell_{\phi_Z} + \ell_{\phi}^{\text{GM}}], \quad (10)$$

$$\ell_{\phi}^{\text{GM}} = \left\| g_{\phi}(x_t, t) - \left(\frac{e^{-J(x_t)}}{Z_{\phi_Z, sg}(x_t, t)} - 1 \right) v_{t|z}(x_t|z) \right\|_2^2,$$

$$\ell_{\phi_Z} = \left\| Z_{\phi_Z}(x_t, t) - e^{-J(x_t)} \right\|_2^2, \quad (11)$$

where sg denotes stopping the gradient in automatic differentiation. We prove in Appendix A.7 that the minimizer of $\mathcal{L}_{\phi_Z, \phi}$ is indeed $Z_{\phi_Z} = Z_t$ and $g_{\phi} = g_t$.

In fact, there are other methods to learn Z_{ϕ_Z} . In the literature on diffusion models, Lu et al. (2023) proposed to use contrastive learning to obtain $Z_t(x_t)$ that can be extended to include the general flow matching path (Appendix A.6). Alternatively, Monte Carlo estimation can be applied to obtain Z_{ϕ_Z} (Appendix A.6). As for learning g_{ϕ} , alternative to the loss ℓ_{ϕ}^{GM} in Eq. (11), there are other losses ℓ^{VGM} , ℓ^{RGM} , and ℓ^{MRGM} which, when substituted into Eq. (10), will produce the same minimizer. We provide detailed analysis and proof for all these losses in Appendix A.7. Notably, ℓ^{MRGM} that we derive is identical to a newly proposed training-based flow matching guidance loss in Zhang et al. (2025). For uncoupled affine Gaussian path flow matching (Assumption 3.3), since the guidance is $g_t \propto \nabla_{x_1} \log Z_t(x_t)$ (Appendix A.4), learning Z_{ϕ_Z} in Eq. (11) is adequate, and Lu et al. (2023) also learns Z_{ϕ_Z} .

These different training-based guidance opens the design space of classifier-guidance (Dhariwal & Nichol, 2021) for flow matching. For diffusion models, one only needs to train a classifier on noisy input to produce the correct guidance,

whereas for general flow matching, one needs the training loss proposed above to train a network to get the accurate guidance.

In summary, we have proposed guidance matching methods for training-based guidance in this subsection. We derive different losses for learning Z_t and g_t , which all provide unbiased estimations of the gradient and can be utilized without specific assumptions on flow matching.

4. Experiments

In the experiment, we benchmark different guidance methods for flow matching models in different tasks including synthetic datasets, generative decision-making, and image inverse problems. These tasks all fall into the category of energy guidance (Janner et al., 2022; Lu et al., 2023) or posterior sampling (Chung et al., 2023). With these experiments, we aim to answer the following questions: **(1)** Can the proposed learning-based exact guidance method g_{ϕ} learn the correct guidance VF g_t for general (including non-uncoupled affine Gaussian path) flow matching? **(2)** For the asymptotically exact MC guidance method, can it produce the correct guidance in non-Gaussian flow matching, and is it exact assuming a sufficiently large sample budget? **(3)** On the practical aspect, how do different types of guidance methods (approximate/exact, training-free/training-based) perform on more realistic tasks, and how do we choose the appropriate flow matching guidance in different tasks?

Table 1: Results of the D4RL Locomotion experiments. Entries with $\geq 95\%$ score of the best results per task (excluding baselines) are highlighted in bold. The standard deviation is deferred to Appendix B.2 due to limited space.

		Baselines		OT-CFM						CFM					
		BC	Diffuser	w/o g	$g^{\text{cov-A}}$	$g^{\text{cov-G}}$	$g^{\text{sim-MC}}$	g^{MC}	g_ϕ	w/o g	$g^{\text{cov-A}}$	$g^{\text{cov-G}}$	$g^{\text{sim-MC}}$	g^{MC}	g_ϕ
Medium-Expert	HalfCheetah	55.2	88.9	61.9	64.8	73.2	78.1	86.4	70.2	46.4	63.4	68.5	83.5	87.7	81.5
	Hopper	52.5	103.3	95.2	101.8	112.3	112.3	112.7	98.1	83.4	93.9	113.3	88.5	113.3	91.5
	Walker2d	107.5	106.9	79.1	97.3	107.2	101.0	107.2	91.3	65.7	100.4	106.9	107.0	107.1	102.3
Medium	HalfCheetah	42.6	42.8	34.7	42.2	42.9	43.1	43.1	43.4	41.8	43.6	43.3	43.8	43.8	43.8
	Hopper	52.9	74.3	63.3	75.1	89.8	76.2	79.8	79.7	73.2	79.1	82.7	82.1	88.0	85.2
	Walker2d	75.3	79.6	72.4	82.7	81.3	83.4	83.0	80.6	72.2	80.7	82.5	81.9	81.9	72.9
Medium-Replay	HalfCheetah	36.6	37.7	25.6	31.7	36.1	36.8	40.0	35.5	22.2	33.4	39.3	37.9	40.6	39.1
	Hopper	18.1	93.6	40.1	57.7	74.1	60.9	88.6	55.3	55.1	63.0	69.3	61.0	80.9	63.5
	Walker2d	26.0	70.6	31.2	62.5	82.5	64.4	88.1	52.4	28.3	64.9	76.6	58.9	70.9	70.3
Average		51.9	77.5	55.9	68.4	77.7	72.9	81.0	67.4	54.3	69.2	75.8	71.6	79.4	72.2

4.1. Synthetic Dataset

We first compare different guidance methods on 2-dimensional synthetic datasets where the source distributions are different from Gaussian. The base flow matching model is trained to learn the flow with source distributions other than Gaussian. During inference, different guidance methods are applied to perform guided sampling with different objective functions J for each dataset. All of the source distributions are non-Gaussian, so traditional diffusion guidance should not be applied. In Figure 2, we compare the performance of g^{MC} , g_ϕ , an exact diffusion guidance called contrastive energy guidance (CEG) proposed by Lu et al. (2023), and approximate guidance $g^{\text{cov-A}}$, $g^{\text{cov-G}}$, and $g^{\text{sim-MC}}$. The original target distributions (w/o g_t) and the J -weighted distributions (ground truth) are shown in the first and second columns. The details of the experiment are provided in Appendix B.

It can be seen from Figure 2 that g^{MC} and g_ϕ generated samples that almost exactly match the ground truth distribution and the performance is consistent across different datasets. Note that the generated samples maintain the correct data manifold instead of concentrating on some points as gradient-based approximate guidance methods do. As has been mentioned in Section 3, CEG is essentially $\nabla_{x_t} \log Z_t$ which is exact under the uncoupled affine Gaussian path assumption 3.3. However, none of the flow matching paths we have here are uncoupled affine Gaussian paths, so exact diffusion guidance performs poorly compared to g_ϕ and g^{MC} , showing a largely distorted generated distribution.

Also, we investigated the asymptotic exactness of g^{MC} . To quantitatively see the increasing guidance precision as the sample size increases, we show the Wasserstein-2 distance between the guided generation distribution and the ground truth target distribution, estimated using 1000 samples. As shown in Figure 4, the error decreases as the number of samples for computing g^{MC} (N in Algorithm 1) increases from 5 to 10^4 , and eventually approaches or surpasses the error

of the learned generative model for the original distribution, which can be seen as an approximate lower-bound of the guided generation error.

4.2. Planning

We also conduct experiments on offline RL tasks where generative models have been used as planners (Janner et al., 2022; Chen & Lipman, 2024). The planning process is realized through sampling from $\frac{1}{Z} p(x_1) e^{R(x_1)}$ (Levine, 2018), which aligns with the goal of our guidance by setting $J = -R$, and R being the return. We report experiment results on the Locomotion tasks in the D4RL dataset (Fu et al., 2020), with experiment setting details and complete results provided in Appendix B.2.

The average normalized scores across five seeded runs of each guidance method are reported in Table 1 where score = 100 corresponds to the scores of the expert. The conclusions for CFM and OT-CFM are consistent: for gradient-based methods, $g^{\text{cov-G}}$ is generally better than $g^{\text{cov-A}}$ with an average increase in score of 8.0. $g^{\text{sim-MC}}$ has a performance between $g^{\text{cov-A}}$ and $g^{\text{cov-G}}$. The improved performance of $g^{\text{cov-G}}$ comes at a higher computation cost of differentiation through the VF model, though. The MC-based guidance, although being gradient-free, outperforms the second-best method $g^{\text{cov-G}}$ by 3.5 on average and is the best method in 7 out of 9 tasks. For g_ϕ , we report the result of the best losses ℓ , but their performance is still relatively weak, falling behind the best by 10.4 on average. We attribute this to the unstable joint training of two networks and provide the complete results in Appendix B.2.

To investigate the effectiveness of g^{MC} , we collect an ensemble of plans that are generated under guidance, compute the critic-predicted objective function value (estimated return R), and then plot the density distribution of the estimated return R . An ideal guidance will result in the R distribution to be $p(R) e^R$ where $p(R)$ is the distribution generated without guidance (Appendix B.2). As can be seen from Figure 3,

Table 2: Image inverse problem results on CelebA-HQ. The best and runner-up results are highlighted in bold and underlined.

		Inpainting-Center, $\sigma_y = 0.05$				Super-Resolution $\times 4$, $\sigma_y = 0.05$				Gaussian Deblurring, $\sigma_y = 0.05$			
		FID ↓	LPIPS ↓	PSNR ↑	SSIM ↑	FID ↓	LPIPS ↓	PSNR ↑	SSIM ↑	FID ↓	LPIPS ↓	PSNR ↑	SSIM ↑
OT-CFM	$g^{\text{cov-A}}$	7.330	0.1904	25.70	0.8432	26.06	0.3016	26.64	0.7292	14.34	0.2968	24.46	0.6982
	$g^{\text{sim-A}}$	10.67	0.1716	25.42	0.8681	31.78	0.3717	23.88	0.6072	11.83	0.2837	24.54	0.7109
	$g^{\text{cov-G}}$	18.43	0.2390	26.96	0.8125	15.42	0.2533	27.40	0.7783	17.04	0.2873	24.96	0.7196
	IIGDM	12.56	<u>0.1717</u>	<u>27.74</u>	0.8723	9.828	0.2322	<u>27.24</u>	0.7891	12.95	0.2226	28.43	0.7952
	g^{MC}	22.75	0.5589	8.67	0.3484	22.92	0.5596	8.59	0.3468	22.82	0.5596	8.64	0.3469
CFM	$g^{\text{cov-A}}$	<u>7.678</u>	0.1920	25.95	0.8414	31.76	0.3770	22.79	0.5899	16.09	0.3052	24.21	0.6825
	$g^{\text{sim-A}}$	11.54	0.1785	28.00	0.8686	33.02	0.3749	23.75	0.6015	13.37	0.2947	24.34	0.6926
	$g^{\text{cov-G}}$	19.65	0.2377	27.03	0.8140	13.89	0.2461	27.15	<u>0.7864</u>	16.89	0.2908	24.84	0.7112
	IIGDM	15.27	0.1753	25.48	<u>0.8700</u>	<u>10.52</u>	<u>0.2435</u>	26.96	0.7755	<u>12.60</u>	<u>0.2244</u>	<u>28.27</u>	<u>0.7893</u>
	g^{MC}	26.37	0.5615	9.06	0.3495	26.75	0.5492	10.05	0.3689	26.84	0.5494	10.04	0.3684

the samples generated with g^{MC} have a density distribution that best matches the ground-truth target distribution.

4.3. Image Inverse Problems

We conduct experiments on the image inverse problems on the CelebA-HQ (256×256) dataset to reflect the guidance performances on higher dimensional generative tasks. We consider three different types of noised inverse problems: box inpainting, super-resolution, and Gaussian deblurring, and compute the metrics FID, LPIPS, PSNR, and SSIM on 3000 samples from the test set. The details of the settings and result visualizations are included in Appendix B.3.

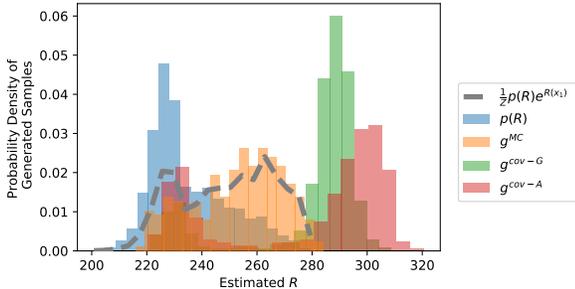


Figure 3: R distribution of generated trajectories in Locomotion. g^{MC} matches the target gray dashed line well.

The results demonstrate that IIGDM is generally better on all tasks, being the best in 8 out of 12 metrics. $g^{\text{cov-G}}$ has a similar but slightly worse performance than IIGDM, being the best or the runner-up in all 4 metrics of the super-resolution task if ranking the results of CFM and OT-CFM separately and 3 out of 4 metrics in the deblurring task. $g^{\text{sim-A}}$ that does not involve the Jacobian shows remarkable performance in inpainting and deblurring, with a lower FID score than IIGDM by 1.5 on average, and LPIPS, PSNR, and SSIM within 3% relative difference compared to the best or the runner-up method in 5 out of 6 metrics, especially when considering the efficiency of $g^{\text{sim-A}}$ that no backprop-

agation through the model is needed. $g^{\text{cov-A}}$ is the worst excluding g^{MC} , being the second worst or the worst in 11 out of 12 metrics and ranking CFM and OT-CFM separately. It should be noted that g^{MC} performs poorly here, as J in the inverse problem is highly complex, thus requiring an infeasibly large sample budget to compute accurate g^{MC} . A more detailed explanation is deferred to Appendix B.3.

5. Limitations and Future Works

The major limitation of our work lies in the assumption that $\mathcal{P} \approx 1$. When the coupling is strong, the guidance VF no longer has the correct direction. In future works, we plan to address this problem by estimating \mathcal{P} , which enables guidance for flow matching models with exact coupling.

In addition, we plan to improve the specific guidance methods. For example, g^{MC} suffers from the low sample efficiency on high-dimensional datasets. Thus, it is worthwhile to comprehensively investigate and improve the effectiveness of $g_t^{\text{MC-IS}}$, or to explore other techniques (such as the control variable method) to further lower the variance.

Besides, guided VF $v'_t(x_t)$ can be chosen to have different properties, such as the straightness of the VF. In this way, there may be add-on VFs that facilitate the sampling efficiency, which we also leave for future work.

6. Conclusion

In this work, we have proposed the first framework for general flow matching guidance, from which new MC-based guidance g^{MC} , many approximate guidance, and learned guidance g_ϕ , are derived, all of them verified by experiments. Many classical guidance methods have been covered as special cases, and we provided a theoretical derivation for each guidance method. We believe this work will facilitate the application of flow matching by opening novel design spaces of its guidance methods.

Acknowledgements

We thank Qianyi Chen, Tengfei Xu, Long Wei, Chuanrui Wang, Jiashu Pan, and Tao Zhang for the discussions and for providing feedback on our manuscript. We gratefully acknowledge the support of Westlake University Research Center for Industries of the Future and Westlake University Center for High-performance Computing. The content is solely the responsibility of the authors and does not necessarily represent the official views of the funding entities.

Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none of which we feel must be specifically highlighted here.

References

- Ajay, A., Du, Y., Gupta, A., Tenenbaum, J., Jaakkola, T., and Agrawal, P. Is conditional generative modeling all you need for decision-making? In *The Eleventh International Conference on Learning Representations*, 2023.
- Andrae, M., Landelius, T., Oskarsson, J., and Lindsten, F. Continuous ensemble weather forecasting with diffusion models. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=ePEZvQNFDW>.
- Ben-Hamu, H., Puny, O., Gat, I., Karrer, B., Singer, U., and Lipman, Y. D-flow: Differentiating through flows for controlled generation. In *Proceedings of the 41st International Conference on Machine Learning*, 2024.
- Boys, B., Girolami, M., Pidstrigach, J., Reich, S., Mosca, A., and Akyildiz, O. D. Tweedie moment projected diffusions for inverse problems. In *The Twelfth International Conference on Learning Representations*, 2023.
- Chen, R. T. Q. and Lipman, Y. Flow matching on general geometries. In *The Twelfth International Conference on Learning Representations*, 2024.
- Chung, H., Kim, J., Mccann, M. T., Klasky, M. L., and Ye, J. C. Diffusion posterior sampling for general noisy inverse problems. In *The Eleventh International Conference on Learning Representations*, 2023.
- Dhariwal, P. and Nichol, A. Diffusion models beat GANs on image synthesis. In *The Thirty-Fifth Annual Conference on Neural Information Processing Systems*, 2021.
- Dou, Z. and Song, Y. Diffusion posterior sampling for linear inverse problem solving: A filtering perspective. In *The Twelfth International Conference on Learning Representations*, 2024.
- Fan, Y., Watkins, O., Du, Y., Liu, H., Ryu, M., Boutilier, C., Abbeel, P., Ghavamzadeh, M., Lee, K., and Lee, K. DPOK: Reinforcement learning for fine-tuning text-to-image diffusion models. In *The Thirty-seventh Annual Conference on Neural Information Processing Systems*, 2023.
- Fu, J., Kumar, A., Nachum, O., Tucker, G., and Levine, S. D4rl: Datasets for deep data-driven reinforcement learning. arXiv preprint arXiv:2004.07219, 2020.
- Gat, I., Remez, T., Shaul, N., Kreuk, F., Chen, R. T. Q., Synnaeve, G., Adi, Y., and Lipman, Y. Discrete flow matching. In *The Thirty-Eighth Annual Conference on Neural Information Processing Systems*, 2024.
- Hiranaka, A., Chen, S.-F., Lai, C.-H., Kim, D., Murata, N., Shibuya, T., Liao, W.-H., Sun, S.-H., and Mitsufuji, Y. HERO: Human-feedback efficient reinforcement learning for online diffusion model finetuning. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=yMHe9SRvXk>.
- Ho, J. and Salimans, T. Classifier-free diffusion guidance. In *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*, 2022.
- Ho, J., Jain, A., and Abbeel, P. Denoising diffusion probabilistic models. In *The Thirty-Fourth Annual Conference on Neural Information Processing Systems*, 2020.
- Janner, M., Du, Y., Tenenbaum, J. B., and Levine, S. Planning with diffusion for flexible behavior synthesis. In *Proceedings of the 39th International Conference on Machine Learning*, 2022.
- Kollovich, M., Lienen, M., Lüdke, D., Schwinn, L., and Günemann, S. Flow matching with gaussian process priors for probabilistic time series forecasting. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=uxVBbSlKQ4>.
- Levine, S. Reinforcement learning and control as probabilistic inference: Tutorial and review. arXiv preprint arXiv:1805.00909, 2018.
- Lipman, Y., Chen, R. T. Q., Ben-Hamu, H., Nickel, M., and Le, M. Flow matching for generative modeling. In *The Eleventh International Conference on Learning Representations*, 2023.
- Lipman, Y., Havasi, M., Holderrieth, P., Shaul, N., Le, M., Karrer, B., Chen, R. T. Q., Lopez-Paz, D., Ben-Hamu, H., and Gat, I. Flow matching guide and code. arXiv preprint arXiv:2412.06264, 2024.

- Liu, X., Gong, C., and Liu, Q. Flow straight and fast: Learning to generate and transfer data with rectified flow. In *The Eleventh International Conference on Learning Representations*, 2023a.
- Liu, X., Wu, L., Zhang, S., Gong, C., Ping, W., and Liu, Q. FlowGrad: Controlling the output of generative ODEs with gradients. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023b.
- Lu, C., Chen, H., Chen, J., Su, H., Li, C., and Zhu, J. Contrastive energy prediction for exact energy-guided diffusion sampling in offline reinforcement learning. In *Proceedings of the 40th International Conference on Machine Learning*, 2023.
- Lu, H., Han, D., Shen, Y., and Li, D. What makes a good diffusion planner for decision making? In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=7BQkXXM8Fy>.
- Ma, N., Goldstein, M., Albergo, M. S., Boffi, N. M., Vanden-Eijnden, E., and Xie, S. SiT: Exploring Flow and Diffusion-based Generative Models with Scalable Interpolant Transformers. In *The 18th European Conference on Computer Vision*, 2024.
- Mathieu, E. and Nickel, M. Riemannian continuous normalizing flows. In *The Thirty-Fourth Annual Conference on Neural Information Processing Systems*, 2020.
- Owen, A. B. *Monte Carlo theory, methods and examples*. <https://artowen.su.domains/mc/>, 2013.
- Papamakarios, G., Pavlakou, T., and Murray, I. Masked autoregressive flow for density estimation. In *The Thirty-First Annual Conference on Neural Information Processing Systems*, 2017.
- Pokle, A., Muckley, M. J., Chen, R. T. Q., and Karrer, B. Training-free linear image inverses via flows. In *The Twelfth International Conference on Learning Representations*, 2024.
- Rozet, F., Andry, G., Lanusse, F., and Louppe, G. Learning diffusion priors from observations by expectation maximization. In *The Thirty-Eighth Annual Conference on Neural Information Processing Systems*, 2024.
- Song, J., Vahdat, A., Mardani, M., and Kautz, J. Pseudoinverse-guided diffusion models for inverse problems. In *The Eleventh International Conference on Learning Representations*, 2023a.
- Song, J., Zhang, Q., Yin, H., Mardani, M., Liu, M.-Y., Kautz, J., Chen, Y., and Vahdat, A. Loss-guided diffusion models for plug-and-play controllable generation. In *Proceedings of the 40th International Conference on Machine Learning*, 2023b.
- Song, Y. and Ermon, S. Generative modeling by estimating gradients of the data distribution. In *The Thirty-third Annual Conference on Neural Information Processing Systems*, 2019.
- Song, Y., Garg, S., Shi, J., and Ermon, S. Sliced score matching: A scalable approach to density and score estimation. In *Proceedings of The 35th Uncertainty in Artificial Intelligence Conference*, 2020.
- Song, Y., Sohl-Dickstein, J., Kingma, D. P., Kumar, A., Ermon, S., and Poole, B. Score-based generative modeling through stochastic differential equations. In *The Ninth International Conference on Learning Representations*, 2021.
- Stimper, V., Schölkopf, B., and Hernández-Lobato, J. M. Resampling base distributions of normalizing flows. In *Proceedings of the 25th International Conference on Artificial Intelligence and Statistics*, 2022.
- Tong, A., Fatras, K., Malkin, N., Huguet, G., Zhang, Y., Rector-Brooks, J., Wolf, G., and Bengio, Y. Improving and generalizing flow-based generative models with minibatch optimal transport. *Transactions on Machine Learning Research*, March 2024.
- Trippe, B. L., Yim, J., Tischer, D., Baker, D., Broderick, T., Barzilay, R., and Jaakkola, T. Diffusion probabilistic modeling of protein backbones in 3d for the motif-scaffolding problem. In *The Eleventh International Conference on Learning Representations*, 2023.
- Wang, L., Cheng, C., Liao, Y., Qu, Y., and Liu, G. Training free guided flow-matching with optimal control. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=61ss5RA1MM>.
- Ye, H., Lin, H., Han, J., Xu, M., Liu, S., Liang, Y., Ma, J., Zou, J., and Ermon, S. TFG: Unified training-free guidance for diffusion models. In *The Thirty-Eighth Annual Conference on Neural Information Processing Systems*, 2024.
- Zhang, S., Zhang, W., and Gu, Q. Energy-weighted flow matching for offline reinforcement learning. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=HA0oLUvuGI>.
- Zheng, H., Chu, W., Wang, A., Kovachki, N., Baptista, R., and Yue, Y. Ensemble kalman diffusion guidance: A derivative-free method for inverse problems. arXiv preprint arXiv:2409.20175v1, 2024.

Zheng, Q., Le, M., Shaul, N., Lipman, Y., Grover, A., and Chen, R. T. Q. Guided flows for generative modeling and decision making. arXiv preprint arXiv:2311.13443, 2023.

Zhou, G., Swaminathan, S., Raju, R. V., Guntupalli, J. S., Lehrach, W., Ortiz, J., Dedieu, A., Lázaro-Gredilla, M., and Murphy, K. Diffusion model predictive control. arXiv preprint arXiv:2410.05364, 2024.

A. Complete Theoretical Derivations and Proofs

Here, we provide the theoretical analysis that is deferred from the main text, including the following subsections: Appendix A.1 includes proof of how energy-guided sampling from $p(x)e^{-J(x)}/Z$ is equivalent to conditional sampling from $p(x|y)$. Appendix A.2 proves Theorem 3.1 by showing $g_t + v_t$ is equal to v'_t which generates the correctly guided probability path. Appendix A.4 Proof of that under uncoupled affine Gaussian paths, our general guidance g_t is equivalent to the diffusion guidance $\nabla_{x_t} \log Z_t$. Appendix A.5 is proof of Proposition 3.5. Appendix A.6 discusses other ways to obtain Z_t , including using contrastive learning and Monte Carlo estimation. In appendix A.7, we propose three other training losses ℓ^{VGM} , ℓ^{RGM} , and ℓ^{MRGM} for g_ϕ and prove that the losses will produce the correct gradient. Appendix A.8 includes a more detailed explanation of g^{MC} and a variant of it under uncoupled paths. Appendix A.9 derives g^{local} and proves its error bound. Appendix A.10 shows how to estimate \hat{x}_1 using x_t and $v_\theta(x_t, t)$ under the affine path assumption. Appendix A.11 proves that g^{local} becomes g^{cov} under affine paths. Appendix A.12 includes the proof of the Jacobian trick (Proposition A.12). Appendix A.13 includes a proof of $g_t^{sim-inv}$ for the image inverse problem, how to derive $g^{sum-inv-A}$, and how to recover IIGDM under the uncoupled affine Gaussian path assumption.

A.1. Energy Guided Sampling as Posterior Sampling

There exists a $J(x)$ such that sampling from $\frac{1}{Z}p(x)e^{-J(x)}$ is equivalent to sampling from $p(x|y)$.

Simply take $J = -\log p(y|x)$. Plug it in to get

$$\frac{1}{Z}p(x)e^{-J(x)} = \frac{p(x)e^{\log p(y|x)}}{\int p(x)e^{\log p(y|x)}dx} = \frac{p(x)e^{\log \frac{p(y|x)}{p(y)}}}{\int p(x)e^{\log \frac{p(y|x)}{p(y)}}dx} = p(x|y). \quad (12)$$

This theorem can be interpreted as, when you have a classifier $p(y|x)$ and an energy guidance algorithm, you can directly use this algorithm to perform conditional generation from $p(y|x)$ by setting the energy $J(x) = p(y|x)$.

Similar approaches have been used in probability inference, reinforcement learning (Levine, 2018), and Diffuser (Janner et al., 2022) uses this to convert return-conditioned sampling into energy-guided sampling.

A.2. General Guidance

We prove Theorem 3.1 here.

Theorem A.1. *Adding the guidance VF $g_t(x_t)$ to the original VF $v_t(x_t)$ will form VF $v'_t(x_t)$ that generates $p'_t(x_t) = \int p_t(x_t|z)p'(z)dz$, as long as $g_t(x_t)$ follows:*

$$g_t(x_t) = \int \left(\mathcal{P} \frac{e^{-J(x_1)}}{Z_t(x_t)} - 1 \right) v_{t|z}(x_t|z)p(z|x_t)dz, \quad (13)$$

$$\text{where } Z_t(x_t) = \int \mathcal{P} e^{-J(x_1)} p(z|x_t)dz, \quad (14)$$

$\mathcal{P} = \frac{\pi'(x_0|x_1)}{\pi(x_0|x_1)}$ is the reverse coupling ratio, where $\pi'(x_0|x_1)$ is the reverse data coupling for the new VF, i.e., the distribution of x_0 given x_1 sampled from the target distribution.

Proof. We can subtract $v_t(x_t)$ from $v'_t(x_t)$ to construct $g_t(x_t)$:

$$g_t(x_t) = v'_t(x_t) - \int v_{t|z}(x_t|z) \frac{p_t(x_t|z)p'(z)}{p_t(x_t)} dz. \quad (15)$$

$v'_t(x_t)$ generates $p_t(x_t) = \int p_t(x_t|z)p'(z)dz$ if

$$v'_t(x_t) = \int v_{t|z}(x_t|z) \frac{p_t(x_t|z)p'(z)}{p'(x_t)} dz, \quad (16)$$

where $p'(z) = \pi'(x_0|x_1) \frac{1}{Z} p(x_1) e^{-J(x_1)}$, which follows from conditional flow matching, i.e., a VF marginalizing a conditional VF will generate the corresponding marginal probability path (Lipman et al., 2023; Tong et al., 2024). Then, we

have a possible form of $g_t(x_t)$

$$\begin{aligned}
 g_t(x_t) &= \int v_{t|z}(x_t|z) \left(\frac{p_t(x_t|z)p'(z)}{p'_t(x_t)} - \frac{p_t(x_t|z)p(z)}{p_t(x_t)} \right) dz \\
 &= \int v_{t|z}(x_t|z) \left(\frac{p_t(x_t|z)\pi'(x_0|x_1)\frac{1}{Z}p(x_1)e^{-J(x_1)}}{p'_t(x_t)} - \frac{p_t(x_t|z)p(z)}{p_t(x_t)} \right) dz \\
 &= \int v_{t|z}(x_t|z) \left(\mathcal{P}(z) \frac{p_t(x_t|z)p(z)\frac{1}{Z}e^{-J(x_1)}}{p'_t(x_t)} - \frac{p_t(x_t|z)p(z)}{p_t(x_t)} \right) dz \\
 &= \int v_{t|z}(x_t|z) \frac{p_t(x_t|z)p(z)}{p_t(x_t)} \left(\mathcal{P}(z) \frac{1}{Z}e^{-J(x_1)} \frac{p_t(x_t)}{p'_t(x_t)} - 1 \right) dz, \tag{17}
 \end{aligned}$$

where $p_t(x_t) = \int p(x_t, z) dz$, $p'(x_t) = \int p'(x_t, z) dz$, and $\mathcal{P}(z) = \frac{\pi'(x_0|x_1)}{\pi(x_0|x_1)}$. Since

$$p(x_t) = \int p(x_t, z) dz = p(x_t) \int p(z|x_t) dz \tag{18}$$

and

$$\begin{aligned}
 p'(x_t) &= \int p'(x_t, z) dz = \int p'(x_t|z)p'(z) dz \\
 &= \int p(x_t|z)\pi'(x_0|x_1)\frac{1}{Z}p(x_1)e^{-J(x_1)} dz \\
 &= \int \mathcal{P}(z)p(x_t|z)p(z)\frac{1}{Z}e^{-J(x_1)} dz \\
 &= \int \mathcal{P}(z)p(x_t, z)\frac{1}{Z}e^{-J(x_1)} dz \\
 &= p(x_t) \int \mathcal{P}(z)p(z|x_t)\frac{1}{Z}e^{-J(x_1)} dz. \tag{19}
 \end{aligned}$$

Plugging Eq. (18) and Eq. (19) into Eq. (17), we get

$$\begin{aligned}
 g_t(x_t) &= \int v_{t|z}(x_t|z) \frac{p_t(x_t|z)p(z)}{p_t(x_t)} \left(\mathcal{P}(z) \frac{1}{Z}e^{-J(x_1)} \frac{p_t(x_t)}{p'_t(x_t)} - 1 \right) dz \\
 &= \int v_{t|z}(x_t|z) \frac{p_t(x_t|z)p(z)}{p_t(x_t)} \left(\mathcal{P}(z) \frac{1}{Z}e^{-J(x_1)} \frac{\cancel{p(x_t)} \int p(z|x_t) dz}{\cancel{p(x_t)} \int \mathcal{P}(z)p(z|x_t)\frac{1}{Z}e^{-J(x_1)} dz} - 1 \right) dz \\
 &= \int v_{t|z}(x_t|z) \frac{p_t(x_t|z)p(z)}{p_t(x_t)} \left(\mathcal{P}(z)e^{-J(x_1)} \frac{1}{\int \mathcal{P}(z)p(z|x_t)e^{-J(x_1)} dz} \underbrace{\int p(z|x_t) dz}_{=1} - 1 \right) dz \\
 &= \int v_{t|z}(x_t|z) \frac{p_t(x_t|z)p(z)}{p_t(x_t)} \left(\frac{\mathcal{P}(z)e^{-J(x_1)}}{\mathbb{E}_{z \sim p(z|x_t)}[\mathcal{P}(z)e^{-J(x_1)}]} - 1 \right) dz. \tag{20}
 \end{aligned}$$

Finally, denote $Z_t = \mathbb{E}_{z \sim p(z|x_t)}[e^{-J(x_1)}]$ to complete the proof.

Remark. The theorem states that $v'_t = g_t + v_t$ not only generates the desired terminal distribution $\frac{1}{Z}p(x_1)e^{-J(x_1)}$ at time $t = 1$, but also generates a probability path $p'_t(x_t)$ that is similar to the original one. Specifically, their "noising process" $p(x_t|x_1)$, and the conditional vector fields $v(x_t|x_1)$, and the reverse coupling $p(x_0|x_1)$ are the same. These are all hyperparameters of flow matching, as one can choose an arbitrary conditional vector field satisfying boundary conditions and the conditional vector field uniquely determines the conditional probability path; the reverse coupling, given target (dataset) distribution $p(x_1)$ or $\frac{1}{Z}p(x_1)e^{-J(x_1)}$, composes the data coupling $p(x_0, x_1) = p(z)$ for flow matching training.

It should be noted that the g_t and v'_t we construct here is only one of infinitely many (Lipman et al., 2023) possible vector fields to generate $\frac{1}{Z}p(x_1)e^{-J(x_1)}$ at $t = 1$. It remains an interesting question whether there exists better v'_t that, for example, simplifies g_t or improves the vector field by straightening the flow.

On the Guidance of Flow Matching

Flow time	0.05	0.25	0.5	0.75	0.95
Relative L2	0.0382 ± 0.0076	0.0297 ± 0.0043	0.0271 ± 0.0032	0.0312 ± 0.0033	0.0717 ± 0.0096

Table 3: The relative L2 of the difference between the VF of OT-CFM and CFM at different time steps on the sampling trajectory.

A.3. The $\mathcal{P} = 1$ Approximation

From the empirical perspective, we found that $P = 1$ is a valid choice for dependent couplings in realistic datasets. As we show in Table 3, the VF of the OT-CFM (batch size 128) in and uncoupled CFM trained on CelebA 256×256 have small relative error at all flow time steps, so their ideal guidance VFs are approximately the same, validating our approximation of $P = 1$.

Next, we give a more intuitive understanding of the approximation. Our framework allows us to choose any $\pi'(x_0|x_1)$ (and hence P) as long as the source distribution is consistent: $p(x_0) = \int \pi'(x_0|x_1) \frac{1}{Z} p(x_1) e^{-J(x_1)} dx_1$. In other words, the error induced by setting $P = 1$ can be characterized by the deviation in either the source distribution or the vector field. In the former case, the guidance VF is essentially considered exact, which implies $\pi'(x_0|x_1) = \pi(x_0|x_1)$. Therefore, the error is induced by the fact that we should have sampled from $\int \pi'(x_0|x_1) \frac{1}{Z} p(x_1) e^{-J(x_1)} dx_1$, rather than the original $p(x_0)$. In the latter case, we actually assumed the source distribution to be unchanged, i.e., we need $\pi'(x_0|x_1) = p(x_0)$ to make the source distributions compatible automatically. In this case, the error is caused by approximating $\mathcal{P} = \frac{p(x_0)}{\pi(x_0|x_1)}$ with 1.

We now discuss the practical effect of the approximation, i.e., when it is a good approximation and when it is not.

In the case of strongly dependent couplings, $P \approx 1$ still holds as long as J varies slowly. This is demonstrated by the small deviation in the error of the source distribution (assuming $\pi'(x_0|x_1) = \pi(x_0|x_1)$) as we discussed above. If J is always near its average value, the new source distribution $\int \pi(x_0|x_1) \frac{1}{Z} p(x_1) e^{-J(x_1)} dx_1$ is almost $\int \pi(x_0|x_1) p(x_1) dx_1 = p(x_0)$ that is the original source.

Nevertheless, when both the coupling is strong and J varies intensively, a more complicated treatment is required for exact guidance. For example, we can try to sample from the new source distribution $\int \pi(x_0|x_1) \frac{1}{Z} p(x_1) e^{-J(x_1)} dx_1$. Although one may argue that this may be equally difficult as sampling from the target distribution $\frac{1}{Z} p(x_1) e^{-J(x_1)}$ exactly, it may be learned more easily as the target distribution is potentially smoothed after being convolved with the "kernel" $\pi(x_0|x_1)$. We leave this more detailed treatment of \mathcal{P} to future work.

A.4. Uncoupled Affine Gaussian Guidance

Here we prove that $\nabla_{x_t} \log Z_t(x_t)$ is proportional to the guidance g_t in Eq. (1). Note that the term $\nabla_{x_t} \log Z_t(x_t)$ is widely used as guidance in the diffusion model literature (Dhariwal & Nichol, 2021; Ho & Salimans, 2022; Song et al., 2021; 2023b;a; Janner et al., 2022; Ajay et al., 2023). Therefore, we prove that our general flow matching guidance exactly covers the original diffusion guidance under the affine Gaussian path assumption, i.e., when flow matching falls back to the diffusion model. Our proof here also elucidates how the gradient ∇_{x_t} emerges from the original expression of the general guidance for flow matching in Eq. (1) where there is no apparent gradient.

We restate Eq. (1) here:

$$g_t(x_t) = \int \left(\frac{e^{-J(x_1)}}{Z_t(x_t)} - 1 \right) v_{t|z}(x_t|z) \frac{p_t(x_t|z)p(z)}{p_t(x_t)} dz \quad (21)$$

$$Z_t(x_t) = \int e^{-J(x_1)} \frac{p_t(x_t|z)p(z)}{p_t(x_t)} dz \quad (22)$$

Assuming the flow matching to be of uncoupled affine path, we have

$$x_t = \sigma_t x_0 + \alpha_t x_1, \quad (23)$$

where σ_t and α_t are schedulers satisfying boundary conditions $\sigma_0 = \alpha_1 = 1$, $\sigma_1 = \alpha_0 = 0$. Thus,

$$\begin{aligned} v_{t|z}(x_t|z) &= \dot{\sigma}_t x_0 + \dot{\alpha}_t x_1 \\ &= \underbrace{\frac{\dot{\sigma}_t}{\sigma_t}}_{a_t} x_t + \underbrace{\frac{1}{\sigma_t}(\dot{\alpha}_t \sigma_t - \dot{\sigma}_t \alpha_t)}_{b_t} x_1 \end{aligned} \quad (24)$$

where $\dot{f}_t := \frac{df}{dt}$ denotes derivative to time t , and we define $a_t := \frac{\dot{\sigma}_t}{\sigma_t}$, $b_t := \frac{1}{\sigma_t}(\dot{\alpha}_t \sigma_t - \dot{\sigma}_t \alpha_t)$.

First, we demonstrate a useful technique for the proof later. Inserting Eq. (24) into $g_t(x_t)$ in Eq. (21) and we have

$$g_t(x_t) = \int \left(\frac{e^{-J(x_1)}}{Z_t} - 1 \right) (a_t x_t + b_t x_1) \frac{p_t(x_t|z)p(z)}{p_t(x_t)} dz. \quad (25)$$

Since $Z_t = \mathbb{E}_{z \sim p(z|x_t)}[e^{-J(x_1)}]$,

$$\int \left(\frac{e^{-J(x_1)}}{Z_t} - 1 \right) a_t x_t \frac{p_t(x_t|z)p(z)}{p_t(x_t)} dz = 0. \quad (26)$$

That is to say, x_t inside the integral of Eq. (26) will integrate to zero, and we can freely remove or add it to construct desired terms.

Recall the assumption of uncoupled Gaussian path, *i.e.* $p(x_0, x_1) = p(x_0)p(x_1)$, $p(x_0) = \mathcal{N}(x_0; 0, I)$. We can utilize the important fact that the conditional probability path for affine Gaussian path flows satisfies $x_t \sim \mathcal{N}(x_t; \alpha_t x_1, \sigma_t^2 I)$, which allows us to connect the conditional score to x_1

$$\nabla_{x_t} \log p(x_t|x_1) = -\frac{x_t - \alpha_t x_1}{\sigma_t^2}. \quad (27)$$

Using Eq. (26) and Eq. (27), Eq. (25) can be further converted to

$$g_t(x_t) = \int \left(\frac{e^{-J(x_1)}}{Z_t} - 1 \right) \underbrace{\left(b_t x_1 - \frac{b_t}{\alpha_t} x_t \right)}_{\text{Eq. (26)}} \frac{p_t(x_t|x_1)p(x_1)}{p_t(x_t)} dx_1 \quad (28)$$

$$\begin{aligned} &= \frac{b_t \sigma_t^2}{\alpha_t} \int \left(\frac{e^{-J(x_1)}}{Z_t} - 1 \right) \nabla_{x_t} \log p(x_t|x_1) \frac{p_t(x_t|x_1)p(x_1)}{p_t(x_t)} dx_1 \\ &= \frac{b_t \sigma_t^2}{\alpha_t} \int \left(\frac{e^{-J(x_1)}}{Z_t} - 1 \right) \underbrace{\left(\nabla_{x_t} \log p(x_t) + \nabla_{x_t} \log p(x_1|x_t) \right)}_{\text{Bayes' rule}} \frac{p_t(x_t|x_1)p(x_1)}{p_t(x_t)} dx_1 \\ &= \frac{b_t \sigma_t^2}{\alpha_t} \int \left(\frac{e^{-J(x_1)}}{Z_t} - 1 \right) \underbrace{\left(\underbrace{\nabla_{x_t} \log p(x_t)}_{\text{Integrates to zero as in Eq. (26)}} + \nabla_{x_t} \log p(x_1|x_t) \right)}_{\text{Integrates to zero as in Eq. (26)}} \frac{p_t(x_t|x_1)p(x_1)}{p_t(x_t)} dx_1 \end{aligned} \quad (29)$$

$$= \frac{b_t \sigma_t^2}{\alpha_t} \int \left(\frac{e^{-J(x_1)}}{Z_t} - 1 \right) \nabla_{x_t} \log p(x_1|x_t) \frac{p_t(x_t|x_1)p(x_1)}{p_t(x_t)} dx_1. \quad (30)$$

Notice in Eq. (30) that

$$\begin{aligned} &\int \nabla_{x_t} \log p(x_1|x_t) \frac{p_t(x_t|x_1)p(x_1)}{p_t(x_t)} dx_1 \\ &= \int \underbrace{p(x_1|x_t) \nabla_{x_t} \log p(x_1|x_t)}_{\text{Since } f \nabla \log f = \nabla f} dx_1 \\ &= \int \nabla_{x_t} p(x_1|x_t) dx_1 \\ &= \nabla_{x_t} \int p(x_1|x_t) dx_1 \\ &= 0. \end{aligned} \quad (31)$$

Therefore

$$\begin{aligned}
 g_t(x_t) &= -\frac{b_t \sigma_t^2}{\alpha_t} \int \left(\frac{e^{-J(x_1)}}{Z_t} \mathcal{I} \right) \nabla_{x_t} \log p(x_1|x_t) \frac{p_t(x_t|x_1)p(x_1)}{p_t(x_t)} dx_1 \\
 &= \frac{b_t \sigma_t^2}{\alpha_t} \int \frac{e^{-J(x_1)}}{Z_t(x_t)} \underbrace{p_t(x_1|x_t) \nabla_{x_t} \log p(x_1|x_t)}_{\text{Using again } f \nabla \log f = \nabla f} dx_1 \\
 &= \frac{b_t \sigma_t^2}{\alpha_t} \frac{1}{Z_t(x_t)} \int e^{-J(x_1)} \nabla_{x_t} p(x_1|x_t) dx_1 \\
 &= \frac{b_t \sigma_t^2}{\alpha_t} \frac{1}{Z_t(x_t)} \underbrace{\nabla_{x_t} \int e^{-J(x_1)} p(x_1|x_t) dx_1}_{\text{Absorb } e^{-J(x_1)} \text{ for it is independent of } x_1, \text{ and exchange with integral}} \\
 &= \frac{b_t \sigma_t^2}{\alpha_t} \frac{1}{Z_t(x_t)} \nabla_{x_t} \underbrace{Z_t(x_t)}_{Z_t \text{'s definition in Eq. (1)}} \\
 &= \frac{b_t \sigma_t^2}{a_t} \nabla_{x_t} \log Z_t(x_t). \tag{32}
 \end{aligned}$$

Another possible way to derive this is to first prove the vector field in affine Gaussian path flow matching to be affine to the marginal score, and we direct interested readers to (Zheng et al., 2023).

Remark A.2. The above derivation opens the possibility of using diffusion guidance into affine Gaussian path flow matching, i.e., by multiplying a scheduler $-\frac{b_t \sigma_t^2}{\alpha_t} = \frac{\sigma_t(\dot{\alpha}_t \sigma_t - \dot{\sigma}_t \alpha_t)}{\alpha_t}$ to the diffusion classifier guidance. The most common scheduler for flow matching is $\sigma_t = 1 - t$, $\alpha_t = t$ (Lipman et al., 2023; Pople et al., 2024; Zheng et al., 2023; Tong et al., 2024; Liu et al., 2023a; Lipman et al., 2024). In this case, the guidance scheduler is $\frac{(1-t)}{t}$. It should be noted that this scheduler explodes near $t = 0$, thus being unstable. The flow matching schedule σ_t and α_t can be chosen as other ways to avoid this instability.

Remark A.3. Note that this guidance cannot be applied to coupled paths. Central to the proof is that in uncoupled affine Gaussian paths, we can convert the conditional vector field to the conditional score. If we could do this in coupled paths, we would require (1) $p_t(x|z)$ is Gaussian $\mathcal{N}(x; \mu_t, \sigma_t I)$, such that $\nabla_{x_t} \log p_t(x_t|z) \propto x_t - \mu_t$. and (2) $v_t|z = \dot{\mu}_t + \dot{\sigma}_t(x_t - \mu_t) \propto \mu_t$, such that in Eq. (28) the conditional vector field can be converted to the conditional score. Therefore, the following equation must hold for any x_t, x_0, x_1

$$\dot{\mu}_t + \dot{\sigma}_t(x_t - \mu_t) \propto \mu_t \tag{33}$$

inside the integral of Eq. (28) where $\mu_t = \xi_t x_0 + \eta_t x_1$. This equivalent to that

$$(\dot{\xi}_t - \dot{\sigma}_t \xi_t)x_0 + (\dot{\eta}_t - \dot{\sigma}_t \eta_t)x_1 + \dot{\sigma}_t x_t \propto \xi_t x_0 + \eta_t x_1 \tag{34}$$

must hold for any x_t, x_0, x_1 inside the integral of Eq. (28). According to Eq. (26), x_t terms will integrate to zero, thus

$$\eta_t(\dot{\xi}_t - \dot{\sigma}_t \xi_t) = \xi_t(\dot{\eta}_t - \dot{\sigma}_t \eta_t) \tag{35}$$

which cannot hold⁸: because of the boundary conditions $\xi_0 = \eta_1 = 1$ and $\xi_1 = \eta_0 = 0$, $\exists t \in (0, 1)$, $\frac{d \log \xi_t}{dt} \neq \frac{d \log \eta_t}{dt}$. It can be observed that the reason why this guidance does not apply to coupled paths is that x_0, x_1 , and x_t are all independent variables here, preventing us from canceling two of x_0 to avoid matching the schedulers' ratio.

A.5. Proof of Proposition 3.5

We prove proposition 3.5 here.

Proposition A.4. Any marginal variable $f(x_t, t) := \mathbb{E}_{z \sim p_t(z|x_t)}[f_t|z(x_t, z, t)]$, $z = (x_0, x_1)$ has an intractable marginal loss

$$\mathcal{L} = \mathbb{E}_{x_t \sim p(x_t)} \left[\left\| f_\theta(x_t, t) - \mathbb{E}_{z \sim p_t(z|x_t)}[f_t|z(x_t, z, t)] \right\|_2^2 \right], \tag{36}$$

whose gradient is identical to the tractable conditional loss

$$\mathcal{L}_{t|z} = \mathbb{E}_{x_t, z \sim p(x_t, z)} \left[\left\| f_\theta(x_t, t) - f_t|z(x_t, z, t) \right\|_2^2 \right]. \tag{37}$$

⁸Unless $\xi_t = 1$ which falls back to the uncoupled case.

Proof. Expand and take gradient w.r.t. Eq. (36) to get

$$\begin{aligned}
 \nabla_{\theta} \mathcal{L}_t &= \nabla_{\theta} \mathbb{E}_{x_t \sim p(x_t)} \left[\left\| f_{\theta}(x_t, t) - \mathbb{E}_{z \sim p_t(z|x_t)} [f(x_t, z, t)] \right\|_2^2 \right] \\
 &= \mathbb{E}_{x_t \sim p(x_t)} \left[\nabla_{\theta} \left\| f_{\theta}(x_t, t) - \mathbb{E}_{z \sim p_t(z|x_t)} [f(x_t, z, t)] \right\|_2^2 \right] \\
 &= \int \nabla_{\theta} p(x_t) \left\| f_{\theta}(x_t, t) - \mathbb{E}_{z \sim p_t(z|x_t)} [f(x_t, z, t)] \right\|_2^2 dx_t \\
 &= \int \nabla_{\theta} p(x_t) \left(\left\| f_{\theta}(x_t, t) \right\|^2 - 2 \langle f_{\theta}(x_t, t), \int p_t(z|x_t) dz f(x_t, z, t) \rangle \right) dx_t \\
 &= \int \nabla_{\theta} p_t(z|x_t) p(x_t) \left(\left\| f_{\theta}(x_t, t) \right\|^2 - 2 \langle f_{\theta}(x_t, t), f(x_t, z, t) \rangle \right) dx_t dz \\
 &= \int p_t(z|x_t) p(x_t) \nabla_{\theta} \left(\left\| f_{\theta}(x_t, t) \right\|^2 - 2 \langle f_{\theta}(x_t, t), f(x_t, z, t) \rangle \right) dx_t dz \\
 &= \mathbb{E}_{z, x_t \sim p_t(z|x_t) p(x_t)} \left[\nabla_{\theta} \left\| f_{\theta}(x_t, t) - f(x_t, z, t) \right\|_2^2 \right] \\
 &= \underbrace{\nabla_{\theta} \mathbb{E}_{z, x_t \sim p_t(z|x_t) p(x_t)} \left[\left\| f_{\theta}(x_t, t) - f(x_t, z, t) \right\|_2^2 \right]}_{\mathcal{L}_{1|z}}. \tag{38}
 \end{aligned}$$

Thus, the gradient of the *marginal loss* \mathcal{L}_t is identical to the gradient of the *conditional loss* $\mathcal{L}_{t|z}$.

A.6. Other Ways to Obtain Z_{ϕ_Z}

Lu et al. (2023) proposed to use contrastive learning to train Z_{ϕ_Z} . The proof already applies to any uncoupled path, and we show that Z_{ϕ_Z} does not depend on the coupling

$$Z_t = \mathbb{E}_{x_1 \sim p(x_0, x_1|x_t)} [e^{-J(x_1)}] = \int e^{-J(x_1)} p(x_0|x_1, x_t) p(x_1|x_t) dx_0 dx_1 \tag{39}$$

$$= \int e^{-J(x_1)} p(x_1|x_t) dx_1 = \mathbb{E}_{x_1 \sim p(x_1|x_t)} [e^{-J(x_1)}]. \tag{40}$$

That is, instead of actually sampling from $p(x_0, x_1|x_t)$, sampling from $p(x_1|x_t)$ will result in the same Z_t . In the case of the coupled path, the marginalized distribution is identical to the uncoupled path case. Therefore, the contrastive learning method can be readily applied to train Z_{ϕ_Z} .

Besides training-based Z_{ϕ_Z} , we can also use Monte Carlo estimation to obtain Z_t . Notice that by using importance sampling, we have

$$Z_t = \mathbb{E}_{x_1 \sim p(x_0, x_1|x_t)} [e^{-J(x_1)}] = \mathbb{E}_{x_1 \sim p(x_0, x_1)} \left[\frac{p(x_t|x_0, x_1)}{p(x_t)} e^{-J(x_1)} \right]. \tag{41}$$

As long as $p(x_t|x_0, x_1)$ is known (which is often the case (Lipman et al., 2023; Tong et al., 2024)), we can estimate Z_t by sampling N pairs of x_0^i, x_1^i from $p(x_0, x_1)$ and estimate

$$\tilde{Z}_t = \sum_i^N \left(\frac{p(x_0^i, x_1^i|x_t)}{\sum_j^N p(x_0^j, x_1^j|x_t)} e^{-J(x_1^i)} \right). \tag{42}$$

A similar technique is used in Section 3.2.

A.7. Guidance Matching Losses

Here, we prove that the loss in guidance matching is correct and show there are three other equivalent training losses $\ell^{\text{VGM}}, \ell^{\text{RGM}}, \ell^{\text{MRGM}}$. The expressions of different losses are summarized below, and their proof follows.

VF-added Guidance Matching (VGM) Loss. By utilizing the learned VF $v_{\theta}(x_t, t)$ into Eq. (11), we have

$$\ell_{\phi}^{\text{VGM}} = \left\| g_{\phi}(x_t, t) + v_{\theta}(x_t, t) - \frac{e^{-J(x_1)}}{Z_{\phi_Z, sg}(x_t, t)} v_{t|z}(x_t|z) \right\|_2^2. \tag{43}$$

Reweight Guidance Matching (RGM) Loss. ℓ_ϕ^{VGM} can be further shown equivalent to

$$\ell_\phi^{\text{RGM}} = \frac{e^{-J(x_1)}}{Z_{\phi_Z, sg}(x_t, t)} \left\| g_\phi(x_t, t) + v_\theta(x_t, t) - v_{t|z}(x_t|z) \right\|_2^2. \quad (44)$$

This training loss steers g_ϕ to where $e^{-J(x_1)}$ is larger by assigning a large loss to steer g_t towards high $e^{-J(x_1)}$ regions.

Marginalized Reweight Guidance Matching (MRGM) Loss. The above loss can be re-assigned a weight, which will result in the same optimal $g_{\phi_2}(x_t, t)$ as in Eq. (11). Specifically, by changing $Z_{\phi_Z, sg}(x_t, t)$ to its expectation under $p_t(x_t)$, we have the following equivalent loss

$$\ell_\phi^{\text{MRGM}} = \frac{e^{-J(x_1)}}{Z} \left\| g_{\phi_2}(x_t, t) + v_\theta(x_t, t) - v_{t|z}(x_t|z) \right\|_2^2, \quad (45)$$

where $Z = \mathbb{E}_{x_1 \sim p(x_1)}[e^{-J(x_1)}]$. ℓ_ϕ^{MRGM} is identical to a newly proposed fine-tuning loss in Zhang et al. (2025). It can also be derived via importance sampling in Eq. 10 and similar reweighting-based fine-tuning losses have been studied in the literature of diffusion models (Fan et al., 2023).

(1) Guidance Matching Loss ℓ_ϕ^{GM} By using proposition 3.5, the following conditional loss

$$\mathcal{L}_\phi^{\text{GM}} = \mathbb{E}_{t \sim \mathcal{U}(0,1), z \sim p(z), x_t \sim p(x_t|z)} \left[\underbrace{\left\| g_\phi(x_t, t) - \left(\frac{e^{-J(x_1)}}{Z_{\phi_Z, sg}(x_t, t)} - 1 \right) v_{t|z}(x_t|z) \right\|_2^2}_{=\ell_\phi^{\text{GM}}} \right] \quad (46)$$

has a gradient that is equivalent to the marginal loss

$$\mathbb{E}_{t \sim \mathcal{U}(0,1), z \sim p(z), x_t \sim p(x_t|z)} \left[\left\| g_\phi(x_t, t) - \underbrace{\mathbb{E}_{z \sim p(z|x_t)} \left[\left(\frac{e^{-J(x_1)}}{Z_{\phi_Z, sg}(x_t, t)} - 1 \right) v_{t|z}(x_t|z) \right]}_{=g_t(x_t)} \right\|_2^2 \right]. \quad (47)$$

Therefore, using the loss \mathcal{L}_{ϕ_Z} we can train g_ϕ to matching g_t . Recall that \mathcal{L} in Eq. (11) is identical to \mathcal{L}_{ϕ_Z} , and we proved the validity of the guidance matching training.

(2) VF-added Guidance Matching Loss ℓ_ϕ^{VGM} . By replacing the learned VF $v_\theta(x_t, t)$ into Eq. (46), we show that

$$\mathcal{L}_\phi^{\text{VGM}} = \mathbb{E}_{t \sim \mathcal{U}(0,1), z \sim p(z), x_t \sim p(x_t|z)} \left[\left\| g_\phi(x_t, t) + v_\theta(x_t, t) - \frac{e^{-J(x_1)}}{Z_{\phi_Z, sg}(x_t, t)} v_{t|z}(x_t|z) \right\|_2^2 \right], \quad (48)$$

has a gradient equal to that of \mathcal{L}_ϕ in Eq. (46).

Expand Eq. (46) to get

$$\begin{aligned} \mathcal{L}_\phi^{\text{GM}} &= \mathbb{E}_{t \sim \mathcal{U}(0,1), z \sim p(z), x_t \sim p(x_t|z)} \left[\left\| g_\phi(x_t, t) - \left(\frac{e^{-J(x_1)}}{Z_{\phi_Z, sg}(x_t, t)} - 1 \right) v_{t|z}(x_t|z) \right\|_2^2 \right] \\ &= \mathbb{E}_{t \sim \mathcal{U}(0,1), z \sim p(z), x_t \sim p(x_t|z)} \left[\underbrace{\left\| g_\phi(x_t, t) \right\|_2^2}_{\text{dependent on } \phi} + \left\| \frac{e^{-J(x_1)}}{Z_{\phi_Z, sg}(x_t, t)} v_{t|z}(x_t|z) \right\|_2^2 + \left\| v_{t|z}(x_t|z) \right\|_2^2 \right. \\ &\quad \left. - 2 \underbrace{\left\langle g_\phi(x_t, t), \frac{e^{-J(x_1)}}{Z_{\phi_Z, sg}(x_t, t)} v_{t|z}(x_t|z) \right\rangle}_{\text{dependent on } \phi} - 2 \underbrace{\left\langle \frac{e^{-J(x_1)}}{Z_{\phi_Z, sg}(x_t, t)} v_{t|z}(x_t|z), v_{t|z}(x_t|z) \right\rangle}_{\text{dependent on } \phi} - 2 \underbrace{\left\langle v_{t|z}(x_t|z), g_\phi(x_t, t) \right\rangle}_{\text{dependent on } \phi} \right]. \quad (49) \end{aligned}$$

After taking gradient w.r.t. ϕ , only the terms

$$\nabla_{\phi} \mathbb{E}_{t \sim \mathcal{U}(0,1), z \sim p(z), x_t \sim p(x_t|z)} \left[\|g_{\phi}(x_t, t)\|_2^2 - 2 \langle g_{\phi}(x_t, t), \frac{e^{-J(x_1)}}{Z_{\phi_Z, sg}(x_t, t)} v_{t|z}(x_t|z) \rangle - 2 \langle v_{t|z}(x_t|z), g_{\phi}(x_t, t) \rangle \right] \quad (50)$$

survive. Therefore, by assuming a perfectly learned $v_{\theta}(x_t, t)$, *i.e.*,

$$v_{\theta}(x_t, t) = \mathbb{E}_{z \sim p(z|x_t)} [v_{t|z}(x_t|z)], \quad (51)$$

we have

$$\begin{aligned} & \mathbb{E}_{t \sim \mathcal{U}(0,1), z \sim p(z), x_t \sim p(x_t|z)} [\langle v_{t|z}(x_t|z), g_{\phi}(x_t, t) \rangle] \\ &= \mathbb{E}_{t \sim \mathcal{U}(0,1), z \sim p(z|x_t), x_t \sim p(x_t)} [\langle v_{t|z}(x_t|z), g_{\phi}(x_t, t) \rangle] \\ &= \mathbb{E}_{\tilde{z} \sim p(\tilde{z}|x_t)} [\mathbb{E}_{t \sim \mathcal{U}(0,1), z \sim p(z|x_t), x_t \sim p(x_t)} [\langle v_{t|z}(x_t|z), g_{\phi}(x_t, t) \rangle]] \\ &= \mathbb{E}_{\tilde{z} \sim p(\tilde{z}|x_t)} [\mathbb{E}_{t \sim \mathcal{U}(0,1), x_t \sim p(x_t)} [\langle \mathbb{E}_{z \sim p(z|x_t)} [v_{t|z}(x_t|z)], g_{\phi}(x_t, t) \rangle]] \\ &= \mathbb{E}_{z \sim p(z|x_t)} [\mathbb{E}_{t \sim \mathcal{U}(0,1), x_t \sim p(x_t)} [\langle \mathbb{E}_{\tilde{z} \sim p(\tilde{z}|x_t)} [v_{t|z}(x_t|z)], g_{\phi}(x_t, t) \rangle]] \\ &= \mathbb{E}_{t \sim \mathcal{U}(0,1), z \sim p(z|x_t), x_t \sim p(x_t)} [\langle v_{\theta}(x_t, t), g_{\phi}(x_t, t) \rangle], \end{aligned} \quad (52)$$

so by adding back terms that the gradient is agnostic to, we can see that the new loss $\mathcal{L}_{\phi}^{(1)}$ in Eq. (48) is equivalent to \mathcal{L}_{ϕ} in Eq. (46)

$$\begin{aligned} \nabla_{\phi} \mathcal{L}_{\phi}^{\text{GM}} &= \nabla_{\phi} \mathbb{E}_{t \sim \mathcal{U}(0,1), z \sim p(z), x_t \sim p(x_t|z)} \left[\left\| g_{\phi}(x_t, t) - \left(\frac{e^{-J(x_1)}}{Z_{\phi_Z, sg}(x_t, t)} - 1 \right) v_{t|z}(x_t|z) \right\|_2^2 \right] \\ &= \nabla_{\phi} \mathbb{E}_{t \sim \mathcal{U}(0,1), z \sim p(z), x_t \sim p(x_t|z)} \left[\|g_{\phi}(x_t, t)\|_2^2 + \underbrace{\left\| \frac{e^{-J(x_1)}}{Z_{\phi_Z, sg}(x_t, t)} v_{t|z}(x_t|z) \right\|_2^2}_{\text{Vanishes after } \nabla_{\phi}} + \underbrace{\|v_{\theta}(x_t, t)\|_2^2}_{\text{Vanishes after } \nabla_{\phi}} \right] \\ &\quad - 2 \langle g_{\phi}(x_t, t), \frac{e^{-J(x_1)}}{Z_{\phi_Z, sg}(x_t, t)} v_{t|z}(x_t|z) \rangle - 2 \langle \frac{e^{-J(x_1)}}{Z_{\phi_Z, sg}(x_t, t)} v_{t|z}(x_t|z), v_{t|z}(x_t|z) \rangle \underbrace{- 2 \langle v_{\theta}(x_t, t), g_{\phi}(x_t, t) \rangle}_{\text{Changed } v_{t|z} \text{ to } v_{\theta} \text{ using Eq. (52)}} \quad (53) \\ &= \nabla_{\phi} \mathbb{E}_{t \sim \mathcal{U}(0,1), z \sim p(z), x_t \sim p(x_t|z)} \left[\underbrace{\left\| g_{\phi}(x_t, t) + v_{\theta}(x_t, t) - \frac{e^{-J(x_1)}}{Z_{\phi_Z, sg}(x_t, t)} v_{t|z}(x_t|z) \right\|_2^2}_{:= \ell^{\text{VGM}}} \right] \\ &= \nabla_{\phi} \mathcal{L}_{\phi}^{\text{VGM}}. \end{aligned} \quad (54)$$

(3) Reweighted Guidance Matching Loss ℓ^{RGM} . Eq. Replacing ℓ^{GM} in (46) with ℓ^{VGM} :

$$\frac{e^{-J(x_1)}}{Z_{\phi_Z, sg}(x_t, t)} \|g_{\phi}(x_t, t) + v_{\theta}(x_t, t) - v_{t|z}(x_t|z)\|_2^2, \quad (55)$$

and the loss $\mathcal{L}_{\phi}^{\text{GM}}$ becomes $\mathcal{L}_{\phi}^{\text{VGM}}$, which are shown equivalent in the following.

Starting from Eq. (49), we can extract $\frac{e^{-J(x_1)}}{Z_{\phi_Z}}$ from the three terms depended on ϕ , and thus showing the resulting loss is indeed ℓ^{RGM} . Notice that because $Z_{\phi_Z} = \mathbb{E}_{z \sim p(z|x_t)} [e^{-J(z)}]$,

$$\mathbb{E}_{z \sim p(z|x_t)} \left[\frac{e^{-J(z)}}{Z_{\phi_Z}(x_t)} f(x_t, t) \right] = f(x_t, t). \quad (56)$$

Thus, we have:

$$\begin{aligned}
 \mathcal{L}_\phi^{\text{GM}} &= \mathbb{E}_{t \sim \mathcal{U}(0,1), z \sim p(z), x_t \sim p(x_t|z)} \left[\underbrace{\|g_\phi(x_t, t)\|_2^2}_{\text{dependent on } \phi} + \left\| \frac{e^{-J(x_1)}}{Z_{\phi_Z, sg}(x_t, t)} v_{t|z}(x_t|z) \right\|_2^2 + \|v_{t|z}(x_t|z)\|_2^2 \right. \\
 &\quad \left. - 2 \left\langle g_\phi(x_t, t), \frac{e^{-J(x_1)}}{Z_{\phi_Z, sg}(x_t, t)} v_{t|z}(x_t|z) \right\rangle - 2 \left\langle \frac{e^{-J(x_1)}}{Z_{\phi_Z, sg}(x_t, t)} v_{t|z}(x_t|z), v_{t|z}(x_t|z) \right\rangle - 2 \underbrace{\langle v_{t|z}(x_t|z), g_\phi(x_t, t) \rangle}_{\text{dependent on } \phi} \right] \\
 &= \mathbb{E}_{t \sim \mathcal{U}(0,1), z \sim p(z), x_t \sim p(x_t|z)} \left[\left\| \frac{e^{-J(x_1)}}{Z_{\phi_Z}(x_t)} g_\phi(x_t, t) \right\|_2^2 + \left\| \frac{e^{-J(x_1)}}{Z_{\phi_Z, sg}(x_t, t)} v_{t|z}(x_t|z) \right\|_2^2 + \|v_{t|z}(x_t|z)\|_2^2 \right. \\
 &\quad \left. - 2 \left\langle g_\phi(x_t, t), \frac{e^{-J(x_1)}}{Z_{\phi_Z, sg}(x_t, t)} v_{t|z}(x_t|z) \right\rangle - 2 \left\langle \frac{e^{-J(x_1)}}{Z_{\phi_Z, sg}(x_t, t)} v_{t|z}(x_t|z), v_{t|z}(x_t|z) \right\rangle - 2 \frac{e^{-J(x_1)}}{Z_{\phi_Z}(x_t)} \langle v_{t|z}(x_t|z), g_\phi(x_t, t) \rangle \right] \\
 &= \mathbb{E}_{t \sim \mathcal{U}(0,1), z \sim p(z), x_t \sim p(x_t|z)} \left[\underbrace{\frac{e^{-J(x_1)}}{Z_{\phi_Z, sg}(x_t, t)} \|g_\phi(x_t, t) + v_\theta(x_t, t) - v_{t|z}(x_t|z)\|_2^2}_{=\ell^{\text{RGM}}} \right] \tag{57} \\
 &= \mathcal{L}_\phi^{\text{RGM}}. \tag{58}
 \end{aligned}$$

Where we used the conclusion of Eq. (52), and inserted terms that vanish after ∇_ϕ to make ℓ^{RGM} .

(3) Marginalized Reweighted Guidance Matching Loss ℓ^{MRGM} . The above loss Eq. (55) can be re-assigned a weight, which will result in the same optimal $g_\phi(x_t, t)$ as in Eq. (11). Specifically, by changing $\frac{1}{Z_{\phi_Z, sg}(x_t, t)}$ to $\frac{1}{\mathbb{E}_{x_t \sim p(x_t)} [Z_{\phi_Z, sg}(x_t, t)]} = \frac{1}{\int p(x_t) Z_{\phi_Z, sg}(x_t, t) dx_t}$, we have

$$\frac{e^{-J(x_1)}}{\int e^{-J(x_1)} p(z) dz} \|g_\phi(x_t, t) + v_\theta(x_t, t) - v_{t|z}(x_t|z)\|_2^2. \tag{59}$$

We only need to prove that

$$\int p(x_t) Z_t(x_t) dx_t = \int e^{-J(x_1)} p(z) dz. \tag{60}$$

Recall that

$$Z_t(x_t) = \int p(z|x_t) e^{-J(x_1)} dz, \tag{61}$$

so

$$\int p(x_t) Z_t(x_t) dx_t = \int p(x_t) \int p(z|x_t) e^{-J(x_1)} dx_t dz \tag{62}$$

$$= \int p(z) e^{-J(x_1)} dz = Z. \tag{63}$$

Eq. (45) can also be derived by applying importance sampling $\mathbb{E}_{z \sim \frac{1}{Z} p(z) e^{-J(x_1)}} [\|\cdot\|_2^2] = \mathbb{E}_{z \sim p(z)} \left[\frac{e^{-J(x_1)}}{Z} \|\cdot\|_2^2 \right]$ to the flow matching objective of the new VF for the new target distribution $p'(x_t) = \frac{1}{Z} p(x_t) e^{-J(x_t)}$.

Discussions The losses have the same expected gradient, but their performance may differ. Among the four losses, ℓ_ϕ^{MRGM} is the only one that does not require the auxiliary model Z_{ϕ_Z} . However, ℓ_ϕ^{RGM} assigns loss weight dependent on x_t . The weight is emphasized when the expectation of $e^{-J(x_1)}$ under $p(x_1|x_t)$ is small. Compared to these two losses, $\ell_\phi^{\text{GM}}, \ell_\phi^{\text{VGM}}$ do not reweight the loss. The variance of ℓ_ϕ^{GM} will be smaller if J is smooth, while ℓ_ϕ^{VGM} is better when v_t is more complex.

A.8. Algorithm Details and Variants of g^{MC}

The pseudocode for computing g^{MC} is as follows.

Algorithm 1 Monte Carlo estimation of the guidance $g_t(x_t)$

Require: Current t, x_t , known $p_t(x_t|z)$.

- 1: Sample $z_i \sim p(z)$, where $i = 1, 2, \dots, N$ // Recall $z_i = (x_1^i, x_0^i)$
 - 2: $\tilde{p}_t(x_t) \leftarrow \frac{1}{N} \sum_i p_t(x_t|z_i)$
 - 3: $\tilde{Z}_t(x_t) \leftarrow \frac{1}{N} \sum_i e^{-J(x_1^i)} \frac{p_t(x_t|z_i)}{\tilde{p}_t(x_t)}$
 - 4: $g_t^{\text{MC}}(x_t) \leftarrow \frac{1}{N} \sum_i \left(\frac{e^{-J(x_1^i)}}{\tilde{Z}_t(x_t)} - 1 \right) v_{t|z}(x_t|z_i) \frac{p_t(x_t|z_i)}{\tilde{p}_t(x_t)}$
 - 5: **return** $g_t^{\text{MC}}(x_t)$
-

Independent Couplings. Although we introduced the flow matching using the condition $z = (x_0, x_1)$, it can also be chosen as x_0 or x_1 (Lipman et al., 2024). When we $z := x_1$, Algorithm 1 can be readily adopted for the x_0 condition. This way, the Monte Carlo estimation reduces the integration region dimensionality to half of the original one, thus becoming more efficient.

In the case where $z = (x_0, x_1)$ and the data coupling is independent $\pi(x_0|x_1) = p(x_0)$, we show here that the MC estimation can be simplified to the x_1 -conditioned case that is more efficient:

$$g_t^{\text{MC-}x_1}(x_t) \triangleq \mathbb{E}_{x_1 \sim p(x_1)} \left[\left(\frac{e^{-J(x_1)}}{Z_t} - 1 \right) v_{t|x_1}(x_t|x_1) \frac{p_t(x_t|x_1)}{p_t(x_t)} \right], \quad (64)$$

$$Z_t^{\text{MC-}x_1}(x_t) = \mathbb{E}_{x_1 \sim p(x_1)} \left[e^{-J(x_1)} \frac{p_t(x_t|x_1)}{p_t(x_t)} \right]. \quad (65)$$

Obviously, as $Z_t = \int p(x_0, x_1|x_t) e^{-J(x_1)} dx_0 dx_1 = \int p(x_0|x_1, x_t) p(x_1|x_t) e^{-J(x_1)} dx_0 dx_1$, integrating out x_0 gives $Z_t = Z_t^{\text{MC-}x_1}$. Therefore, to prove the above simplification, we only need to prove that:

$$\mathbb{E}_{x_0, x_1 \sim p(x_0, x_1|x_t)} \left[\left(\frac{e^{-J(x_1)}}{Z_t} - 1 \right) v(x_t|x_0, x_1) \right] = \mathbb{E}_{x_1 \sim p(x_1|x_t)} \left[\left(\frac{e^{-J(x_1)}}{Z_t} - 1 \right) v(x_t|x_1) \right]. \quad (66)$$

The proof is simply integrating out x_0 :

$$\begin{aligned} & \int p(x_0, x_1|x_t) \left(\frac{e^{-J(x_1)}}{Z_t} - 1 \right) v(x_t|x_0, x_1) dx_0 dx_1 \\ &= \int \underbrace{\int p(x_0|x_1, x_t) v(x_t|x_0, x_1) dx_0}_{:=v(x_t|x_1)} \left(\frac{e^{-J(x_1)}}{Z_t} - 1 \right) p(x_1|x_t) dx_1 \\ &= \int v(x_t|x_1) \left(\frac{e^{-J(x_1)}}{Z_t} - 1 \right) p(x_1|x_t) dx_1. \end{aligned} \quad (67)$$

It should be noted that $v(x_t|x_1)$ is defined to be generally different from $v(x_t|x_0, x_1)$, and to do MC estimation via importance sampling, we need the forward probability path $p(x_t|x_1)$ to have a known density. The variance-reducing variant of g^{MC} is summarized in Algorithm 2.

Algorithm 2 Monte Carlo estimation of the guidance $g_t(x_t)$

Require: Current t, x_t , known $p_t(x_t|x_1)$.

- 1: Sample $x_1^i \sim p(x_1)$, where $i = 1, 2, \dots, N$
 - 2: $\tilde{p}_t(x_t) \leftarrow \frac{1}{N} \sum_i p_t(x_t|x_1^i)$
 - 3: $\tilde{Z}_t(x_t) \leftarrow \frac{1}{N} \sum_i e^{-J(x_1^i)} \frac{p_t(x_t|x_1^i)}{\tilde{p}_t(x_t)}$
 - 4: $\tilde{g}_t(x_t) \leftarrow \frac{1}{N} \sum_i \left(\frac{e^{-J(x_1^i)}}{\tilde{Z}_t(x_t)} - 1 \right) v_{t|z}(x_t|x_1^i) \frac{p_t(x_t|x_1^i)}{\tilde{p}_t(x_t)}$
 - 5: **return** $\tilde{g}_t(x_t)$
-

A.9. Localized Approximation

To get g^{local} , we presume $p(z|x_t)$ is localized, and we can use a point estimation to approximate Z_t :

$$Z_t(x_t) = \int p(z|x_t) e^{-J(x_1)} dz \approx e^{-J(\hat{x}_1)} \quad (68)$$

where $\hat{x}_1 := \mathbb{E}_{x_0, x_1 \sim p(z|x_t)}[x_1]$, and then expanding g_t to the first order

$$\begin{aligned} g_t(x_t) &\approx g_t(x_t)^{\text{local}} = \mathbb{E}_{z \sim p(z|x_t)} \left[\left(\frac{e^{-J(x_1)}}{e^{-J(\hat{x}_1)}} - 1 \right) v_{t|z}(x_t|z) \right] \\ &\approx \mathbb{E}_{z \sim p(z|x_t)} \left[\left(\frac{e^{-J(\hat{x}_1)} (1 - \nabla_{\hat{x}_1} J(\hat{x}_1)(x_1 - \hat{x}_1))}{e^{-J(\hat{x}_1)}} - 1 \right) v_{t|z}(x_t|z) \right] \\ &= -\mathbb{E}_{x_1 \sim p(x_1|x_t)} [(x_1 - \hat{x}_1) v_{t|z}(x_t|z)] \nabla_{\hat{x}_1} J(\hat{x}_1). \end{aligned} \quad (69)$$

To quantify the approximation error, we have

$$\begin{aligned} \|\delta g\|^2 &:= \|g_t - g^{\text{local}}\|_2^2 \\ &= \left\| \mathbb{E}_{z \sim p(z|x_t)} \left[\left(\frac{e^{-J(x_1)}}{Z_t(x_t)} - 1 \right) v_{t|z}(x_t|z) \right] \right. \\ &\quad \left. - \mathbb{E}_{z \sim p(z|x_t)} \left[\left(\frac{e^{-J(\hat{x}_1)} (1 - \nabla_{\hat{x}_1} J(\hat{x}_1)(x_1 - \hat{x}_1))}{e^{-J(\hat{x}_1)}} - 1 \right) v_{t|z}(x_t|z) \right] \right\|_2^2 \\ &= \left\| \mathbb{E}_{z \sim p(z|x_t)} \left[\left(\frac{e^{-J(x_1)}}{Z_t(x_t)} - \frac{e^{-J(\hat{x}_1)} (1 - \nabla_{\hat{x}_1} J(\hat{x}_1)(x_1 - \hat{x}_1))}{e^{-J(\hat{x}_1)}} \right) v_{t|z}(x_t|z) \right] \right\|_2^2 \end{aligned} \quad (70)$$

where

$$Z_t(x_t) = \mathbb{E}_{z \sim p(z|x_t)} [e^{-J(x_1)}]. \quad (71)$$

We start by computing the error bound of approximating Z_t with $e^{-J(\hat{x}_1)}$. Using Taylor expansion and the Taylor Remainder Theorem⁹,

$$\begin{aligned} \left\| Z_t(x_t) - e^{-J(\hat{x}_1)} \right\|_2^2 &= \left\| \mathbb{E}_{z \sim p(z|x_t)} \left[\sum_{k=2} \frac{1}{k!} D_{x_1}^k e^{-J(x)} \Big|_{x_1=\hat{x}_1} (x_1 - \hat{x}_1)^k \right] \right\|_2^2 \\ &\leq \mathbb{E}_{z \sim p(z|x_t)} \left[\left\| \frac{1}{2} (x_1 - \hat{x}_1)^T \underbrace{\nabla_{x_1} \nabla_{x_1} e^{-J(x)} \Big|_{x_1=\hat{x}_1 + t(x_1 - \hat{x}_1)}}_{:=h_t^{(J)}} (x_1 - \hat{x}_1) \right\|_2^2 \right], \end{aligned} \quad (72)$$

where $t \in [0, 1]$.

If we set the L2 norm of the covariance matrix $\mathbb{E}_{z \sim p(z|x_t)} [(x_1 - \hat{x}_1)(x_1 - \hat{x}_1)^T]$ as σ_1 , and the eigenvalue with the largest absolute value of $\max_{t, x_1} |h_t^{(J)}|$ to be λ_h , we can show that

$$\begin{aligned} \left\| Z_t(x_t) - e^{-J(\hat{x}_1)} \right\|_2 &\leq \mathbb{E}_{z \sim p(z|x_t)} [(x_1 - \hat{x}_1)^T h_t^{(J)} (x_1 - \hat{x}_1)] \\ &\leq \mathbb{E}_{z \sim p(z|x_t)} [(x_1 - \hat{x}_1)^T \lambda_h (x_1 - \hat{x}_1)] \\ &\leq \lambda_h \mathbb{E}_{z \sim p(z|x_t)} [(x_1 - \hat{x}_1)^T (x_1 - \hat{x}_1)] \\ &\leq \lambda_h \text{tr} \Sigma_{11} \\ &\leq \lambda_h \sigma_1 d, \end{aligned} \quad (73)$$

⁹The notations here neglect the order of vector/matrix products, but this does not matter as all of them will be scaled using the triangle inequality.

where Σ_{11} is the covariance matrix of $p(x_1|x_t)$, d is the dimensionality of $x \in \mathbb{R}^d$. The last inequality follows from $\text{tr}A = \sum_i^n \lambda_i \leq n \max_i \lambda_i$, and the L2 norm of a matrix is its largest singular value, *i.e.*, for the covariance matrix, that is the largest eigenvalue.

Then,

$$\begin{aligned} \delta g &= \mathbb{E}_{z \sim p(z|x_t)} \left[\left(\frac{e^{-J(x_1)}}{Z_t(x_t)} - \frac{e^{-J(\hat{x}_1)}(1 - \nabla_{\hat{x}_1} J(\hat{x}_1)(x_1 - \hat{x}_1))}{e^{-J(\hat{x}_1)}} \right) v_{t|z} \right] \\ &= \mathbb{E}_{z \sim p(z|x_t)} \left[\left(\underbrace{\frac{e^{-J(x_1)}}{Z_t(x_t)} - \frac{e^{-J(\hat{x}_1)}}{e^{-J(\hat{x}_1)}}}_{\text{Using the error bound of } Z_t} + \frac{e^{-J(x_1)} - e^{-J(\hat{x}_1)}(1 - \nabla_{\hat{x}_1} J(\hat{x}_1)(x_1 - \hat{x}_1))}{e^{-J(\hat{x}_1)}} \right) v_{t|z} \right]. \end{aligned} \quad (74)$$

Therefore,

$$\begin{aligned} \|\delta g\|_2^2 &\leq \left\| -\frac{\lambda_h \sigma_1 d}{Z_t(x_t) e^{-J(\hat{x}_1)}} \mathbb{E}_{z \sim p(z|x_t)} [e^{-J(x_1)} v_{t|z}] \right\|_2^2 \\ &+ \left\| \mathbb{E}_{z \sim p(z|x_t)} \left[\frac{e^{-J(x_1)} - e^{-J(\hat{x}_1)}(1 - \nabla_{\hat{x}_1} J(\hat{x}_1)(x_1 - \hat{x}_1))}{e^{-J(\hat{x}_1)}} v_{t|z} \right] \right\|_2^2 \\ &= \left\| -\frac{\lambda_h \sigma_1 d}{Z_t(x_t) e^{-J(\hat{x}_1)}} \mathbb{E}_{z \sim p(z|x_t)} [e^{-J(x_1)} v_{t|z}] \right\|_2^2 \\ &+ \left\| \mathbb{E}_{z \sim p(z|x_t)} \left[\frac{e^{-J(\hat{x}_1)}(1 - \nabla_{\hat{x}_1} J(\hat{x}_1)(x_1 - \hat{x}_1)) + R_2 - e^{-J(\hat{x}_1)}(1 - \nabla_{\hat{x}_1} J(\hat{x}_1)(x_1 - \hat{x}_1))}{e^{-J(\hat{x}_1)}} v_{t|z} \right] \right\|_2^2. \end{aligned} \quad (75)$$

By using the Taylor Remainder Theorem again, we have

$$R_2 = \frac{1}{2} (x_1 - \hat{x}_1)^T \underbrace{\nabla_{\xi} \nabla_{\xi} e^{-J(\xi)} \Big|_{\xi = \hat{x}_1 + t(x_1 - \hat{x}_1)}}_{=h_t^{(J)}} (x_1 - \hat{x}_1). \quad (76)$$

Thus,

$$\begin{aligned} \|\delta g\|_2^2 &\leq \left\| \lambda_h \sigma_1 d \frac{\mathbb{E}[e^{-J(x_1)} v(x_t|z)]}{\mathbb{E}[e^{-J(x_1)}] e^{-J(\hat{x}_1)}} \right\|_2^2 + \left\| \mathbb{E}_{z \sim p(z|x_t)} \left[\frac{1}{2e^{-J(\hat{x}_1)}} (x_1 - \hat{x}_1)^T h_t^{(J)} (x_1 - \hat{x}_1) v_{t|z} \right] \right\|_2^2 \\ &\leq \left\| \lambda_h \sigma_1 d \frac{\mathbb{E}[e^{-J(x_1)} v(x_t|z)]}{\mathbb{E}[e^{-J(x_1)}] e^{-J(\hat{x}_1)}} \right\|_2^2 \\ &+ \left[\mathbb{E}_{z \sim p(z|x_t)} \left\| \frac{1}{2e^{-J(\hat{x}_1)}} (x_1 - \hat{x}_1)^T h_t^{(J)} (x_1 - \hat{x}_1) \right\|_2^2 \right] \left[\mathbb{E}_{z \sim p(z|x_t)} \|v(x_t|z)\|_2^2 \right] \\ &\leq \left\| \lambda_h \sigma_1 d \frac{\mathbb{E}[e^{-J(x_1)} v(x_t|z)]}{\mathbb{E}[e^{-J(x_1)}] e^{-J(\hat{x}_1)}} \right\|_2^2 + \left\| \frac{\lambda_h \sigma_1 d}{2e^{-J(\hat{x}_1)}} \right\|_2^2 \mathbb{E} [\|v(x_t|z)\|_2^2]. \end{aligned} \quad (77)$$

Then we have

$$\begin{aligned} \|\delta g\|_2^2 &\leq \left\| \lambda_h \sigma_1 d \frac{\mathbb{E}[e^{-J(x_1)} v(x_t|z)]}{\mathbb{E}[e^{-J(x_1)}] e^{-J(\hat{x}_1)}} \right\|_2^2 + \left\| \frac{\lambda_h \sigma_1 d}{2e^{-J(\hat{x}_1)}} \right\|_2^2 \mathbb{E} [\|v(x_t|z)\|_2^2] \\ &= \left| \frac{\lambda_h \sigma_1 d}{e^{-J(\hat{x}_1)}} \right|^2 \left(\underbrace{\left\| \frac{\mathbb{E}[e^{-J(x_1)} v(x_t|z)]}{\mathbb{E}[e^{-J(x_1)}]} \right\|_2^2}_{C_1} + \underbrace{\frac{1}{4} \mathbb{E} [\|v(x_t|z)\|_2^2]}_{C_2} \right), \end{aligned} \quad (78)$$

where we omit $z \sim p(z|x_t)$ in $\mathbb{E}_{z \sim p(z|x_t)}[\cdot]$ and simplify the notation to $\mathbb{E}[\cdot]$. Therefore, the approximation error of g^{local} is bounded by $(\lambda_h \sigma_1 d)^2 (C_1 + C_2) / |e^{-J(\mathbb{E}[x_1])}|$, where λ_h is the largest eigenvalue of $h_t^{(J)}$, the Hessian matrix of the objective function e^{-J} , σ_1 is the L2 norm of the covariance matrix, d is the sample dimensionality, C_1 is a constant that has to do with the norm of the new VF, and C_2 is the variance of the original *conditional* VF. Some intuitions can be emphasized:

1. The error is small when J is smooth, in which case the Hessian of e^{-J} will approach zero. This corresponds to the mild guidance, where approximation-based g^{local} works well.
2. The error is small when σ_1 is small, *i.e.* the covariance matrix Σ_{11} has a small Frobenius norm. This is the case when the flow time $t \rightarrow 1$ (and $\sigma_t = 0$), where x_t predicts x_1 well.
3. The magnitude of $|e^{-J(\mathbb{E}[x_1])}|$ reflects how well the *estimated generated sample* $\mathbb{E}[x_1]$ matches the objective J given the current x_t . If $\mathbb{E}[x_1]$ lies inside the region where J is small, *i.e.*, $\mathbb{E}[x_1]$ is a good sample, then the approximate guidance will be more accurate as the optimization is conducted locally, and gradient can reflect the landscape well. If $J(\mathbb{E}[x_1])$ is large, the gradient is almost randomly exploring the sample space, producing a larger approximation error.
4. The cases where C_1 and C_2 are small are not necessarily those where the guidance is of better accuracy. Because of the small norm of the VF, the error in the guidance VF will likely cause a larger deviation due to increased relative error.

A.10. Estimation of \hat{x}_1

Under the *affine path assumption* (Assumption 3.2), we can estimate the expectation of x_1 under the distribution $p(z|x_t)$. This is a well-known trick (Lipman et al., 2024; Pokle et al., 2024), but our analysis includes the dependent coupling case.

Since the flow matching model learns

$$v_\theta(x_t, t) \approx v_t(x_t) = \mathbb{E}_{z \sim p(z|x_t)} [v(x_t|z)], \quad (79)$$

using the *affine path assumption* ($x_t = \alpha_t x_1 + \beta_t x_0 + \sigma_t \dot{\sigma}_t \varepsilon$),

$$v(x_t|z) = \frac{d}{dt} x_t = (\dot{\alpha}_t x_1 + \dot{\beta}_t x_0 + \dot{\sigma}_t \varepsilon), \quad (80)$$

so

$$v_t(x_t) = \mathbb{E}_{z \sim p(z|x_t)} [\dot{\alpha}_t x_1 + \dot{\beta}_t x_0 + \dot{\sigma}_t \varepsilon]. \quad (81)$$

Meanwhile, taking the expectation of x_t under $p(z|x_t)$ yields

$$\underbrace{\mathbb{E}_{z \sim p(z|x_t)} [x_t]}_{\text{because } \int z p(z|x_t) dz = 1} = \mathbb{E}_{z \sim p(z|x_t)} [\alpha_t x_1 + \beta_t x_0 + \sigma_t \varepsilon]. \quad (82)$$

Then, by using Eq. 81 and 82, we can eliminate either \hat{x}_0 or \hat{x}_1 in each other's expression:

$$\hat{x}_0 := \mathbb{E}_{z \sim p(z|x_t)} [x_0] = \frac{\dot{\alpha}_t x_t - \alpha_t v_t(x_t)}{\beta_t \dot{\alpha}_t - \dot{\beta}_t \alpha_t} + \underbrace{\mathbb{E}_{z \sim p(z|x_t)} \left[\frac{-\dot{\alpha}_t \sigma_t \varepsilon + \alpha_t \dot{\sigma}_t \varepsilon}{\beta_t \dot{\alpha}_t - \dot{\beta}_t \alpha_t} \right]}_{:= \zeta_t^0} \quad (83)$$

$$\hat{x}_1 := \mathbb{E}_{z \sim p(z|x_t)} [x_1] = \frac{-\dot{\beta}_t x_t + \beta_t v_t(x_t)}{\beta_t \dot{\alpha}_t - \dot{\beta}_t \alpha_t} + \underbrace{\mathbb{E}_{z \sim p(z|x_t)} \left[\frac{\dot{\beta}_t \sigma_t \varepsilon - \beta_t \dot{\sigma}_t \varepsilon}{\beta_t \dot{\alpha}_t - \dot{\beta}_t \alpha_t} \right]}_{:= \zeta_t^1}. \quad (84)$$

It should be noted that we have assumed that σ_t is small and thus ζ_t^0 and ζ_t^1 are also small in the *affine path assumption* (Assumption 3.2):

$$\begin{aligned} \zeta_t^0 &= \frac{-\dot{\alpha}_t \sigma_t + \alpha_t \dot{\sigma}_t}{\beta_t \dot{\alpha}_t - \dot{\beta}_t \alpha_t} \mathbb{E}_{z \sim p(z|x_t)} [\varepsilon] \\ &= \frac{-\dot{\alpha}_t \sigma_t + \alpha_t \dot{\sigma}_t}{\beta_t \dot{\alpha}_t - \dot{\beta}_t \alpha_t} \int \frac{p(x_t|x_0, x_1) \pi(x_0, x_1)}{p(x_t)} \varepsilon dx_0 dx_1 \\ &= \frac{-\dot{\alpha}_t \sigma_t + \alpha_t \dot{\sigma}_t}{\beta_t \dot{\alpha}_t - \dot{\beta}_t \alpha_t} \int \frac{1}{p(x_t)} \mathbb{E}_{\varepsilon \sim p_\varepsilon(\varepsilon)} [\pi(x_0(x_t, x_1, \varepsilon), x_1) \varepsilon] dx_1. \end{aligned} \quad (85)$$

Since $\pi(x_0(x_t, x_1, \varepsilon), x_1)$ is a probability distribution that is assumed to be bounded, we denote $\max_\varepsilon \|\pi(x_0(x_t, x_1, \varepsilon), x_1)\| \leq \mathcal{M}(x_1, x_t)$, and thus

$$\lim_{\sigma_t \rightarrow 0, \dot{\sigma}_t \rightarrow 0} \zeta_t^0 \leq \lim_{\sigma_t \rightarrow 0, \dot{\sigma}_t \rightarrow 0} \left| \frac{-\dot{\alpha}_t \sigma_t + \alpha_t \dot{\sigma}_t}{\beta_t \dot{\alpha}_t - \dot{\beta}_t \alpha_t} \right| \cdot \int \left\| \frac{1}{p(x_t)} \right\|_2^2 \cdot \|\mathbb{E}_{\varepsilon \sim p_\varepsilon(\varepsilon)} \|\varepsilon\|_2 \mathcal{M}(x_1, x_t)\|_2^2 dx_1 = 0. \quad (86)$$

Since everything in the integral is independent of ε , x_0 , or σ_t , as $\sigma_t \rightarrow 0$ ζ_t^0 simply converges to zero. A similar approach can prove that $\lim_{\sigma_t \rightarrow 0, \dot{\sigma}_t \rightarrow 0} \zeta_t^1$ is also zero.

Next, we explain why we specifically care about the case where the small σ_t assumption holds. In independent coupling flow matching, $\sigma_t \varepsilon$ is exactly zero since we can use two of x_t, x_0 , and x_1 to express the third one. In dependent coupling flow matching, this assumption also holds for famous methods such as optimal transport conditional flow matching or Schrodinger Bridge conditional flow matching (Tong et al., 2024), where $\varepsilon \sim \mathcal{N}(0, I)$ and σ_t is set as a small constant. Therefore, the assumption that $\sigma_t \varepsilon$ is small in *affine path assumption* is general and applies to many existing flow matching methods. Hence, by approximating ζ_t^0 and ζ_t^1 as zero, we have the final estimation of \hat{x}_1

$$\hat{x}_0 \approx \frac{\dot{\alpha}_t x_t - \alpha_t v_\theta(x_t, t)}{\beta_t \dot{\alpha}_t - \dot{\beta}_t \alpha_t} \quad (87)$$

$$\hat{x}_1 \approx \frac{-\dot{\beta}_t x_t + \beta_t v_\theta(x_t)}{\beta_t \dot{\alpha}_t - \dot{\beta}_t \alpha_t}, \quad (88)$$

where Eq. (88) is just Eq. (5). Note the approximations become exact under Assumption 3.3.

A.11. Proof of g^{cov}

Here, we prove that under the affine path assumption (Assumption 3.2), Eq. (6)

$$g_t^{\text{local}} \approx g_t^{\text{cov}} = - \underbrace{\frac{\dot{\alpha}_t \beta_t - \dot{\beta}_t \alpha_t}{\beta_t}}_{\text{schedule}} \Sigma_{1|t} \nabla_{\hat{x}_1} J(\hat{x}_1), \quad (89)$$

where

$$g_t^{\text{local}} = -\mathbb{E}_{z \sim p(z|x_t)} [(x_1 - \hat{x}_1) v_{1|t}(x_t|z)] \nabla_{\hat{x}_1} J(\hat{x}_1). \quad (90)$$

Under the affine path $x_t = \alpha_t x_1 + \beta_t x_0 + \sigma_t \varepsilon$ the conditional vector field $v_{1|t}$ follows

$$v_{1|t}(x_t) = \dot{\alpha}_t x_1 + \dot{\beta}_t x_0 + \dot{\sigma}_t \varepsilon. \quad (91)$$

Plugging this into the definition of g^{local} and we get

$$\begin{aligned} g_t^{\text{local}} &= -\mathbb{E}_{z \sim p(z|x_t)} \underbrace{[(x_1 - \hat{x}_1)(\dot{\alpha}_t x_1 + \dot{\beta}_t x_0 + \dot{\sigma}_t \varepsilon)]}_{\text{substitute } x_0 \text{ with } x_1, \sigma_t \varepsilon, \text{ and } x_t} \nabla_{\hat{x}_1} J(\hat{x}_1) \\ &= -\mathbb{E}_{z \sim p(z|x_t)} [(x_1 - \hat{x}_1)(\dot{\alpha}_t x_1 + \frac{\dot{\beta}_t}{\beta_t}(x_t - \alpha_t x_1 - \sigma_t \varepsilon) + \dot{\sigma}_t \varepsilon)] \nabla_{\hat{x}_1} J(\hat{x}_1) \\ &= -\mathbb{E}_{z \sim p(z|x_t)} \left[(x_1 - \hat{x}_1) \left(\left(\frac{\beta_t \dot{\alpha}_t - \alpha_t \dot{\beta}_t}{\beta_t} \right) x_1 + \cancel{\mathcal{X}} + (\dot{\sigma}_t - \sigma_t) \varepsilon \right) \right] \nabla_{\hat{x}_1} J(\hat{x}_1) \\ &= -\frac{\dot{\alpha}_t \beta_t - \dot{\beta}_t \alpha_t}{\beta_t} \Sigma_{1|t} \nabla_{\hat{x}_1} J(\hat{x}_1) + \underbrace{(\sigma_t - \dot{\sigma}_t) \mathbb{E}_{z \sim p(z|x_t)} [(x_1 - \hat{x}_1) \varepsilon] \nabla_{\hat{x}_1} J(\hat{x}_1)}_{:= \Upsilon, \lim_{\sigma_t \rightarrow 0, \dot{\sigma}_t \rightarrow 0} \|\Upsilon\|_2^2 = 0}, \end{aligned} \quad (92)$$

where the x_t term is canceled out because $\int p(z|x_t)(x_1 - \mathbb{E}_{z \sim p(z|x_t)}[x_1]) dz = \int p(z|x_t)x_1 dz - \mathbb{E}_{z \sim p(z|x_t)}[x_1] = 0$, $\Sigma_{1|t} := \mathbb{E}_{z \sim p(z|x_t)} [(x_1 - \hat{x}_1)(x_1 - \hat{x}_1)]$, and the residual term that characterizes the approximation error (denoted as $\|\Upsilon\|_2^2$)

in Eq. (6) (restated in Eq. (89)) is

$$\begin{aligned}
 \Upsilon &= (\sigma_t - \dot{\sigma}_t) \mathbb{E}_{z \sim p(z|x_t)} [(x_1 - \hat{x}_1) \varepsilon] \nabla_{\hat{x}_1} J(\hat{x}_1) \\
 &= (\sigma_t - \dot{\sigma}_t) \nabla_{\hat{x}_1} J(\hat{x}_1) \int \frac{p(x_t|z)p(z)}{p(x_t)} (x_1 - \hat{x}_1) \varepsilon dx_1 dx_0 \\
 &= (\sigma_t - \dot{\sigma}_t) \nabla_{\hat{x}_1} J(\hat{x}_1) \int \frac{p(\sigma_t \varepsilon) \pi(x_0|x_1) p(x_1)}{p(x_t)} (x_1 - \hat{x}_1) \frac{1}{\sigma_t} (x_t - \alpha_t x_1 - \beta_t x_0) dx_1 dx_0 \\
 &= (\sigma_t - \dot{\sigma}_t) \nabla_{\hat{x}_1} J(\hat{x}_1) \int \frac{p(x_1)}{p(x_t)} (x_1 - \hat{x}_1) dx_1 \int p(\sigma_t \varepsilon) \pi(x_0|x_1) \varepsilon dx_0 \\
 &= (\sigma_t - \dot{\sigma}_t) \nabla_{\hat{x}_1} J(\hat{x}_1) \int \frac{p(x_1)}{p(x_t)} (x_1 - \hat{x}_1) \mathbb{E}_{\varepsilon \sim p_\varepsilon(\varepsilon)} \left[\pi \left(\frac{1}{\beta_t} (x_t - \alpha_t x_1 - \sigma_t \varepsilon) \mid x_1 \right) \varepsilon \right] dx_1, \tag{93}
 \end{aligned}$$

where $p_\varepsilon(\varepsilon)$ is the marginal distribution of ε . Suppose $\|\pi \left(\frac{1}{\beta_t} (x_t - \alpha_t x_1 - \sigma_t \varepsilon) \mid x_1 \right)\|_2^2 = \|\pi(x_0 \mid x_1)\|_2^2 \leq \mathcal{M}(x_1, x_t)$ (which is a function independent of ε , then

$$\begin{aligned}
 &\|\Upsilon\|_2^2 \\
 &\leq \|(\sigma_t - \dot{\sigma}_t) \nabla_{\hat{x}_1} J(\hat{x}_1)\|_2^2 \cdot \left\| \int \frac{p(x_1)}{p(x_t)} (x_1 - \hat{x}_1) \mathbb{E}_{\varepsilon \sim p_\varepsilon(\varepsilon)} \left[\pi \left(\frac{1}{\beta_t} (x_t - \alpha_t x_1 - \sigma_t \varepsilon) \mid x_1 \right) \varepsilon \right] dx_1 \right\|_2^2 \\
 &\leq \|(\sigma_t - \dot{\sigma}_t)\| \underbrace{\|\nabla_{\hat{x}_1} J(\hat{x}_1)\|_2^2}_{:=\mathcal{G}} \cdot \underbrace{\left\| \int \frac{p(x_1)}{p(x_t)} (x_1 - \hat{x}_1) \right\|_2^2}_{:=\mathcal{Q}} \cdot \underbrace{\left\| \mathbb{E}_{\varepsilon \sim p_\varepsilon(\varepsilon)} \left[\pi \left(\frac{1}{\beta_t} (x_t - \alpha_t x_1 - \sigma_t \varepsilon) \mid x_1 \right) \varepsilon \right] \right\|_2^2}_{\leq (\mathcal{M} \mathbb{E}_{\varepsilon \sim p_\varepsilon(\varepsilon)} [\|\varepsilon\|_2])^2 \leq \mathcal{M}^2 \text{Var}_{p_\varepsilon}} \\
 &\leq \|(\sigma_t - \dot{\sigma}_t)\| \mathcal{G}(x_t) \int \mathcal{Q}(x_1, x_t) \mathcal{M}^2(x_1, x_t) \text{Var}_{p_\varepsilon} dx_1, \tag{94}
 \end{aligned}$$

all of which are independent on x_0 or σ_t . Thus,

$$\lim_{\sigma \rightarrow 0, \dot{\sigma}_t \rightarrow 0} \|\Upsilon\|_2^2 = 0. \tag{95}$$

A.12. The Jacobian Trick

We prove the Jacobian Trick here.

Proposition A.5. *The Jacobian trick. Under Assumption 3.3, the inverse covariance matrix of $p(x_1|x_t)$, $\Sigma_{1|t}$, is affine to the Jacobian of the VF $\frac{\partial v_t}{\partial x_t}$, and is proportional to the Jacobian $\frac{\partial \hat{x}_1}{\partial x_t}$:*

$$\Sigma_{1|t} = \frac{\beta_t^2}{\alpha_t(\dot{\alpha}_t \beta_t - \dot{\beta}_t \alpha_t)} (-\dot{\beta}_t + \beta_t \frac{\partial v_t}{\partial x_t}) = \frac{\beta_t^2}{\alpha_t} \frac{\partial \hat{x}_1}{\partial x_t}.$$

Proof.

To begin with, we prove $\Sigma_{1|t} = \frac{\beta_t^2}{\alpha_t} \frac{\partial \hat{x}_1}{\partial x_t}$. A similar conclusion has been proved in [Ye et al. \(2024\)](#). We generalize their proof to affine Gaussian path flow matching:

Recall from Eq. (88) that

$$\hat{x}_1 = -\frac{\dot{\beta}_t}{\dot{\alpha}_t \beta_t - \dot{\beta}_t \alpha_t} x_t + \frac{\beta_t}{\dot{\alpha}_t \beta_t - \dot{\beta}_t \alpha_t} v_t \tag{96}$$

and

$$\hat{x}_0 = \frac{\dot{\alpha}_t}{\dot{\alpha}_t \beta_t - \dot{\beta}_t \alpha_t} x_t - \frac{\alpha_t}{\dot{\alpha}_t \beta_t - \dot{\beta}_t \alpha_t} v_t. \tag{97}$$

So we have the Jacobian trick

$$\frac{\partial \hat{x}_1}{\partial x_t} = -\frac{\dot{\beta}_t}{\dot{\alpha}_t \beta_t - \dot{\beta}_t \alpha_t} + \frac{\beta_t}{\dot{\alpha}_t \beta_t - \dot{\beta}_t \alpha_t} \frac{\partial v_t(x_t)}{\partial x_t}, \tag{98}$$

and because the VF is associated with the score

$$v_t(x_t) = \frac{\beta_t(\dot{\alpha}_t\beta_t - \dot{\beta}_t\alpha_t)}{\alpha_t} \nabla_{x_t} \log p_t(x_t) + \frac{\dot{\alpha}_t}{\alpha_t} x_t. \quad (99)$$

Next, we try to prove

$$\nabla_{x_t} \nabla_{x_t} \log p_t(x_t) = -\frac{1}{\beta_t^2} + \frac{\alpha_t^2}{\beta_t^4} \Sigma_{x_1 x_1}, \quad (100)$$

which allows us to connect the derivative of the score $\nabla_{x_t}^2 \log p(x_t)$ ¹⁰ with the covariance matrix $\Sigma_{1|t}$.

$$\begin{aligned} \nabla_{x_t}^2 \log p(x_t) &= \frac{\nabla_{x_t}^2 p(x_t)}{p(x_t)} - \nabla_{x_t} \log p(x_t) \nabla_{x_t} \log p(x_t) \\ &= \frac{1}{p(x_t)} \int p(x_1) \underbrace{\nabla_{x_t}^2 p(x_t|x_1)}_{\text{using } \nabla^2 p = p \nabla^2 \log p + p (\nabla \log p)^2} dx_1 - \nabla_{x_t} \log p(x_t) \nabla_{x_t} \log p(x_t) \\ &= \frac{1}{p(x_t)} \int p(x_1) (p(x_t|x_1) \nabla_{x_t}^2 \log p(x_t|x_1) + p(x_t|x_1) \nabla_{x_t} \log p(x_t|x_1) \nabla_{x_t} \log p(x_t|x_1)) dx_1 \\ &\quad - \nabla_{x_t} \log p(x_t) \nabla_{x_t} \log p(x_t) \\ &= \mathbb{E}_{x_1 \sim p(x_1|x_t)} [\nabla_{x_t}^2 \log p(x_t|x_1) + \nabla_{x_t} \log p(x_t|x_1) \nabla_{x_t} \log p(x_t|x_1)] - \nabla_{x_t} \log p(x_t) \nabla_{x_t} \log p(x_t) \\ &= \mathbb{E}_{x_1 \sim p(x_1|x_t)} \left[-\frac{1}{\beta_t^2} + \left(\frac{x_t - \alpha_t x_1}{\beta_t^2} \right)^2 \right] - \left(\frac{x_t - \alpha_t \mathbb{E}_{x_1 \sim p(x_1|x_t)}[x_1]}{\beta_t^2} \right)^2 \\ &= -\frac{1}{\beta_t^2} + \frac{\alpha_t^2}{\beta_t^4} (\mathbb{E}[x_1 x_1^T] - \mathbb{E}[x_1] \mathbb{E}[x_1]^T) \\ &= -\frac{1}{\beta_t^2} + \frac{\alpha_t^2}{\beta_t^4} \Sigma_{x_1 x_1}. \end{aligned} \quad (101)$$

Then by combining Eq. (99) and (100) we have

$$\begin{aligned} \frac{\partial \hat{x}_1}{\partial x_t} &= -\frac{\dot{\beta}_t}{\dot{\alpha}_t \beta_t - \dot{\beta}_t \alpha_t} + \frac{\beta_t}{\dot{\alpha}_t \beta_t - \dot{\beta}_t \alpha_t} \left(\frac{\beta_t(\dot{\alpha}_t \beta_t - \dot{\beta}_t \alpha_t)}{\alpha_t} \nabla_{x_t} \nabla_{x_t} \log p_t(x_t) + \frac{\dot{\alpha}_t}{\alpha_t} \right) \\ &= -\frac{\dot{\beta}_t}{\dot{\alpha}_t \beta_t - \dot{\beta}_t \alpha_t} + \frac{\beta_t}{\dot{\alpha}_t \beta_t - \dot{\beta}_t \alpha_t} \left(\frac{\beta_t(\dot{\alpha}_t \beta_t - \dot{\beta}_t \alpha_t)}{\alpha_t} \left(-\frac{1}{\beta_t^2} + \frac{\alpha_t^2}{\beta_t^4} \Sigma_{x_1 x_1} \right) + \frac{\dot{\alpha}_t}{\alpha_t} \right) \\ &= \frac{\alpha_t}{\beta_t^2} \Sigma_{x_1 x_1}. \end{aligned} \quad (102)$$

Inserting back Eq. (88) and we prove

$$\Sigma_{1|t} = \frac{\beta_t^2}{\alpha_t(\dot{\alpha}_t \beta_t - \dot{\beta}_t \alpha_t)} \left(-\dot{\beta}_t + \beta_t \frac{\partial v_t}{\partial x_t} \right). \quad (103)$$

A.13. Proof for $g_t^{\text{sim-inv}}$

We begin with

$$g_t^{\text{sim-inv}}(x_t) = \int \left(\frac{e^{-J(x_1)}}{\tilde{Z}_t} - 1 \right) v_{t|z}(x_t|z) \tilde{p}(z|x_t) dz, \quad (104)$$

$$\text{where } \tilde{Z}_t = \int e^{-J(x_1)} \tilde{p}(z|x_t) dz, \quad (105)$$

¹⁰We use $\nabla \nabla$ and ∇^2 interchangeably, with a little abuse of notation. It should not cause confusion since the size of the terms in the equations must match.

and approximate $p(x_1|x_t)$ with $\mathcal{N}(x_1; \hat{x}_1, \Sigma_t)$ where $\hat{x}_1 := \mathbb{E}_{z \sim p(z|x_t)}[x_1]$ and Σ_t is already known.

With Assumption 3.2 and assuming $e^{-J(x_1)} = \mathcal{N}(Hx_1; y, \sigma_y I)$, we have

$$\begin{aligned} \tilde{Z}_t &= \int e^{-J(x_1)} \tilde{p}(z|x_t) dz \\ &= \int e^{-\frac{1}{2\sigma_y^2} \|y - Hx_1\|_2^2 - \frac{1}{2} (z - \hat{z})^T \Sigma_t^{-1} (z - \hat{z})} dz \end{aligned} \quad (106)$$

where $\hat{z} = (\hat{x}_0, \hat{x}_1)$ is the expectation of z under $p(z|x_t)$. Note H operates on x_1 only, and we pad the blocks related to x_0 with zero in H .

Then, by inserting the Gaussian approximation

$$g_t^{\text{sim-inv}}(x_t) \approx \int \left(\frac{e^{-J(x_1)}}{\tilde{Z}_t} - 1 \right) (\dot{\alpha}_t x_1 + \dot{\beta}_t x_0) \underbrace{\tilde{p}(z|x_t)}_{\text{Gaussian}} dz \quad (107)$$

$$= \int \underbrace{\frac{1}{\tilde{Z}_t} \exp\left(-\frac{1}{2\sigma_y^2} \|y - Hx_1\|_2^2 - \frac{1}{2} (z - \hat{z})^T \Sigma_t^{-1} (z - \hat{z})\right)}_{:= \tilde{p}(z|x_t)} (\dot{\alpha}_t x_1 + \dot{\beta}_t x_0) dz - v_t(x_t). \quad (108)$$

Remark A.6. Note that Σ_t^{-1} couples x_0 and x_1 . This is a fundamental feature of dependent couplings $\pi(x_0, x_1)$. However, it may seem tempting to further assume that Σ_t^{-1} is diagonal or even a scalar. It should be noted that this assumption completely discards the dependency of x_0 and x_1 in the coupling, and thus, we try to avoid that in the dependent coupling case.

For clarity, we need to express Σ_t^{-1} with

$$\Sigma_t^{-1} \triangleq \begin{pmatrix} \Xi_{00} & \Xi_{01} \\ \Xi_{10} & \Xi_{11} \end{pmatrix}. \quad (109)$$

Then, the distribution $\exp\left(-\frac{1}{2\sigma_y^2} \|y - Hx_1\|_2^2 - \frac{1}{2} (z - \hat{z})^T \Sigma_t^{-1} (z - \hat{z})\right)$ is still a Gaussian, and to estimate the expectation of $z = (x_0, x_1)$ we need to simply the probability density of this Gaussian into a standard form.

$$\begin{aligned} \tilde{Z}_t \tilde{p}(z|x_t) &= \exp\left(-\frac{1}{2\sigma_y^2} \|y - Hx_1\|_2^2 - \frac{1}{2} (z - \hat{z})^T \Sigma_t^{-1} (z - \hat{z})\right) \\ &= \exp\left(-\frac{1}{2\sigma_y^2} (\|y\|^2 - 2\langle y, Hx_1 \rangle + \|Hx_1\|_2^2) - \frac{1}{2} z^T \Sigma_t^{-1} z - \frac{1}{2} \hat{z}^T \Sigma_t^{-1} \hat{z} + \underbrace{\langle z, \Sigma_t^{-1} \hat{z} \rangle}_{\text{since } \Sigma_t^{-1} = \Sigma_t^{-1T}}\right) \\ &= \exp\left(-\frac{1}{2} \left(z^T \left(\frac{H^T H}{\sigma_y^2} + \Sigma_t^{-1} \right) z - 2\langle z, \frac{H^T}{\sigma_y^2} y + \Sigma_t^{-1} \hat{z} \rangle + \left(\frac{1}{2\sigma_y^2} \|y\|^2 + \frac{1}{2} \hat{z}^T \Sigma_t^{-1} \hat{z} \right) \right)\right), \end{aligned} \quad (110)$$

It is obvious that the mean of this Gaussian is

$$\mu = \underbrace{\left(\frac{H^T H}{\sigma_y^2} + \Sigma_t^{-1} \right)^{-1}}_{:= P} \left(\frac{H^T}{\sigma_y^2} y + \Sigma_t^{-1} \hat{z} \right), \quad (111)$$

where we can find P 's blocks to be

$$P = \begin{pmatrix} \Xi_{00} & \Xi_{01} \\ \Xi_{10} & \Xi_{11} + \frac{H^T H}{\sigma_y^2} \end{pmatrix}. \quad (112)$$

Then, by computing $\mu = (\hat{x}_0, \hat{x}_1)$ we can compute $g_t^{\text{sim-inv}} + v_t$, where \hat{x}_0, \hat{x}_1 are the x_0 and x_1 term in the integral of Eq. (108), because

$$g_t^{\text{sim-inv}} + v_t = \mathbb{E}_{z \sim \tilde{p}(z|x_t)}[\dot{\alpha}_t x_1 + \dot{\beta}_t x_0]. \quad (113)$$

To simplify, insert back v_t to get

$$\begin{aligned}
 g_t^{\text{sim-inv}} &= \mathbb{E}_{z \sim \tilde{p}(z|x_t)}[\dot{\alpha}_t x_1 + \dot{\beta}_t x_0] - \mathbb{E}_{z \sim \tilde{p}(z|x_t)}[\dot{\alpha}_t x_1 + \dot{\beta}_t x_0] \\
 &= (\dot{\beta}_t I \quad \dot{\alpha}_t I) P^{-1} \left(\begin{pmatrix} 0 \\ \frac{H^T}{\sigma_y^2} y \end{pmatrix} + \left(\begin{pmatrix} \Xi_{00} & \Xi_{01} \\ \Xi_{10} & \Xi_{11} \end{pmatrix} - P \right) \begin{pmatrix} \hat{x}_0 \\ \hat{x}_1 \end{pmatrix} \right) \\
 &= (\dot{\beta}_t I \quad \dot{\alpha}_t I) P^{-1} \left(\begin{pmatrix} 0 \\ \frac{H^T}{\sigma_y^2} y \end{pmatrix} + \begin{pmatrix} 0 & 0 \\ 0 & -\frac{H^T H}{\sigma_y^2} \end{pmatrix} \begin{pmatrix} \hat{x}_0 \\ \hat{x}_1 \end{pmatrix} \right) \\
 &= (\dot{\beta}_t I \quad \dot{\alpha}_t I) P^{-1} \left(\frac{H^T}{\sigma_y^2} y - \frac{H^T H}{\sigma_y^2} \hat{x}_1 \right). \tag{114}
 \end{aligned}$$

Usually, $(\dot{\beta}_t I \quad \dot{\alpha}_t I) P^{-1}$ is difficult to obtain:

$$g_t^{\text{sim-inv}} = (\dot{\beta}_t P_{01}^{-1} + \dot{\alpha}_t P_{11}^{-1}) \left(\frac{H^T}{\sigma_y^2} y - \frac{H^T H}{\sigma_y^2} \hat{x}_1 \right), \tag{115}$$

where P_{01}^{-1} and P_{11}^{-1} requires computing the inversion of P and thus in general intractable. Using block matrix inversion, we have

$$g_t^{\text{sim-inv}} = \left(-\dot{\beta}_t \Xi_{11}^{-1} \Xi_{01} \left(\Xi_{11} + \frac{H^T H}{\sigma_y^2} - \Xi_{10} \Xi_{11}^{-1} \Xi_{01} \right)^{-1} + \dot{\alpha}_t \left(\Xi_{11} + \frac{H^T H}{\sigma_y^2} - \Xi_{10} \Xi_{11}^{-1} \Xi_{01} \right)^{-1} \right) \left(\frac{H^T}{\sigma_y^2} y - \frac{H^T H}{\sigma_y^2} \hat{x}_1 \right). \tag{116}$$

For general (possibly coupled) affine path flow matching, we can make approximations and set the *blocks* in Σ_t^{-1} to scalars. It should be noted that this Gaussian assumption can still capture some coupling between x_0 and x_1 since the off-diagonal blocks Ξ_{01} and Ξ_{10} are not set to zero. Specifically, we have

$$g_t^{\text{sim-inv-A}} = -\lambda_t \left(\frac{\sigma_y^2}{r_t^2} + H^T H \right)^{-1} H^T (y - H \hat{x}_1), \tag{117}$$

where λ_t and r_t^2 are hyperparameters. λ_t approximates $\dot{\alpha}_t - \dot{\beta}_t \Xi_{11}^{-1} \Xi_{01}$, absorbing the flow schedule.

Special case: the non-coupled affine Gaussian path

Next, we prove that $g_t^{\text{sim-inv}}$ covers IIGDM (Song et al., 2023a) and OT-ODE (Pokle et al., 2024) as special cases. Under the uncoupled affine Gaussian path assumption (Assumption 3.3), one may think that the covariance matrix is block diagonal, but it is false: x_0 and x_1 are still dependent on each other in the distribution $p(z|x_t) = p(x_0, x_1|x_t)$ even if the coupling is independent. In the uncoupled case, the probability graph is $x_0 \rightarrow x_t \leftarrow x_1$, so although x_0 and x_1 are marginally independent ($\pi(x_0, x_1) = p(x_0)p(x_1)$), their conditional can be dependent $p(x_0, x_1|x_t) \neq p(x_0|x_t)p(x_1|x_t)$. Then, we notice the uncoupled path is

$$x_t = \alpha_t x_1 + \beta_t x_0, \tag{118}$$

so we actually should not have approximated the distribution $p_{z|x_t}$ as a Gaussian in the uncoupled case. Fortunately, there is a workaround to make x_0 almost entirely dependent on x_1 . We can set $x_0 = -\frac{\alpha_t}{\beta_t} x_1 + \frac{1}{\beta_t} x_t + \xi \epsilon$, where $\epsilon \sim \mathcal{N}(0, I)$, and setting $\xi \rightarrow 0$ gives our desired uncoupled path results. The covariance matrix of x_0 and x_1 to:

$$\Sigma_t = \begin{pmatrix} \frac{\alpha_t^2}{\beta_t^2} \Sigma_{x_1 x_1} + \xi^2 I & -\frac{\alpha_t}{\beta_t} \Sigma_{x_1 x_1} \\ -\frac{\alpha_t}{\beta_t} \Sigma_{x_1 x_1} & \Sigma_{x_1 x_1} \end{pmatrix} \tag{119}$$

Note that $\Sigma_{x_1 x_1}^{-1} \neq \Xi_{11}$ as Ξ_{11} is a block in the inversion of the larger matrix Σ_t . Next we compute Σ_t^{-1} :

$$\begin{aligned}
 \Sigma_t^{-1} &= \begin{pmatrix} \frac{1}{\xi^2} I & -\frac{\alpha}{\xi^2} I \\ -\frac{\alpha}{\xi^2} I & \left(\Sigma_{x_1 x_1} - \alpha^2 \Sigma_{x_1 x_1} (\alpha^2 \Sigma_{x_1 x_1} + \xi^2 I)^{-1} \Sigma_{x_1 x_1} \right)^{-1} \end{pmatrix} \\
 &= \begin{pmatrix} \frac{1}{\xi^2} I & \frac{\alpha_t}{\beta_t \xi^2} I \\ \frac{\alpha_t}{\beta_t \xi^2} I & \left(\Sigma_{x_1 x_1} - \Sigma_{x_1 x_1} (\Sigma_{x_1 x_1} + \frac{\beta_t^2}{\alpha_t^2} \xi^2 I)^{-1} \Sigma_{x_1 x_1} \right)^{-1} \end{pmatrix}. \tag{120}
 \end{aligned}$$

where $a = -\frac{\alpha_t}{\beta_t}$. Therefore, $\Xi_{11} \rightarrow \infty$, $\Xi_{01} = \Xi_{10} = \frac{\alpha_t}{\beta_t \xi^2} I \rightarrow \infty$, and $\Xi_{00} = \frac{1}{\xi^2} I \rightarrow \infty$. Thus, we need more detailed calculations to get the result:

$$g_t^{\text{sim-inv-diffusion}} = \left(-\dot{\beta}_t \underbrace{\Xi_{11}^{-1} \Xi_{01}}_{\rightarrow \frac{\alpha_t}{\beta_t} I} \left(\underbrace{\Xi_{11}}_{\rightarrow \infty} + \frac{H^T H}{\sigma_y^2} - \underbrace{\Xi_{10} \Xi_{11}^{-1} \Xi_{01}}_{\rightarrow \infty} \right)^{-1} + \dot{\alpha}_t \left(\underbrace{\Xi_{11}}_{\rightarrow \infty} + \frac{H^T H}{\sigma_y^2} - \underbrace{\Xi_{10} \Xi_{11}^{-1} \Xi_{01}}_{\rightarrow \infty} \right)^{-1} \right) \left(\frac{H^T}{\sigma_y^2} y - \frac{H^T H}{\sigma_y^2} \hat{x}_1 \right). \quad (121)$$

Obviously we want to find the finite term left in $\Xi_{11} - \Xi_{10} \Xi_{11}^{-1} \Xi_{01}$:

$$\begin{aligned} & \lim_{\xi \rightarrow 0} \Xi_{11} - \Xi_{10} \Xi_{11}^{-1} \Xi_{01} \\ &= \lim_{\xi \rightarrow 0} \left(\Sigma_{x_1 x_1} - \Sigma_{x_1 x_1} \left(\Sigma_{x_1 x_1} + \frac{\beta_t^2}{\alpha_t^2} \xi^2 I \right)^{-1} \Sigma_{x_1 x_1} \right)^{-1} - \Xi_{10} \Xi_{11}^{-1} \Xi_{01} \\ &= \lim_{\xi \rightarrow 0} \left(\Sigma_{x_1 x_1} - \Sigma_{x_1 x_1} \left(\Sigma_{x_1 x_1} + \frac{\beta_t^2}{\alpha_t^2} \xi^2 I \right)^{-1} \Sigma_{x_1 x_1} \right)^{-1} - \frac{\alpha_t^2}{\beta_t^2 \xi^2} \\ &= \lim_{\xi \rightarrow 0} \left(\Sigma_{x_1 x_1} \left(\Sigma_{x_1 x_1} + \frac{\beta_t^2}{\alpha_t^2} \xi^2 I \right)^{-1} \left(\left(\Sigma_{x_1 x_1} + \frac{\beta_t^2}{\alpha_t^2} \xi^2 I \right) - \Sigma_{x_1 x_1} \right) \right)^{-1} - \frac{\alpha_t^2}{\beta_t^2 \xi^2} \\ &= \lim_{\xi \rightarrow 0} \frac{\alpha_t^2}{\beta_t^2 \xi^2} \left(\Sigma_{x_1 x_1} \left(\Sigma_{x_1 x_1} + \frac{\beta_t^2}{\alpha_t^2} \xi^2 I \right)^{-1} - 1 \right) \\ &= \Sigma_{x_1 x_1}^{-1}. \end{aligned} \quad (122)$$

Now we have

$$g_t^{\text{sim-inv-diffusion}} = \frac{\dot{\alpha}_t \beta_t - \dot{\beta}_t \alpha_t}{\beta_t} \left(\Sigma_{x_1 x_1}^{-1} + \frac{H^T H}{\sigma_y^2} \right)^{-1} \left(\frac{H^T}{\sigma_y^2} y - \frac{H^T H}{\sigma_y^2} \hat{x}_1 \right). \quad (123)$$

This is essentially the same formulation as in IIGDM (Song et al., 2023a) and OT-ODE (Pokle et al., 2024). Next, we will make some trivial conversions to cover the formulations exactly.

In diffusion paths (Assumption 3.4) we proved that $\frac{\partial \hat{x}_t}{\partial x_t} = \frac{\alpha_t}{\beta_t^2} \Sigma_{1|t}$ where $\Sigma_{1|t}$ is just what we denote $\Sigma_{x_1 x_1}$ here. Equivalently,

$$\frac{\partial x_t}{\partial \hat{x}_1} = \frac{\beta_t^2}{\alpha_t} \Sigma_{x_1 x_1}^{-1}. \quad (124)$$

Thus,

$$\begin{aligned} g_t^{\text{sim-inv-diffusion}} &= \frac{\dot{\alpha}_t \beta_t - \dot{\beta}_t \alpha_t}{\beta_t} \left(\Sigma_{x_1 x_1}^{-1} + \frac{H^T H}{\sigma_y^2} \right)^{-1} \left(\frac{H^T}{\sigma_y^2} y - \frac{H^T H}{\sigma_y^2} \hat{x}_1 \right) \\ &= \frac{\dot{\alpha}_t \beta_t - \dot{\beta}_t \alpha_t}{\beta_t} \Sigma_{x_1 x_1} \left(\sigma_y^2 I + \Sigma_{x_1 x_1} H^T H \right)^{-1} H^T (y - H \hat{x}_1) \\ &= \frac{\dot{\alpha}_t \beta_t - \dot{\beta}_t \alpha_t}{\beta_t} \frac{\beta_t^2}{\alpha_t} \frac{\partial \hat{x}_1}{\partial x_t} \left(\sigma_y^2 I + \Sigma_{x_1 x_1} H^T H \right)^{-1} (H^T y - H^T H \hat{x}_1) \\ &= \frac{\beta_t (\dot{\alpha}_t \beta_t - \dot{\beta}_t \alpha_t)}{\alpha_t} \left(\frac{\partial \hat{x}_1}{\partial x_t} \left(\sigma_y^2 I + \Sigma_{x_1 x_1} H^T H \right)^{-1} (H^T y - H^T H \hat{x}_1) \right) \\ &= \frac{\beta_t (\dot{\alpha}_t \beta_t - \dot{\beta}_t \alpha_t)}{\alpha_t} \left((y - H \hat{x}_1)^T H \left(\sigma_y^2 I + \Sigma_{x_1 x_1} H^T H \right)^{-1} \frac{\partial \hat{x}_1}{\partial x_t} \right)^T. \end{aligned} \quad (125)$$

Now we make the same approximation in Π GDM that $\Sigma_{x_1x_1} = r_t^2 I$. Then by noticing that

$$\begin{aligned} (\sigma_y^2 I + r_t^2 H H^T) H &= H (\sigma_y^2 I + r_t^2 H^T H) \\ H (\sigma_y^2 I + r_t^2 H^T H)^{-1} &= (\sigma_y^2 I + r_t^2 H H^T)^{-1} H \end{aligned} \quad (126)$$

We exactly cover

$$g_t^{\text{sim-inv-}\Pi\text{GDM}} = \frac{\beta_t(\dot{\alpha}_t\beta_t - \dot{\beta}_t\alpha_t)}{\alpha_t} \left((y - H\hat{x}_1)^T (\sigma_y^2 I + r_t^2 H^T H)^{-1} H \frac{\partial \hat{x}_1}{\partial x_t} \right)^T, \quad (127)$$

and the scheduler $\frac{\beta_t(\dot{\alpha}_t\beta_t - \dot{\beta}_t\alpha_t)}{\alpha_t}$ in the path $\alpha_t = t, \beta_t = 1 - t$ becomes $\frac{1-t}{t}$, which exactly covers the schedule in OT-ODE which takes the same path. In addition, we can also directly compute $\Sigma_{x_1x_1}$ using $\frac{\partial \hat{x}_1}{\partial x_t}$ instead of approximating it with r_t . This corresponds to the approach in [Boys et al. \(2023\)](#), which uses the Jacobian to acquire the covariance and then remove the approximation error in computing $(\sigma_y^2 I + \Sigma_{x_1x_1} H^T H)^{-1}$.

Remark A.7. Starting from the more general assumption of affine path flow matching, we derived the guidance compatible with dependent coupling flow matching, including OT-CFM. The fact that our guidance can exactly cover classical diffusion guidance like Π GDM and affine Gaussian path flow matching guidance like OT-ODE verifies the validity of our theory.

B. Experiment Details

B.1. Synthetic Dataset Experiment Details

The training of flow matching models involves sampling x_0 from a source distribution of {Circle,8 Gaussians,Uniform,Gaussian} and sampling x_1 from the target distribution of {S-Curve,Moons,8 Gaussians}. The model backbone is an MLP of 4 layers with a hidden dimension of 256. The models are trained 1e5 steps.

To evaluate the asymptotic exactness of g^{MC} , we compute the Wasserstein-2 (\mathcal{W}_2) distance of the samples generated under guidance, with the ground truth energy-weighted distribution $p(x_1)e^{-J(x_1)}/Z$. Since the source distribution $p(x_1)$ is learned, the flow matching model itself has a small error w , which can also be quantified using the \mathcal{W}_2 distance. In principle, this error w characterizes the performance upper bound of the guided distribution: the \mathcal{W}_2 distance of the guided distribution will, in principle, not be significantly lower than w . The result is shown in Figure 4, where w is demonstrated using the dashed line.

B.2. Planning Experiment Details

Settings. The experiment leverages the D4RL (Datasets for Deep Data-Driven Reinforcement Learning) dataset ([Fu et al., 2020](#)), specifically the Locomotion datasets, which is a common choice to evaluate offline reinforcement learning methods, as well as offline planning methods ([Janner et al., 2022](#); [Dou & Song, 2024](#)). The datasets contain non-expert behaviors from which the model is required to learn the optimal policy, such as a mixture of expert and medium-level experts, or the training replay buffer of an RL agent.

To evaluate the performance of different guidance methods, we conduct experiments on offline RL tasks where generative models have been used as planners ([Janner et al., 2022](#); [Chen & Lipman, 2024](#)). Our setting is based on that of a classical generative planner called Diffuser ([Janner et al., 2022](#)), where a generative model generates a state-action pair sequence of multiple future steps, and another critic model that predicts the future reward¹¹ of the generated plans. The generative model then optimizes its plans using guidance for higher future rewards. Following the formulation in [Levine \(2018\)](#); [Janner et al. \(2022\)](#), the optimization is realized through sampling from $\frac{1}{Z}p(x_1)e^{R(x_1)}$, where R is the critic model. This aligns with the goal of guidance which we discussed in this paper, so we chose this experiment to evaluate different guidance methods.

Baseline Methods. The results of the two baselines are collected from the literature ([Janner et al., 2022](#)). Behavior cloning refers to using a Gaussian distribution to fit the offline behavior distribution. Diffuser refers to using a diffusion model to learn the offline behavior and then guiding the model to generate plans with higher expected future returns. Note that the Diffuser also adopts training-based guidance, which requires re-training the guidance model when switching to a new objective function. On the contrary, the training-free guidances $g^{\text{cov-A}}$, $g^{\text{cov-G}}$, $g^{\text{sim-MC}}$, and g^{MC} has zero-shot generalization

¹¹Formally, it predicts the discounted return-to-go.

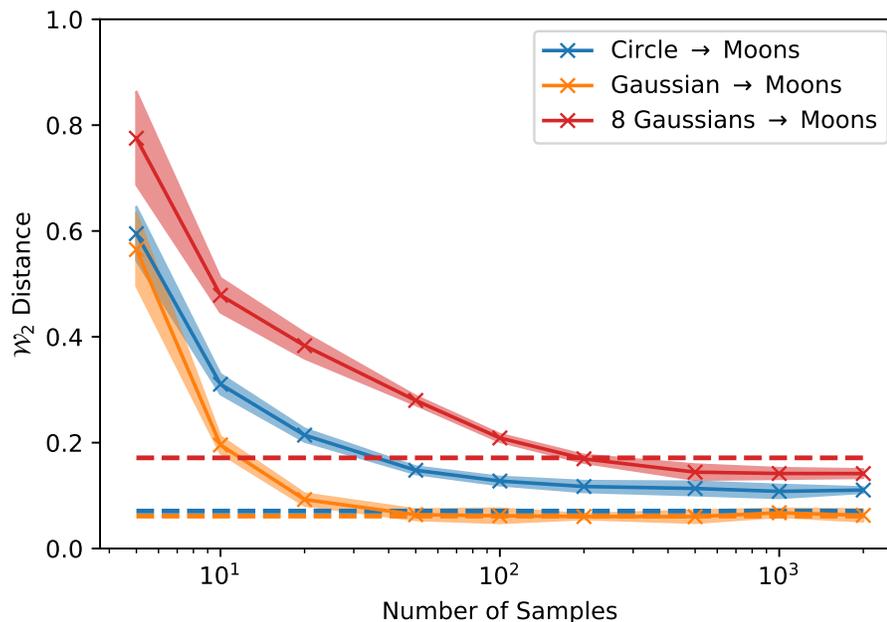


Figure 4: Error scaling with Monte Carlo sample number. In the synthetic dataset, the guidance performance (\mathcal{W}_2 distance between the generated distribution and the ground truth energy weighted distribution $p(x_1)e^{-J(x_1)}/Z$) decreases as the number of Monte Carlo samples increases. The dashed lines denote the \mathcal{W}_2 distance between the learned unguided distribution and the original ground truth distribution $p(x_1)$. The reason why the guided generation errors (crosses) do not converge to the dashed lines is that they measure the \mathcal{W}_2 distance of $p(x_1)$ and $p(x_1)e^{-J(x_1)}/Z$, respectively.

ability for (Zhou et al., 2024) the objective function. In the experiment results, we do not highlight the results of baselines as we focus on the comparison between different guidance methods.

Hyperparameters. As we mentioned, a generative model is first trained on the offline dataset as a behavior-cloning method but captures the actual action distribution rather than approximating it with a Gaussian. The generative model we consider here is the CFM or mini-batch optimal transport CFM with affine paths $\alpha_t = t, \beta_t = 1 - t$, and whose backbones are an 8-layer Transformer with a hidden dimension of 256. The models are trained with $1e5$ steps, a batch size of 32, a learning rate of $2e-4$, and the cosine annealing learning rate scheduler. As for the critic model, it is trained with the same backbone model using the last token as the value output and trained $1e4$ steps, batch size of 64, and a learning rate of $2e-4$. The value discount factor is set to 0.99 for all 3 datasets. We use a planning horizon of 20 steps and the planning stride 1. We exclude tricks such as using the inverse dynamics model, planning with stride, and using sample-and-select methods (Lu et al., 2025). During evaluation, the same base model is utilized for different guidance methods to ensure a fair comparison. We report the normalized score (Fu et al., 2020) where 100 is the expert RL agent’s return.

For different guidance methods, different hyperparameter combinations are tuned. We elaborate on them here.

- $g^{\text{cov-A}}$: We tune λ_t in {constant, cosine decay, exponential decay, linear decay} with a scaler {0.01, 0.1, 1.0, 10.0}, where the schedule functions are normalized to [0, 1].
- $g^{\text{cov-G}}$: The same hyperparameters are tuned as in $g^{\text{cov-A}}$.
- g^{MC} : We tune the scale of J in {0.2, 1, 2, 3, 5}. The Monte Carlo sample number is limited to be smaller than {128}. We also include a small number ϵ to enhance numerical stability. We conduct an ablation study to show the performance is insensitive to ϵ , as shown in Table 6.
- $g^{\text{sim-MC}}$: We tune scale before J in {0.1, 1, 10} and the assumed standard deviation of $p(x_1|x_t)$ in {0.1, 0.5, 1, 10}, and do extra schedule and scale of the estimated guidance with schedule tuned in {linear decay, constant} and the

scale tuned in $\{0.1, 1, 10\}$. It is worth noting that if the objective function J is properly normalized, the scale does not require extensive tuning.

- g_ϕ : Training-based methods have many hyperparameters involving model architecture and the training settings. We switch between different model depth and hidden dimensions, and different training losses. The best results for each loss are provided in Table 4 and 5.

Finally, for each pre-trained VF (CFM or OT-CFM), we try adding different guidance VFs (also CFM or OT-CFM), and then the best result is reported.

Estimation of the ground truth distribution of J . Suppose the unguided model generates $x \sim p(x)$, and $J(x)$ follows the distribution $p_J(J)$. Then, if $x' \sim \frac{1}{Z}p(x)e^{-J(x)}$, then J' follows

$$p'(J) = p'(x)\det\left(\frac{\partial x}{\partial J(x)}\right) = \frac{1}{Z}p(x)e^{-J(x)}\det\left(\frac{\partial x}{\partial J(x)}\right). \tag{128}$$

Since for the original distribution

$$p(J) = p(x)\det\left(\frac{\partial x}{\partial J(x)}\right), \tag{129}$$

so

$$p'(J) = \frac{1}{Z}e^{-J}p(J). \tag{130}$$

Therefore, by sampling from the unguided model and then reweighting the distribution of J , we can compute the ground truth distribution $p'(J)$ for the J under ideal guidance. Note that in the planning experiment, $J = -R$.

It should be noted that although gradient-based guidances $g^{\text{cov-A}}$ and $g^{\text{cov-G}}$ result in distributions where the estimated return R is higher, it does not necessarily mean that their performance is better: first, the goal of the guidance is the gray line, what we assume here is that the methods produce a distribution close to the gray line is better; second, practically speaking, the high return is predicted by the critic model, but gradient methods may produce plans that the critic has not seen during training, which is called distribution shift, thus cheating the critic. On the contrary, the target guided distribution $p(x_1)e^{R(x_1)}/Z$ regularizes the guided distribution on the support of $p(x_1)$, alleviating the problem of distribution shift.

Additional Results on the Distribution of Generated R . The additional results of the distribution of R in different environments with different guidance scales (the α in $p'(R) = p(R)e^{-\alpha R}/Z$) are shown in Figure 5.

Additional Results. The complete results, including standard deviations, are provided in Table 4 and 5.

Table 4: Full experiment results on D4RL Locomotion datasets. The base model is mini-batch optimal transport conditional flow matching. Entries with $\geq 95\%$ score than the best per task are highlighted in bold. Baselines are excluded from the ranking.

		w/o g	$g^{\text{cov-A}}$	$g^{\text{cov-G}}$	$g^{\text{sim-MC}}$	g^{MC}	g_ϕ GM	g_ϕ VGM	g_ϕ RGM	g_ϕ MRGM
Medium-Expert	HalfCheetah	61.9 ± 13.3	64.8 ± 12.7	73.2 ± 9.5	78.1 ± 3.2	86.4 ± 0.8	59.5 ± 18.4	57.7 ± 14.1	57.5 ± 13.1	70.2 ± 18.1
	Hopper	95.2 ± 20.4	101.8 ± 22.2	112.3 ± 1.8	112.3 ± 0.8	112.7 ± 0.9	85.2 ± 23.3	98.1 ± 16.3	90.3 ± 24.1	89.3 ± 18.7
	Walker2d	79.1 ± 35.2	97.3 ± 9.4	107.2 ± 1.4	101.0 ± 10.2	107.5 ± 1.0	87.0 ± 16.7	90.5 ± 10.0	88.0 ± 17.2	91.3 ± 11.7
Medium	HalfCheetah	34.7 ± 9.6	42.2 ± 0.8	42.9 ± 0.9	43.1 ± 1.7	43.1 ± 0.4	42.7 ± 1.4	43.0 ± 1.2	42.7 ± 0.8	43.4 ± 0.9
	Hopper	63.3 ± 4.6	75.1 ± 14.9	89.8 ± 13.6	76.2 ± 13.2	79.8 ± 14.8	79.7 ± 12.4	71.6 ± 9.2	77.4 ± 6.6	72.5 ± 6.0
	Walker2d	72.4 ± 13.3	82.7 ± 5.3	81.3 ± 2.0	83.4 ± 1.9	83.0 ± 3.4	80.6 ± 2.2	80.2 ± 2.0	78.4 ± 3.9	76.6 ± 6.1
Medium-Replay	HalfCheetah	25.6 ± 13.0	31.7 ± 3.4	36.1 ± 5.1	36.8 ± 1.8	40.0 ± 1.6	33.4 ± 2.6	35.5 ± 1.8	32.9 ± 1.8	34.7 ± 3.1
	Hopper	40.1 ± 3.7	57.7 ± 15.4	74.1 ± 5.1	60.9 ± 13.1	88.6 ± 11.6	54.6 ± 14.8	48.1 ± 15.2	46.6 ± 11.9	55.3 ± 19.4
	Walker2d	31.2 ± 6.0	62.5 ± 16.8	82.5 ± 10.8	64.4 ± 9.7	88.1 ± 2.1	45.6 ± 17.2	37.8 ± 14.5	44.3 ± 23.1	52.4 ± 20.6

B.3. Image Experiment Details

We pre-trained a CFM and mini-batch optimal transport CFM model with affine path $\alpha_t = t, \beta_t = 1 - t$ on CelebA-HQ 256×256 dataset. The flow matching model utilizes the backbone of a U-Net following (Pokle et al., 2024). The pretraining was conducted with a learning rate of $1e-4$ and a batch size of 128 for 500 epochs. The run time was roughly 3 days on two

On the Guidance of Flow Matching

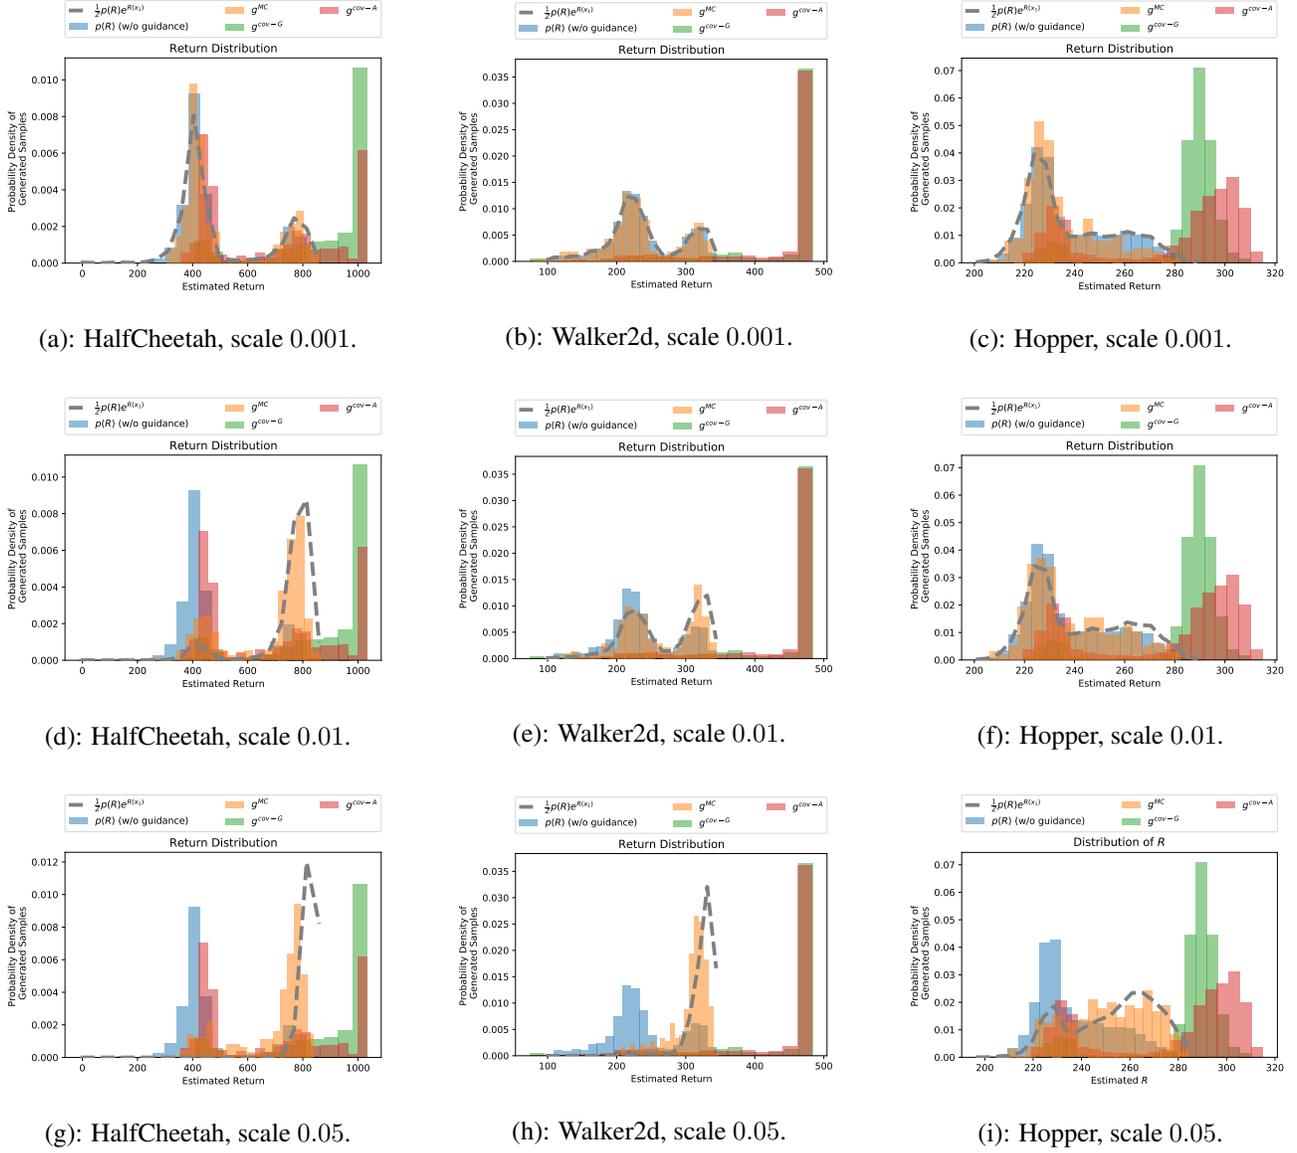


Figure 5: The complete results of the distribution of R . The distribution of J under g^{MC} matches the ground truth value (gray dashed line) very well.

Table 5: Full experiment results on D4RL Locomotion datasets. The base model is conditional flow matching. Entries with $\geq 95\%$ score than the best per task are highlighted in bold. Baselines are excluded from the ranking.

		w/o g	$g^{\text{cov-A}}$	$g^{\text{cov-G}}$	$g^{\text{sim-MC}}$	g^{MC}	g_ϕ GM	g_ϕ VGM	g_ϕ RGM	g_ϕ MRGM
Medium-Expert	HalfCheetah	46.4 \pm 10.1	63.4 \pm 17.8	68.5 \pm 6.2	83.5 \pm 4.2	87.7 \pm 1.8	61.2 \pm 17.1	81.5 \pm 7.7	64.2 \pm 20.7	66.4 \pm 19.2
	Hopper	83.4 \pm 19.2	93.9 \pm 22.5	113.3 \pm 1.8	88.5 \pm 28.0	113.3 \pm 0.9	91.4 \pm 16.1	84.2 \pm 22.2	86.2 \pm 18.7	86.5 \pm 22.4
	Walker2d	65.7 \pm 12.1	100.4 \pm 10.7	106.9 \pm 0.8	107.0 \pm 0.8	107.1 \pm 0.4	96.2 \pm 13.7	102.3 \pm 9.2	100.9 \pm 9.4	98.2 \pm 12.3
Medium	HalfCheetah	41.8 \pm 0.9	43.6 \pm 0.6	43.3 \pm 1.2	43.8 \pm 1.0	43.8 \pm 1.0	43.7 \pm 0.7	43.7 \pm 1.0	42.9 \pm 0.7	43.8 \pm 1.1
	Hopper	73.2 \pm 5.4	79.1 \pm 9.8	82.7 \pm 11.9	82.1 \pm 9.2	88.0 \pm 11.3	81.4 \pm 14.0	85.2 \pm 15.5	74.0 \pm 16.3	71.6 \pm 14.7
	Walker2d	72.2 \pm 5.9	80.7 \pm 1.1	82.5 \pm 2.8	81.9 \pm 5.4	81.9 \pm 1.4	67.5 \pm 23.7	72.9 \pm 2.1	63.3 \pm 21.4	55.9 \pm 28.6
Medium-Replay	HalfCheetah	22.2 \pm 14.9	33.4 \pm 7.1	39.3 \pm 2.0	37.9 \pm 1.4	40.6 \pm 2.0	36.9 \pm 2.3	39.1 \pm 1.7	37.3 \pm 2.1	36.8 \pm 4.4
	Hopper	55.1 \pm 17.2	63.0 \pm 17.5	69.3 \pm 16.4	61.0 \pm 22.6	80.9 \pm 15.7	51.4 \pm 22.6	63.5 \pm 12.7	48.6 \pm 10.9	57.3 \pm 24.0
	Walker2d	28.3 \pm 7.3	64.9 \pm 20.2	76.6 \pm 10.7	58.9 \pm 18.9	70.9 \pm 21.7	57.5 \pm 20.6	70.3 \pm 4.9	53.1 \pm 22.0	54.6 \pm 18.7

Table 6: Ablation study of the impact of the epsilon to the performance of g^{MC} . The best results and the second best per task are highlighted in bold and underlined.

	1	$1e^{-3}$	$5e^{-3}$	$1e^{-2}$	$5e^{-2}$
HalfCheetah					
Medium	42.5 ± 1.6	43.1 ± 0.6	39.8 ± 8.4	41.7 ± 2.9	43.2 ± 1.5
Medium-Expert	68.2 ± 15.5	66.9 ± 17.1	75.6 ± 12.6	74.6 ± 12.1	72.7 ± 16.6
Medium-Replay	33.5 ± 9.9	31.7 ± 12.6	39.7 ± 2.0	34.7 ± 10.9	37.3 ± 6.3
Hopper					
Medium	69.0 ± 10.6	67.6 ± 4.9	73.1 ± 9.4	72.2 ± 10.0	72.6 ± 12.3
Medium-Expert	95.4 ± 19.5	103.2 ± 19.2	108.0 ± 12.3	107.1 ± 14.3	101.5 ± 20.8
Medium-Replay	53.8 ± 18.7	59.1 ± 17.0	68.2 ± 18.7	68.1 ± 18.1	74.3 ± 18.8
Walker2d					
Medium	74.6 ± 6.5	75.1 ± 15.2	74.0 ± 10.6	71.1 ± 13.8	74.4 ± 11.0
Medium-Expert	79.0 ± 26.4	95.7 ± 17.8	94.2 ± 22.0	100.9 ± 17.6	103.0 ± 7.9
Medium-Replay	48.7 ± 21.5	49.5 ± 16.7	54.7 ± 20.6	57.6 ± 22.3	56.7 ± 18.6

H800 GPUs. For the CelebA dataset, we employed a train-validation-test split of 8:1:1. The test data was subsequently used for three downstream tasks: central box inpainting, superresolution by four times, and Gaussian deblurring which is all common benchmarks.

Settings for the experiments. We evaluated the guidance methods using 3,000 images randomly sampled from the test set across the three inverse problems. Specifically, for deblurring, we apply a 61×61 Gaussian kernel with a standard deviation of $\sigma_b = 1.0$. For super-resolution, we perform $4 \times$ downsampling on the CelebA images. In the case of box-inpainting, we use a centered 40×40 mask. Furthermore, for all three tasks, we add Gaussian noise after the degradation operation with a standard deviation of $\sigma = 0.05$ to the images.

Metrics. In this paper, we use four commonly adopted metrics for image quality assessment: FID (Fréchet Inception Distance), which measures the distance between generated and real image distributions; LPIPS (Learned Perceptual Image Patch Similarity), which evaluates the perceptual similarity between images; and PSNR (Peak Signal-to-Noise Ratio) and SSIM (Structural Similarity Index) to quantify image quality in terms of signal preservation and structural consistency, respectively.

Why is g^{MC} bad at image inverse problems? As can be seen from Figure 6 and Figure 7, the images generated by g^{MC} do not respect the reference degraded image. This is largely due to the variance of the MC estimation being too high given the limited number of samples. Specifically, to estimate g_t , one needs to obtain samples from regions where e^{-J} is significantly higher than average, which corresponds to the images that already look like the degraded image. Sampling such images requires an infeasibly large number of samples. More advanced MC sampling techniques may help address this shortcoming, such as the control variable method (Owen, 2013). Combining g^{MC} and methods that are biased but with lower variance, such as g^{local} or g^{sim} , may also boost the performance.

However, on tasks such as conditional generation, as long as the condition often appears in the dataset, it will be easier to obtain an accurate estimation of g_t using g_t^{MC} . Such scenarios include property-conditioned molecular structure generation, label-conditioned image generation, and decision-making tasks, which are included in our experiments.

Visualizations. We provide visualizations of the results of the inverse problems in Figure 6 and 7.

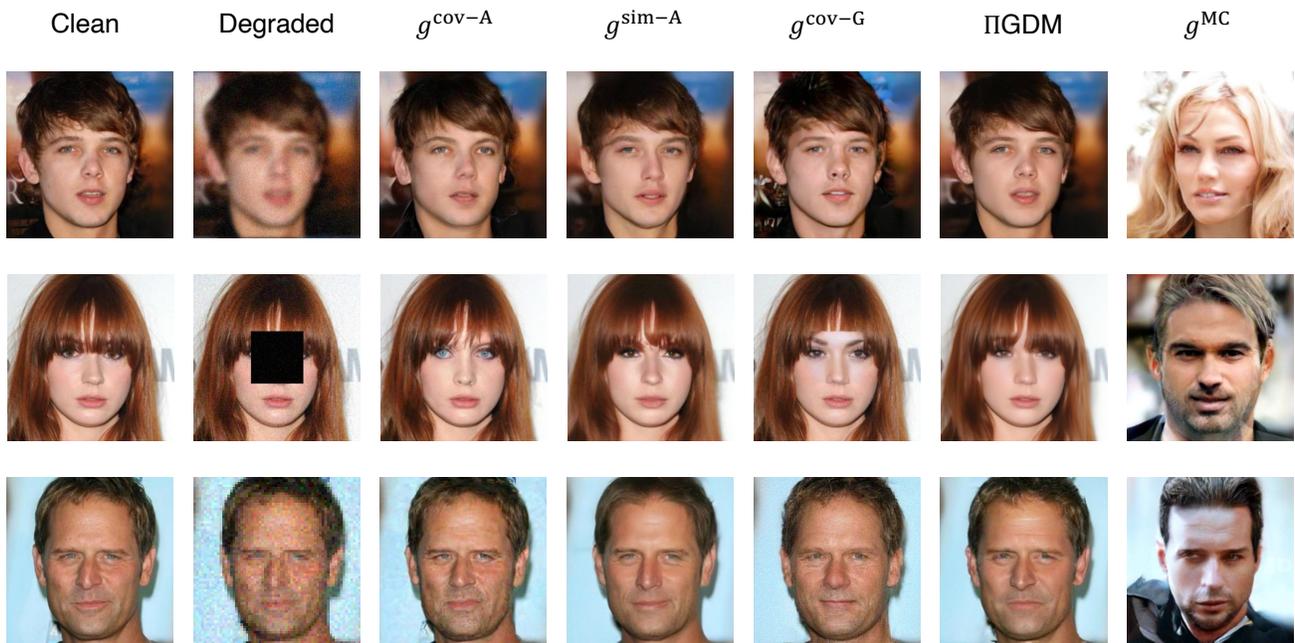


Figure 6: The visualization of the image inverse problems with the base flow matching model of mini-batch optimal transport conditional flow matching (OT-CFM). Three rows show the results of Gaussian deblurring, box-inpainting, and super-resolution from top to bottom.

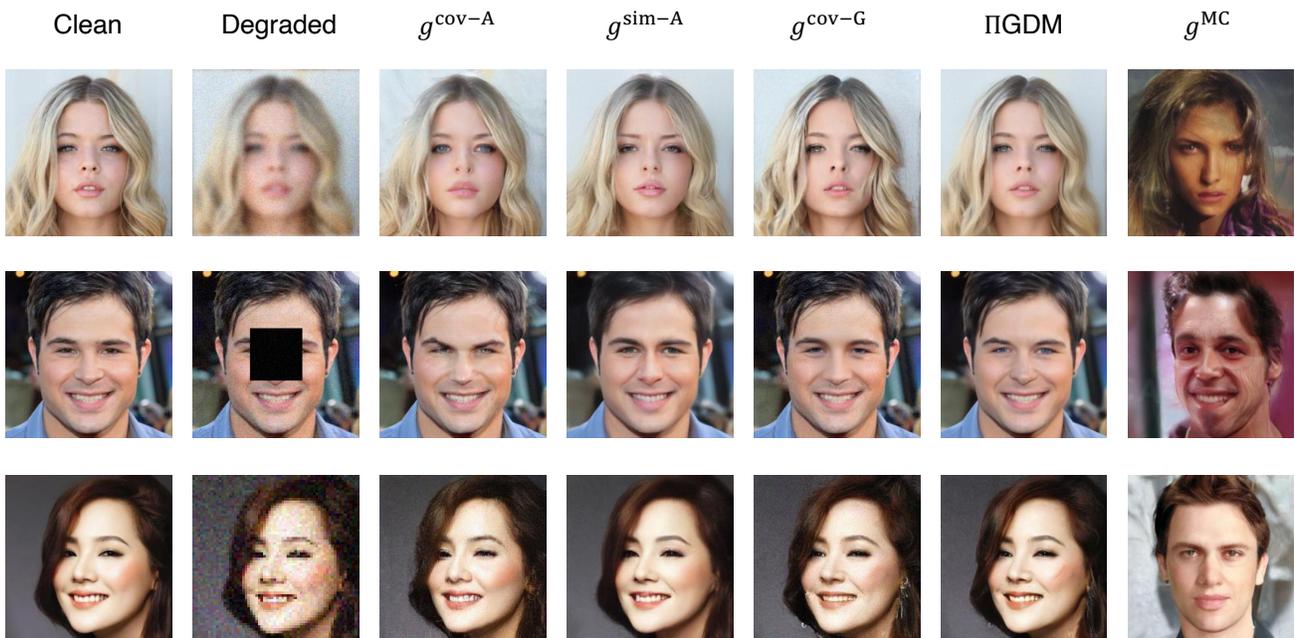


Figure 7: The visualization of the image inverse problems with the base flow matching model of conditional flow matching (CFM). Three rows show the results of Gaussian deblurring, box-inpainting, and super-resolution from top to bottom.