ENERGY-WEIGHTED FLOW MATCHING FOR OFFLINE REINFORCEMENT LEARNING

Shiyuan Zhang[1] **Weitong Zhang**[2] **Quanquan Gu**[3]
[1]Tsinghua University [2]UNC-Chapel Hill [3]University of California, Los Angeles
shiyuan-21@mails.tsinghua.edu.cn, weitongz@unc.edu, qgu@cs.ucla.edu

ABSTRACT

This paper investigates energy guidance in generative modeling, where the target distribution is defined as $q(\mathbf{x}) \propto p(\mathbf{x})\exp(-\beta\mathcal{E}(\mathbf{x}))$, with $p(\mathbf{x})$ being the data distribution and $\mathcal{E}(\mathbf{x})$ as the energy function. To comply with energy guidance, existing methods often require auxiliary procedures to learn intermediate guidance during the diffusion process. To overcome this limitation, we explore energy-guided flow matching, a generalized form of the diffusion process. We introduce energy-weighted flow matching (EFM), a method that directly learns the energy-guided flow without the need for auxiliary models. Theoretical analysis shows that energy-weighted flow matching accurately captures the guided flow. Additionally, we extend this methodology to energy-weighted diffusion models and apply it to offline reinforcement learning (RL) by proposing the Q-weighted Iterative Policy Optimization (QIPO). Empirically, we demonstrate that the proposed QIPO algorithm improves performance in offline RL tasks. Notably, our algorithm is the first energy-guided diffusion model that operates independently of auxiliary models and the first exact energy-guided flow matching model in the literature.

1 INTRODUCTION

Recent years have witnessed the success of applying diffusion models (Ho et al., 2020; Song et al., 2020) and flow matching models (Chen et al., 2018; Lipman et al., 2022) to generative models. Given this success, another important aspect is to *guide* generative models to achieve specific, controlled outputs, such as generating images for a certain class (Ho & Salimans, 2021; Dhariwal & Nichol, 2021), designing molecular structures with desired properties (Wang et al., 2024; Hoogeboom et al., 2022), or improving policies for reinforcement learning (Wang et al., 2022; Lu et al., 2023). Guidance can come from various sources, such as classifiers, including both classifier guidance (Dhariwal & Nichol, 2021) and classifier-free guidance (Ho & Salimans, 2021). In addition, Lu et al. (2023) proposed using guidance from an *energy function*, where the distribution is generated from $q(\mathbf{x}) \propto p(\mathbf{x})\exp(-\beta\mathcal{E}(\mathbf{x}))$, where the model is guided to generate data $\mathbf{x}$ with lower energy $\mathcal{E}(\mathbf{x})$ from the original data distribution.

Several recent efforts have been made to learn and sample from the guided distribution $q(\mathbf{x})$ using diffusion models. Chen et al. (2022) performed rejection sampling from the learned data distribution $p(\mathbf{x})$. Lu et al. (2023) introduced an *intermediate energy function* $\mathcal{E}_t(\cdot)$, allowing the score function of $q_t(\mathbf{x})$ to be decomposed as $\nabla_{\mathbf{x}} \log q_t(\mathbf{x}) = \nabla_{\mathbf{x}} \log p_t(\mathbf{x}) - \nabla_{\mathbf{x}}\mathcal{E}_t(\mathbf{x})$ within the diffusion process. Lu et al. (2023) also proposed contrastive energy prediction for training the intermediate energy $\mathcal{E}_t$, relying on back-propagation to calculate its gradient with respect to $\mathbf{x}$. Wang et al. (2024) proposed to directly approximate the gradient of this intermediate energy function $\mathcal{E}_t$ as the 'force-field' guidance. However, all these methods require either additional neural networks, back-propagation, or post-processing to compose the guided distribution $q(\mathbf{x})$, which introduces unnecessary errors and complexity. Therefore, the following question arises:

Q1. *Can we directly obtain an energy-guided diffusion model without auxiliary models?*

Another challenge for energy-guided generative models lies in providing guidance in flow matching models (Chen et al., 2018; Lipman et al., 2022), which is a more general, simulation-free counterpart to diffusion models. Zheng et al. (2023) explored the use of classifier-free guidance for flow matching in offline RL. However, since flow matching models approximate the velocity field $\mathbf{u}_t(\mathbf{x})$ for the dynamics of the probability density path $p_t(\mathbf{x})$, it is highly non-trivial to obtain the guided velocity field $\widehat{\mathbf{u}}_t(\mathbf{x})$ for the distribution $q_t(\mathbf{x})$ under energy guidance. This presents the second key question:

Q2. *Can we inject exact energy guidance into general flow matching models?*

In this paper, we answer the aforementioned two questions affirmatively by proposing an energy-guided velocity field and an *energy-weighted* flow matching objective, with extensions to *energy-weighted* diffusion models and applications in offline reinforcement learning. Our contributions are summarized as follows:

- In response to **Q2.**, for general flow matching, we propose the energy-guided velocity field $\widehat{\mathbf{u}}_t(\mathbf{x})$, based on the conditional velocity field $\mathbf{u}_{t0}(\mathbf{x}|\mathbf{x}_0)$. The proposed $\widehat{\mathbf{u}}_t(\mathbf{x})$ is theoretically guaranteed to generate the energy-guided distribution $q(\mathbf{x}) \propto p(\mathbf{x}) \exp(-\beta\mathcal{E}(\mathbf{x}))$.
- We introduce the *energy-weighted* flow matching loss to train a neural network $\mathbf{v}_t^{\boldsymbol{\theta}}$ that approximates $\widehat{\mathbf{u}}_t(\mathbf{x})$. The energy-weighted flow matching only requires the conditional vector field $\mathbf{u}_t(\mathbf{x}|\mathbf{x}_0)$ and the energy $\mathcal{E}(\mathbf{x}_0)$ for $\mathbf{x}_0$ from the dataset. As the answer to **Q1.**, we extend this approach to diffusion models, proposing the energy-weighted diffusion model. Energy-weighted diffusion model learns an energy-guided diffusion model directly without any auxiliary model.
- We apply these methods to offline reinforcement learning tasks to evaluate the performance of the energy-weighted flow matching and diffusion models. Under this framework, we introduce an *iterative policy refinement* technique for offline reinforcement learning. Empirically, we demonstrate that the proposed method achieves superior performance across various offline RL tasks.

Notations. Vectors are denoted by lowercase boldface letters, such as $\mathbf{x}$, and matrices by uppercase boldface letters, such as $\mathbf{A}$. For any positive integer k, the set $1, 2, \ldots, k$ is denoted by $[k]$, and we define $\overline{[k]} = [k] \cup \{0\}$. The natural logarithm of x is denoted by $\log x$. We use p_t to represent the marginal distribution of $\mathbf{x}$ at time t, and p_{0t} to represent the conditional distribution of $\mathbf{x}_0$ given $\mathbf{x}_t$. Similarly, p_0 denotes the original data distribution for the diffusion model at $t = 0$, while p_{t0} represents the conditional distribution of $\mathbf{x}_t$ given $\mathbf{x}_0$ in the forward process of the diffusion model.

2 Related work

Diffusion Models and Flow Matching Models. Diffusion models (Ho et al., 2020) and score matching (Song et al., 2020) have emerged as powerful generative modeling techniques in tasks such as image synthesis (Dhariwal & Nichol, 2021), text-to-image generation (Podell et al., 2023), and video generation (Ho et al., 2022). In addition to these frontier applications, the success of diffusion models has been further enhanced by accelerated sampling processes (Lu et al., 2022a;b) and the extension of diffusion models to discrete value spaces (Austin et al., 2021). Alongside the success of diffusion models, flow matching models (Lipman et al., 2022; Chen et al., 2018) provide an alternative for simulation-free generation. Unlike score-based approaches, flow matching models aim to learn the velocity field that transports data points from the initial noise distribution to the target data distribution. This velocity field can be viewed as a generalized form of the reverse process in diffusion models and can be extended to optimal transport (Villani et al., 2009), rectified flow (Liu, 2022), and more complex flows.

Guidance in Diffusion and Flow Matching Models. Beyond learning and generating the original data distribution with diffusion or flow matching models, significant efforts have been made to control the generation process to produce data with specific desired properties. Dhariwal & Nichol (2021) introduced classifier guidance, which decomposes the conditional score function $\nabla \log p(\mathbf{x}|y)$ into the sum of the data distribution gradient $\nabla \log p(\mathbf{x})$ and the gradient from a classifier $\nabla \log p(y|\mathbf{x})$. To simplify this, Ho & Salimans (2021) proposed classifier-free guidance, which directly integrates $\nabla \log p(y|\mathbf{x})$ into the score function. Sendera et al. (2024) studied diffusion-structured samplers by introducing the inductive bias in Langevin process. Lu et al. (2023); Chen et al. (2022) further explored energy-based guidance, where the target distribution is defined as $q(\mathbf{x}) \propto p(\mathbf{x}) \exp(-\beta\mathcal{E}(\mathbf{x}))$. Unlike classifier guidance, energy-based guidance extends to real-valued energy functions $\mathcal{E}$, making it particularly relevant for tasks such as molecular structure generation. Specifically, Chen et al. (2022); Cremer et al. (2024) employed rejection sampling to implement energy guidance, while Lu et al. (2023); Wang et al. (2024) used auxiliary models to esti-

Table 1: Comparison between guidance methods. *Exact Guidance?* means if the model can generate $p(\mathbf{x})p^{\beta}(c|\mathbf{x})$ when $\beta \neq 1$. *w/o Auxiliary Model?* means if the method can direct learn the guidance without auxiliary model (✓) or not (×).

Guidance	Exact Guidance?	w/o Auxiliary Model?
Classifier-guidance (Dhariwal & Nichol, 2021)	×	×
Classifier-free guidance (Ho & Salimans, 2021)	×	✓
Contrastive energy prediction (Lu et al., 2023)	✓	×
Energy-weighted diffusion (ours)	✓	✓

mate the guidance from the energy function. We defer a more formal, technical comparison between the energy-based guidance and classifier-based guidance in Table 1 in Section 4.3. In the context of flow matching, Zheng et al. (2023) introduced classifier-free guidance for flow matching in the domain of offline reinforcement learning.

Diffusion and Flow Matching Models in Reinforcement Learning. Recent advances in diffusion models and flow matching models have enabled a range of applications in reinforcement learning (RL). Janner et al. (2022); Wang et al. (2022) explore modeling behavior policies using diffusion models. Building on these results, Chen et al. (2022); Lu et al. (2023) model the offline RL objective as an energy-guided diffusion process, while Ajay et al. (2022); Zheng et al. (2023) apply the same policy optimization using classifier-free diffusion and flow matching models. Chen et al. (2023); Hansen-Estruch et al. (2023) use diffusion models to regularize the distance between the optimal policy and the behavioral policy, and Fang & Lan (2024); He et al. (2023) leverage diffusion models for constrained policy optimization. Another line of research (Jackson et al., 2024; Lee et al., 2024; Lu et al., 2024) focuses on using generative models to augment synthetic datasets.

3 Preliminaries

3.1 Conditional Flow Matching for Generative Modeling

Continuous Normalizing Flows (CNFs) (Chen et al., 2018) considers the dynamic of the probability density function by *probability density path* $p : [0, 1] \times \mathbb{R}^d \mapsto \mathbb{R}_{\geq 0}$ which transmits between the data distribution p_0 and the initial distribution (e.g., Gaussian distribution) p_1. The *flow* $\boldsymbol{\phi} : [0, 1] \times \mathbb{R}^d \mapsto \mathbb{R}^d$ is constructed by a *vector field* $\mathbf{v} : [0, 1] \times \mathbb{R}^d \mapsto \mathbb{R}^d$ describing the velocity of the particle at position $\mathbf{x}$, i.e., $\frac{\mathrm{d}}{\mathrm{d}t}\boldsymbol{\phi}_t(\mathbf{x}) = \mathbf{v}_t(\boldsymbol{\phi}_t(\mathbf{x}))$ where $\boldsymbol{\phi}_1(\mathbf{x}) = \mathbf{x}$. [1]

In order to ensure that the vector field $\mathbf{v}$ *generates* the probability density path p_t, the following *continuity equation* (Villani et al., 2009) is required:

$$\frac{\mathrm{d}}{\mathrm{d}t}p_t(\mathbf{x}) + \mathrm{div} \cdot [p_t(\mathbf{x})\mathbf{v}_t(\mathbf{x})] = 0, \quad \forall \mathbf{x} \in \mathbb{R}^d. \tag{3.1}$$

The objective of flow matching is to learn a neural network $\mathbf{v}_t^{\boldsymbol{\theta}}$ to learn the ground truth vector field $\mathbf{u}_t$ by minimizing their differences, i.e., $\mathcal{L}_{\mathrm{FM}}(\boldsymbol{\theta}) = \mathbb{E}_{t,p_t(\mathbf{x})}\|\mathbf{v}_t^{\boldsymbol{\theta}}(\mathbf{x}) - \mathbf{u}_t(\mathbf{x})\|_2^2$ with respect to the network parameter $\boldsymbol{\theta}$. However, it is infeasible to calculate the ground truth vector field $\mathbf{u}_t$. To address this issue, Lipman et al. (2022) suggests to match the *conditional vector field* $\mathbf{u}_{t0}(\mathbf{x}|\mathbf{x}_0)$ instead of the vector field $\mathbf{u}_t(\mathbf{x})$, as presented by the following theorem:

Theorem 3.1 (Theorem 1, 2; Lipman et al. 2022)**.** Given the conditional vector field $\mathbf{u}_{t0}(\mathbf{x}|\mathbf{x}_0)$ that generates the conditional distribution $p_{t0}(\mathbf{x}|\mathbf{x}_0)$, then the "marginal" vector field $\mathbf{u}_t(\mathbf{x}) = \int_{\mathbf{x}_0} p_{0t}(\mathbf{x}_0|\mathbf{x})\mathbf{u}_{t0}(\mathbf{x}|\mathbf{x}_0)\mathrm{d}\mathbf{x}_0$ generates the marginal distribution $p_t(\mathbf{x})$. In addition, up to a constant factor independent of $\boldsymbol{\theta}$, *Flow Matching* loss $\mathcal{L}_{\mathrm{FM}}(\boldsymbol{\theta})$ and *Conditional Flow Matching* loss $\mathcal{L}_{\mathrm{CFM}}(\boldsymbol{\theta})$ are equal, where

$$\mathcal{L}_{\mathrm{FM}}(\boldsymbol{\theta}) = \mathbb{E}_{t,\mathbf{x}}\|\mathbf{v}_t^{\boldsymbol{\theta}}(\mathbf{x}) - \mathbf{u}_t(\mathbf{x})\|_2^2, \ \mathcal{L}_{\mathrm{CFM}}(\boldsymbol{\theta}) = \mathbb{E}_{t,\mathbf{x}_0,\mathbf{x}}\|\mathbf{v}_t^{\boldsymbol{\theta}}(\mathbf{x}) - \mathbf{u}_{t0}(\mathbf{x}|\mathbf{x}_0)\|_2^2, \tag{3.2}$$

where $t \sim \lambda(t)$, $\mathbf{x}_0$ follows the data distribution $p_0(\cdot)$ and $\mathbf{x} \sim p_{t0}(\cdot|\mathbf{x}_0)$ where p_{t0} is generated by conditional vector field $\mathbf{u}_{t0}$. Hence $\nabla_{\boldsymbol{\theta}}\mathcal{L}_{\mathrm{FM}}(\boldsymbol{\theta}) = \nabla_{\boldsymbol{\theta}}\mathcal{L}_{\mathrm{CFM}}(\boldsymbol{\theta})$.

[1] We adapt the notation of diffusion to unify the diffusion and flow matching. The notation here is different from flow matching notations in Chen et al. (2018); Lipman et al. (2022), where p_1 represent the data distribution. The flow matching result remains unchanged except switching $t = 1$ with $t = 0$.

In practice, the conditional distribution $p_{t0}(\mathbf{x}|\mathbf{x}_0)$ is usually modeled as a *Gaussian path* with $p_{t0}(\mathbf{x}|\mathbf{x}_0) = \mathcal{N}(\mu_t \mathbf{x}_0, \sigma_t^2 \mathbf{I})$. Zheng et al. (2023) suggests that in this case, conditional flow matching is equivalent to the score matching (Song et al., 2020):

Lemma 3.2 (Lemma 1, Zheng et al. 2023)**.** Let $p_{t0}(\mathbf{x}|\mathbf{x}_0)$ be a Gaussian path with scheduler (μ_t, σ_t), i.e., $p_{t0}(\mathbf{x}|\mathbf{x}_0) = \mathcal{N}(\mu_t \mathbf{x}_0, \sigma_t^2 \mathbf{I})$, then the velocity field $\mathbf{u}_{t0}(\mathbf{x}|\mathbf{x}_0)$ is related to the score function $\nabla_{\mathbf{x}} \log p_{t0}(\mathbf{x}|\mathbf{x}_0)$ by

$$\mathbf{u}_{t0}(\mathbf{x}|\mathbf{x}_1) = \dot{\mu}_t \mu_t^{-1} \mathbf{x} + (\dot{\mu}_t \sigma_t - \mu_t \dot{\sigma}_t) \sigma_t \mu_t^{-1} \nabla_{\mathbf{x}} \log p_{t0}(\mathbf{x}|\mathbf{x}_0), \tag{3.3}$$

where $\dot{\mu}_t$ and $\dot{\sigma}_t$ are both the derivative of μ_t and σ_t with respect to time t.

In addition, Zheng et al. (2023) proved that the reverse process of this diffusion process with Gaussian path μ_t, σ_t can be written by

$$\frac{\mathrm{d}\mathbf{x}}{\mathrm{d}t} = \frac{\dot{\mu}_t}{\mu_t} \mathbf{x} + (\dot{\mu}_t \sigma_t - \mu_t \dot{\sigma}_t) \frac{\sigma_t}{\mu_t} \nabla_{\mathbf{x}} \log p_t(\mathbf{x}) = \mathbf{u}_t(\mathbf{x}). \tag{3.4}$$

3.2 Energy-Guided Diffusion Models

The standard diffusion model aims to learn and generate from the data distribution p_0. However, instead of generating from p_0, there are a series of applications consider sampling from an energy-guided distribution $q_0(\mathbf{x}) \propto p_0(\mathbf{x}) \exp(-\beta \mathcal{E}(\mathbf{x}))$ where $\mathcal{E} : \mathbb{R}^d \mapsto \mathbb{R}$ is the energy function and $\beta \in \mathbb{R}^+$ is the strength of the guidance. Lu et al. (2023) suggested to construct the score function $\nabla_{\mathbf{x}} \log q_t(\mathbf{x})$ from the original score function $\nabla_{\mathbf{x}} \log p_t(\mathbf{x})$ by introducing the *intermediate energy function* $\mathcal{E}_t(\mathbf{x})$ through the following theorem:

Theorem 3.3 (Theorem 3.1, Lu et al. (2023)[2])**.** Let $q_0(\mathbf{x}) \propto p_0(\mathbf{x}) \exp(-\beta \mathcal{E}(\mathbf{x}))$ and define the forward process as $q_{t0}(\mathbf{x}|\mathbf{x}_0) = p_{t0}(\mathbf{x}|\mathbf{x}_0) = \mathcal{N}(\mu_t \mathbf{x}_0, \sigma_t^2 \mathbf{I})$, and the marginal distribution $q_t(\mathbf{x}), p_t(\mathbf{x})$ at time t defined by

$$q_t(\mathbf{x}) = \int_{\mathbf{x}_0} q_{t0}(\mathbf{x}|\mathbf{x}_0) q_0(\mathbf{x}_0) \mathrm{d}\mathbf{x}_0, \qquad p_t(\mathbf{x}) = \int_{\mathbf{x}_0} p_{t0}(\mathbf{x}|\mathbf{x}_0) p_0(\mathbf{x}_0) \mathrm{d}\mathbf{x}_0.$$

Let the intermediate energy function be

$$\mathcal{E}_t(\mathbf{x}) = -\log \mathbb{E}_{p_{0t}(\mathbf{x}_0|\mathbf{x})} [\exp(-\beta \mathcal{E}(\mathbf{x}_0)], \tag{3.5}$$

then the marginal distribution p_t and q_t satisfy

$$q_t(\mathbf{x}) \propto p_t(\mathbf{x}) \exp(-\mathcal{E}_t(\mathbf{x})), \qquad \nabla_{\mathbf{x}} \log q_t(\mathbf{x}) = \nabla_{\mathbf{x}} \log p_t(\mathbf{x}) - \nabla_{\mathbf{x}} \mathcal{E}_t(\mathbf{x}). \tag{3.6}$$

Therefore, Lu et al. (2023) suggests to firstly learn the intermediate energy function $\mathcal{E}_t$ using *contrastive energy prediction* (CEP) and to learn the score function $\nabla_{\mathbf{x}} \log p_t(\mathbf{x})$ using standard diffusion models (e.g., DDPM (Ho et al., 2020)). Then the score function of the energy-guided distribution $\nabla_{\mathbf{x}} \log q_t(\mathbf{x})$ can therefore be composed according to (3.6).

4 Methodology

In this section, we propose a *energy-weighted* method for training both CNFs and diffusion models to generate the energy-guided distribution $q(\mathbf{x}) \propto p(\mathbf{x}) \exp(-\beta \mathcal{E}(\mathbf{x}))$. Compared with Lu et al. (2023); Wang et al. (2024), the energy-weighted method provide a more straightforward way to obtain the energy-guided generative models and removes the necessaries of estimating the intermediate energy function $\mathcal{E}_t(\mathbf{x})$ and its gradient $\nabla_{\mathbf{x}} \mathcal{E}_t(\mathbf{x})$.

4.1 Energy-Weighted Flow Matching

In this subsection, we construct a new energy guided flow to generate the energy-guided probability distribution. We also proposed two equivalent loss function to train the neural networks for approximating the energy-guided flow. We start by the first theorem suggesting a energy-guided flow to generate the energy-guided probability distribution $q_t(\mathbf{x}) \propto p_t(\mathbf{x}) \exp(-\mathcal{E}_t(\mathbf{x}))$.

[2] We swap the notation p and q to align with our notation systems.

Theorem 4.1. Given an energy function $\mathcal{E}(\cdot)$ and a conditional flow $\mathbf{u}_{t0}(\mathbf{x}|\mathbf{x}_0)$ that generates the probability distribution $p_{t0}(\mathbf{x}|\mathbf{x}_0)$, the energy guided distribution $q_t(\mathbf{x}) \propto p_t(\mathbf{x})\exp(-\mathcal{E}_t(\mathbf{x}))$ is generated by the flow

$$\widehat{\mathbf{u}}_t(\mathbf{x}) = \int_{\mathbf{x}_0} p_{0t}(\mathbf{x}_0|\mathbf{x})\mathbf{u}_t(\mathbf{x}|\mathbf{x}_0)\frac{\exp(-\beta\mathcal{E}(\mathbf{x}_0))}{\exp(-\mathcal{E}_t(\mathbf{x}))}\mathrm{d}\mathbf{x}_0, \tag{4.1}$$

which will generate distribution $q_0(\mathbf{x}) \propto p_0(\mathbf{x})\exp(-\beta\mathcal{E}(\mathbf{x}))$. The intermediate energy function is defined in (3.5).

Remark 4.2. Theorem 4.1 suggests a method to construct the vector field $\widehat{\mathbf{u}}_t(\mathbf{x})$ from the conditional vector field $\mathbf{u}_{t0}(\mathbf{x}|\mathbf{x}_0)$ and the intermediate energy function $\mathcal{E}_t(\mathbf{x})$ in the closed-form solution. It holds universally to any conditional flow including the the optimal transport, Gaussian path or rectify flow. We will extend the discussion on the diffusion models in the next subsection.

Despite the closed-form expression for the energy-guided flow, it remains challenging to learn a neural network $\mathbf{v}_t^{\boldsymbol{\theta}}$ to match $\widehat{\mathbf{u}}_t$ since the following two reasons. First, $\widehat{\mathbf{u}}_t$ in (4.1) requires to sample over data distribution $\mathbf{x}_0$. Secondly, the expression of $\widehat{\mathbf{u}}_t$ still requires the estimation of the intermediate energy function $\mathcal{E}_t$. Previous methods are both using auxiliary neural networks to approximate either $\mathcal{E}_t$ (Lu et al., 2023) or its gradient $\nabla_{\mathbf{x}}\mathcal{E}_t(\mathbf{x})$ (Wang et al., 2024). To overcome these two challenges, the following theorem suggests a *weighted* flow matching objective which can be directly used to learn $\widehat{\mathbf{u}}_t$ without the aforementioned procedures.

Theorem 4.3. Define the **E**nergy-weighted **F**low **M**atching loss $\mathcal{L}_{\mathrm{EFM}}$ as

$$\mathcal{L}_{\mathrm{EFM}}(\boldsymbol{\theta}) = \mathbb{E}_{t,\mathbf{x}}\left[\frac{\exp(-\mathcal{E}_t(\mathbf{x}))}{\mathbb{E}_{p_t(\widetilde{\mathbf{x}})}[\exp(-\mathcal{E}_t(\widetilde{\mathbf{x}}))]}\|\mathbf{v}_t^{\boldsymbol{\theta}}(\mathbf{x}) - \widehat{\mathbf{u}}_t(\mathbf{x})\|_2^2\right], \tag{4.2}$$

and the **C**onditional **E**nergy-weighted **F**low **M**atching loss $\mathcal{L}_{\mathrm{CEFM}}$ as

$$\mathcal{L}_{\mathrm{CEFM}}(\boldsymbol{\theta}) = \mathbb{E}_{t,\mathbf{x},\mathbf{x}_0}\left[\frac{\exp(-\beta\mathcal{E}(\mathbf{x}_0))}{\mathbb{E}_{p_0(\widetilde{\mathbf{x}_0})}[\exp(-\beta\mathcal{E}(\widetilde{\mathbf{x}_0}))]}\|\mathbf{v}_t^{\boldsymbol{\theta}}(\mathbf{x}) - \mathbf{u}_{t0}(\mathbf{x}|\mathbf{x}_0)\|_2^2\right], \tag{4.3}$$

where the expectation on t is taken over some predefined distribution $\lambda(t)$, $\mathbf{x}_0$ is sampled from the data distribution $p_0(\cdot)$ and $\mathbf{x}$ at time t is sampled by $p_t(\mathbf{x})$ with conditional distribution $p_{t0}(\mathbf{x}|\mathbf{x}_0)$ generated by the flow $\mathbf{u}_{t0}(\mathbf{x}|\mathbf{x}_0)$. $\mathcal{L}_{\mathrm{EFM}}(\boldsymbol{\theta})$ and $\mathcal{L}_{\mathrm{CEFM}}(\boldsymbol{\theta})$ are equal up to a constant factor. Hence $\nabla_{\boldsymbol{\theta}}\mathcal{L}_{\mathrm{EFM}}(\boldsymbol{\theta}) = \nabla_{\boldsymbol{\theta}}\mathcal{L}_{\mathrm{CEFM}}(\boldsymbol{\theta})$.

Theorem 4.3 suggests that minimizing $\mathcal{L}_{\mathrm{CEFM}}$ is equivalent to minimizing $\mathcal{L}_{\mathrm{EFM}}$. It is obvious that the global minimum of $\mathcal{L}_{\mathrm{CEFM}}$ is $\mathbf{v}_t^{\boldsymbol{\theta}}(\mathbf{x}) = \widehat{\mathbf{u}}_t(\mathbf{x})$, given enough neural network capacity and infinite data. Therefore, one can use $\mathcal{L}_{\mathrm{CEFM}}$ to directly learn the guided flow $\widehat{\mathbf{u}}_t(\mathbf{x})$, without calculating the intermediate energy function $\mathcal{E}_t(\mathbf{x})$ or its gradient.

Besides the aforementioned message, Theorem 4.3 suggests several understandings and intuitions in training the neural network $\mathbf{v}_t^{\boldsymbol{\theta}}$ which are discussed as follows

Remark 4.4 (Regarding the weighted energy guided loss $\mathcal{L}_{\mathrm{EFM}}$)**.** Instead of directly minimizing $\mathbb{E}_{t,\mathbf{x}}\|\mathbf{v}_t^{\boldsymbol{\theta}}(\mathbf{x}) - \widehat{\mathbf{u}}_t(\mathbf{x})\|_2^2$, $\mathcal{L}_{\mathrm{EFM}}$ places higher weight on the input $\mathbf{x}$ with a lower intermediate energy $\mathcal{E}_t(\mathbf{x})$. Intuitively speaking, $\exp(-\mathcal{E}(\mathbf{x}))$ can be viewed as a prior distribution in generating $q_t(\mathbf{x}) \propto p_t(\mathbf{x})\exp(-\mathcal{E}_t(\mathbf{x}))$. Therefore, for all time t, areas with higher $\exp(-\mathcal{E}(\mathbf{x}))$ will be more likely to be visited. As a result, it would be more efficient placing more importance on $\mathbf{x}$ in these areas instead of learning $\widehat{\mathbf{u}}_t(\mathbf{x})$ uniformly for all $\mathbf{x} \in \mathbb{R}^d$.

Remark 4.5 (Regarding the conditional weighted energy guided loss $\mathcal{L}_{\mathrm{CEFM}}$)**.** The weight $\exp(-\beta\mathcal{E}(\mathbf{x}_0))$ suggests how the energy "guides" the conditional flow matching. Fixing t and $\mathbf{x}$ and changing the form of expectations in (4.3), $\mathcal{L}_{\mathrm{CEFM}}(\boldsymbol{\theta})$ becomes

$$\mathcal{L}_{\mathrm{CEFM}}(\boldsymbol{\theta}; t, \mathbf{x}) = \mathbb{E}_{p_{0t}(\mathbf{x}_0|\mathbf{x})}\left[\frac{\exp(-\beta\mathcal{E}(\mathbf{x}_0))}{\mathbb{E}_{p_0(\widetilde{\mathbf{x}_0})}[\exp(-\beta\mathcal{E}(\widetilde{\mathbf{x}_0}))]}\|\mathbf{v}_t^{\boldsymbol{\theta}}(\mathbf{x}) - \mathbf{u}_{t0}(\mathbf{x}|\mathbf{x}_0)\|_2^2\right].$$

Intuitively speaking, velocity field $\mathbf{u}_{t0}(\mathbf{x}|\mathbf{x}_0)$ will move the particle $\mathbf{x}$ to $\mathbf{x}_0$. Therefore, when the energy guidance does not exist (i.e, $\mathcal{E}(\mathbf{x}) = 0$), $\mathbf{v}_t^{\boldsymbol{\theta}}(\mathbf{x})$ is essentially finding the "center" of all $\mathbf{x}_0$ possibly generated from $\mathbf{x}$ following $p(\mathbf{x}_0|\mathbf{x})$. In the presence of the energy function $\mathcal{E}(\mathbf{x}_0)$, the learnt vector field $\mathbf{v}_t^{\boldsymbol{\theta}}(\mathbf{x})$ is biased to the conditional vector field $\mathbf{u}_{t0}(\mathbf{x}|\mathbf{x}_0)$ with higher weight $\exp(-\beta\mathcal{E}(\mathbf{x}_0))$. As a result, the learnt velocity field $\mathbf{v}_t^{\boldsymbol{\theta}}(\mathbf{x})$ will generate $\mathbf{x}_0$ with lower energy $\mathcal{E}(\mathbf{x}_0)$.

Remark 4.6 (Connection with the importance sampling)**.** The conditional weighted energy guided loss $\mathcal{L}_{\text{CEFM}}$ can be also interpreted from the importance sampling techniques. Suppose we can sample directly from the data $q_0(\mathbf{x}) \propto p_0(\mathbf{x})\exp(-\beta\mathcal{E}(\mathbf{x}))$, minimizing the following loss $\mathcal{L}_q$ will get a velocity field $\mathbf{v}_t$ for generating distribution q_0

$$\mathcal{L}_q(\theta) = \mathbb{E}_{t,\mathbf{x}_0\sim q_0(\mathbf{x}),\mathbf{x}\sim q_{t0}(\mathbf{x}|\mathbf{x}_0)}[\|\mathbf{v}_t^{\boldsymbol{\theta}}(\mathbf{x}) - \mathbf{u}_{t0}(\mathbf{x}|\mathbf{x}_0)\|_2^2],$$

where $q_{t0}(\mathbf{x}|\mathbf{x}_0) = p_{t0}(\mathbf{x}|\mathbf{x}_0)$. Since , where Z is a constant, changing the data distribution from q_0 to p_0 yields that

$$\begin{aligned}
\mathcal{L}_q(\theta) &= \mathbb{E}_{t,\mathbf{x}_0\sim p_0(\mathbf{x}),\mathbf{x}\sim q_{t0}(\mathbf{x}|\mathbf{x}_0)}\left[\frac{q_0(\mathbf{x})}{p_0(\mathbf{x})}\|\mathbf{v}_t^{\boldsymbol{\theta}}(\mathbf{x}) - \mathbf{u}_{t0}(\mathbf{x}|\mathbf{x}_0)\|_2^2\right] \\
&= \mathbb{E}_{t,\mathbf{x}_0\sim p_0(\mathbf{x}),\mathbf{x}\sim p_{t0}(\mathbf{x}|\mathbf{x}_0)}\left[\frac{\exp(-\beta\mathcal{E}(\mathbf{x}_0))}{\mathbb{E}_{p_0(\widetilde{\mathbf{x}}_0)}[\exp(-\beta\mathcal{E}(\widetilde{\mathbf{x}}_0)]}\|\mathbf{v}_t^{\boldsymbol{\theta}}(\mathbf{x}) - \mathbf{u}_{t0}(\mathbf{x}|\mathbf{x}_0)\|_2^2\right] = \mathcal{L}_{\text{CEFM}}(\theta),
\end{aligned}$$

where the second equation is given by $q_0(\mathbf{x}) = p_0(\mathbf{x})\exp(-\beta\mathcal{E}(\mathbf{x}))/\mathbb{E}_{\mathbf{x}_0}[\exp(-\beta\mathcal{E}(\mathbf{x}_0)]$ according to Lemma B.1.

4.2 Weighted Diffusion Models

Theorem 4.3 suggests a general method to learn an energy-guided flow $\mathbf{v}^{\boldsymbol{\theta}}$ given any condition flow $\mathbf{u}_{t0}(\mathbf{x}|\mathbf{x}_0)$, including diffusion flow (Song et al., 2020), optimal transport (Lipman et al., 2022), rectified flow (Liu, 2022) or even more complicated $\mathbf{u}_{t0}(\mathbf{x}|\mathbf{x}_0)$. In this subsection, we restrict the analysis to the diffusion flow and present several useful analysis for the diffusion and score matching models. The first corollary provides the closed-form score function $\nabla_{\mathbf{x}}\log q_t(\mathbf{x})$ for the energy-guided distribution $q_t(\mathbf{x}) \propto p_t(\mathbf{x})\exp(-\mathcal{E}_t(\mathbf{x}))$:

Corollary 4.7. Under the assumptions claimed in Lemma 3.2, when $p_{t0}(\mathbf{x}|\mathbf{x}_0)$ is a Gaussian path with scheduler (μ_t, σ_t), we have $\nabla_{\mathbf{x}}\log q_t(\mathbf{x}) = \nabla_{\mathbf{x}}\log p_t(\mathbf{x}) - \nabla_{\mathbf{x}}\mathcal{E}_t(\mathbf{x})$ and

$$\nabla_{\mathbf{x}}\log q_t(\mathbf{x}) = \int_{\mathbf{x}_0} p_{0t}(\mathbf{x}_0|\mathbf{x})\nabla_{\mathbf{x}}\log p_{t0}(\mathbf{x}|\mathbf{x}_0)\frac{\exp(-\beta\mathcal{E}(\mathbf{x}_0))}{\exp(-\mathcal{E}_t(\mathbf{x}))}\mathrm{d}\mathbf{x}_0, \tag{4.4}$$

where $\nabla_{\mathbf{x}}\log p_{t0}(\mathbf{x}|\mathbf{x}_0) = -(\mathbf{x} - \mu_t\mathbf{x}_0)/\sigma_t^2 = -\boldsymbol{\epsilon}/\sigma_t$, $\boldsymbol{\epsilon} \sim \mathcal{N}(0, \mathbf{I}_d)$.

Corollary 4.7 suggests a method to estimate the guided score function without calculating the gradient of the intermediate energy function $\nabla_{\mathbf{x}}\mathcal{E}_t(\mathbf{x})$ as conducted in Lu et al. (2023). Then the following corollary suggests a similar energy-weighted diffusion model to train this score function $\nabla_{\mathbf{x}}\log q_t(\mathbf{x})$ in practice.

Corollary 4.8. Define the **E**nergy-weighted **D**iffusion loss $\mathcal{L}_{\text{ED}}$ and the **C**onditional **E**nergy-weighted **D**iffusion loss $\mathcal{L}_{\text{CED}}$ separately as

$$\begin{aligned}
\mathcal{L}_{\text{ED}}(\boldsymbol{\theta}) &= \mathbb{E}_{t,\mathbf{x}}\left[\frac{\exp(-\mathcal{E}_t(\mathbf{x}))}{\mathbb{E}_{p_t(\widetilde{\mathbf{x}})}[\exp(-\mathcal{E}_t(\widetilde{\mathbf{x}}))]}\|\mathbf{s}_t^{\boldsymbol{\theta}}(\mathbf{x}) - \nabla_{\mathbf{x}}\log q_t(\mathbf{x})\|_2^2\right], \\
\mathcal{L}_{\text{CED}}(\boldsymbol{\theta}) &= \mathbb{E}_{t,\mathbf{x},\mathbf{x}_0}\left[\frac{\exp(-\beta\mathcal{E}(\mathbf{x}_0))}{\mathbb{E}_{p_0(\widetilde{\mathbf{x}_0})}[\exp(-\beta\mathcal{E}(\widetilde{\mathbf{x}_0}))]}\left\|\mathbf{s}_t^{\boldsymbol{\theta}}(\mathbf{x}) - \nabla_{\mathbf{x}}\log p_{t0}(\mathbf{x}|\mathbf{x}_0)\right\|_2^2\right],
\end{aligned}$$

where the expectation is taken from $t \sim \lambda(t), \mathbf{x}_0 \sim p_0(\mathbf{x}_0)$ and $\mathbf{x} \sim p_{t0}(\mathbf{x}|\mathbf{x}_0)$. Thus the marginal distribution of $\mathbf{x}$ is is $p_t(\mathbf{x})$. $\mathcal{L}_{\text{ED}}(\boldsymbol{\theta})$ is equal with $\mathcal{L}_{\text{CED}}(\boldsymbol{\theta})$ up to a constant and thus $\nabla_{\boldsymbol{\theta}}\mathcal{L}_{\text{ED}}(\boldsymbol{\theta}) = \nabla_{\boldsymbol{\theta}}\mathcal{L}_{\text{CED}}(\boldsymbol{\theta})$.

Remark 4.9. A similar approach is proposed in Wang et al. (2024) for estimating $\nabla_{\mathbf{x}}\mathcal{E}_t(\mathbf{x})$ using a neural network by

$$\nabla_{\mathbf{x}}\mathcal{E}_t(\mathbf{x}) = \frac{\mathbb{E}_{p_{0t}(\mathbf{x}_0|\mathbf{x})}\left[\exp(-\beta\mathcal{E}(\mathbf{x}_0))\left(\nabla_{\mathbf{x}}\log p_t(\mathbf{x}) - \nabla_{\mathbf{x}}\log p_{t0}(\mathbf{x}|\mathbf{x}_0)\right)\right]}{\exp(-\mathcal{E}_t(\mathbf{x}))}, \tag{4.5}$$

and then plugging (4.5) back to $\nabla_{\mathbf{x}}\log q_t(\mathbf{x}) = \nabla_{\mathbf{x}}\log p_t(\mathbf{x}) - \nabla_{\mathbf{x}}\mathcal{E}_t(\mathbf{x})$ for generating the score function $\nabla_{\mathbf{x}}\log q_t(\mathbf{x})$. However, in order to obtain $\nabla_{\mathbf{x}}\mathcal{E}_t(\mathbf{x})$, Wang et al. (2024) requires to estimate $\exp(\mathcal{E}_t(\mathbf{x}))$ by sampling and approximate $\nabla_{\mathbf{x}}\mathcal{E}_t(\mathbf{x})$ via another neural network. In contrast, using $\mathcal{L}_{\text{CED}}$ as discussed in Corollary 4.8 removes the necessity of estimating both $\nabla_{\mathbf{x}}\mathcal{E}_t(\mathbf{x})$ and $\mathcal{E}_t(\mathbf{x})$. Therefore, Energy-weighted diffusion can directly obtain the score function for guided distribution without additional sampling or back-propagation.

Algorithm 1 Training Energy-Weighted Diffusion Model

Input: Score function $\mathbf{s}_t^{\boldsymbol{\theta}}(\cdot)$, schedule (μ_t, σ_t), guidance scale β, batch size B, time weight $\lambda(t)$
1: **for** batch $\{\mathbf{x}_0^i, \mathcal{E}(\mathbf{x}_0^i)\}_i$ **do**
2: **for** index $i \in [B]$ **do**
3: Calculate guidance $g_i = \mathrm{softmax}(-\beta\mathcal{E}(\mathbf{x}_0^i)) = \exp(-\beta\mathcal{E}(\mathbf{x}_0^i))/\sum_j \exp(-\beta\mathcal{E}(\mathbf{x}_0^j))$
4: Sample $t_i \sim U(0,1)$, calculate μ_{t_i}, σ_{t_i}, sample $\boldsymbol{\epsilon}_i \sim \mathcal{N}(0, \mathbf{I}_d)$ and $\mathbf{x}_{t_i} = \mu_{t_i}\mathbf{x}_0^i + \sigma_{t_i}\boldsymbol{\epsilon}_i$
5: **end for**
6: Calculate and take a gradient step using $\mathcal{L}_{\mathrm{CED}}(\boldsymbol{\theta}) = \sum_i \lambda(t_i) g_i \|\mathbf{s}_{t_i}^{\boldsymbol{\theta}}(\mathbf{x}_{t_i}) + \boldsymbol{\epsilon}_i/\sigma_{t_i}\|_2^2$.
7: **end for**

In the implementation of the diffusion models, since $\mathbf{x} \sim \mathcal{N}(\mu_t\mathbf{x}_0, \sigma_t^2\mathbf{I})$, the conditional score function $\nabla_{\mathbf{x}} \log p_{t0}(\mathbf{x}|\mathbf{x}_0) = -\boldsymbol{\epsilon}/\sigma_t$ where $\boldsymbol{\epsilon} = (\mathbf{x}_t - \mu\mathbf{x}_0)/\sigma_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. In addition, the denominator $\mathbb{E}_{p_0(\mathbf{x}_0)}[\exp(-\beta\mathcal{E}(\mathbf{x}_0))]$ can be approximated by the empirical average $\sum_{i\in N} \exp(-\beta\mathcal{E}(\mathbf{x}_0^i))/N$ in a batch. A practical algorithm is presented in Algorithm 1. The training process is similar with the standard DDPM (Ho et al., 2020) or score matching (Song et al., 2020). The only difference is that we incorporate the weight by calculating a energy guidance g_i using the softmax value of using $-\beta\mathcal{E}(\mathbf{x}_0^i)$ from the current batch in Line 3.

4.3 Comparison between CEP and classifier (free) guidance

In this subsection, we compare our method with Contrastive Energy Prediction (CEP, Lu et al. 2023), Classifier-Guidance (CG, Dhariwal & Nichol 2021) and Classifier-Free Guidance (CFG, Ho & Salimans 2021). We consider the guided distribution $q_0(\mathbf{x}) \propto p_0(\mathbf{x})p^\beta(c|\mathbf{x})$ where $p_0(c|\mathbf{x})$ is the classifier, β is the guidance scale, c is the desired class which we fix during the analysis. Comparing with the formulation of the energy guidance $q_1(\mathbf{x}) \propto p_1(\mathbf{x})\exp(-\beta\mathcal{E}(\mathbf{x}))$, the "energy function" can be interpreted as $\mathcal{E}(\mathbf{x}) = -\log p(c|\mathbf{x})$. To begin with, the following lemma provides the closed-form solution for the energy-guided diffusion and classifier-guided diffusion

Lemma 4.10. Given the same guidance scale β and the same diffusion process, let the energy function be defined by $\mathcal{E}(\mathbf{x}) = -\log p(c|\mathbf{x})$, the score function for CG and CFG are both:

$$\nabla_{\mathbf{x}} \log q_t(\mathbf{x}) = \nabla_{\mathbf{x}} \log p_t(\mathbf{x}) + \nabla_{\mathbf{x}} \log \left[\mathbb{E}_{p_{0t}(\mathbf{x}_0|\mathbf{x})} p(c|\mathbf{x}_0)\right]^\beta, \tag{4.6}$$

while the score function for energy-weighted diffusion and CEP are both

$$\nabla_{\mathbf{x}} \log q_t(\mathbf{x}) = \nabla_{\mathbf{x}} \log p_t(\mathbf{x}) + \nabla_{\mathbf{x}} \log \mathbb{E}_{p_{0t}(\mathbf{x}_0|\mathbf{x})} p^\beta(c|\mathbf{x}_0). \tag{4.7}$$

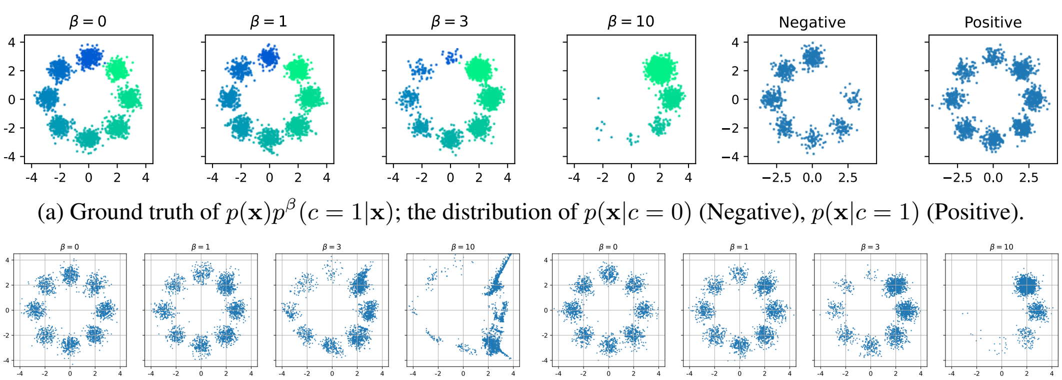

(a) Ground truth of $p(\mathbf{x})p^\beta(c=1|\mathbf{x})$; the distribution of $p(\mathbf{x}|c=0)$ (Negative), $p(\mathbf{x}|c=1)$ (Positive).

(b) Data sampled from classifier-free diffusion.

(c) Data sampled from energy-weighted diffusion.

Figure 1: Visualization of the ground-truth distribution $p(\mathbf{x})p^\beta(c=1|\mathbf{x})$ with different values of β, the posterior distribution $p(\mathbf{x}|c)$ with $c \in \{0,1\}$, and the data sampled from classifier-free diffusion and energy-weighted diffusion. The energy-weighted diffusion process demonstrates better performance when $\beta > 1$. More examples and details of this experiments are provided in Appendix C.

The score function (4.7) is exactly the score function that generates $q_0(\mathbf{x}) \propto p_0(\mathbf{x})\exp(-\beta\mathcal{E}(\mathbf{x})) = p_0(\mathbf{x})p(c|\mathbf{x})^\beta$ according to Theorem 3.3. In a sharp contrast, (4.6) is not guaranteed to generate the desired distribution q_0 when $\beta \neq 1$ because $[\mathbb{E}_{p_{0t}(\mathbf{x}_0|\mathbf{x})} p(c|\mathbf{x}_0)]^\beta \neq \mathbb{E}_{p_{0t}(\mathbf{x}_0|\mathbf{x})} p^\beta(c|\mathbf{x}_0)$. As demonstrated in Figure 1, when $\beta = 1$, the distributions generated by CFG and energy-guided diffusion are both similar to the ground-truth distribution. However, when $\beta > 1$, the distribution generated by CFG differs from the ground-truth distribution, whereas the energy-guided diffusion can still generate the ground-truth distribution. Finally, the following lemma also verifies that when $\beta = 1$, the energy-weighted diffusion process is the same with the classifier-free guidance to learn the score function of the posterior distribution $\nabla_{\mathbf{x}} \log p_t(\mathbf{x}|c)$:

Lemma 4.11. Let $\beta = 1$ and assume that $\mathcal{E}(\mathbf{x}) = -\log p(c|\mathbf{x})$ for some fixed c, then $\mathcal{L}_{\text{CED}}(\boldsymbol{\theta})$ is

$$\mathcal{L}_{\text{CED}}(\boldsymbol{\theta}) = \mathbb{E}_{t,\mathbf{x},\mathbf{x}_0|c} \|\mathbf{s}_t^{\boldsymbol{\theta}}(\mathbf{x}) - \nabla_{\mathbf{x}} \log p_{t0}(\mathbf{x}|\mathbf{x}_0)\|_2^2, \tag{4.8}$$

where the expectation is taken over $t \sim \lambda(t), \mathbf{x}_0 \sim p_0(\cdot|c)$ and $\mathbf{x} \sim p_{t0}(\cdot|\mathbf{x}_0)$. The global optimal for (4.8) is $\mathbf{s}_t^{\boldsymbol{\theta}}(\mathbf{x}) = \nabla_{\mathbf{x}} \log p_t(\mathbf{x}|c)$.

In addition to the *exact* guidance provided by the CEP and energy-weighted diffusion, it is important to highlight that the energy-weighted diffusion model eliminates the necessity for the intermediate energy model. This advantage is similar to the simplicity provided by the CFG, compared with CG. We summarized the difference and connection between energy-weighted diffusion, CEP, CFG and CG in Table 1.

5 Q-WEIGHTED ITERATIVE POLICY OPTIMIZATION FOR OFFLINE RL

We consider the episodic Markov Decision Processes denoted by $\mathcal{M}(\mathcal{S}, \mathcal{A}, r, \mathbb{P}, \gamma)$ with $\mathcal{S}, \mathcal{A}$ denoting the state and action space respectively. r is the reward function, $\mathbb{P}(\cdot|\mathbf{s}, \mathbf{a})$ is the transition kernel, and γ is the discount factor. In offline RL, the data is collected by a *behavioral policy* μ. The policy optimization with KL regularization (Peters et al., 2010; Peng et al., 2019) is formulated as

$$\operatorname*{argmax}_{\pi^{\boldsymbol{\theta}}} \mathbb{E}_{\mathbf{a}\sim\pi^{\boldsymbol{\theta}}(\cdot|\mathbf{x})} Q(\mathbf{x}, \mathbf{a}) - \tfrac{1}{\beta} \mathrm{KL}(\pi^{\boldsymbol{\theta}} \parallel \mu), \tag{5.1}$$

where $\mathbf{x}$ denotes the state and $\mathbf{a}$ denotes the action. $Q(\mathbf{x}, \mathbf{a})$ is the estimation of the state-action value function $Q^\pi(\mathbf{x}, \mathbf{a}) = \mathbb{E}[\sum_{t=0}^\infty \gamma_t r(\mathbf{x}_t, \mathbf{a}_t)|\mathbf{x}_0 = xb, \mathbf{a}_0 = \mathbf{a}, \pi]$. The closed-form solution to (5.1) is $\pi(\mathbf{a}|\mathbf{x}) \propto \mu(\mathbf{a}|\mathbf{x}) \exp(\beta Q(\mathbf{x}, \mathbf{a}))$. Following the procedure of the energy-weighted diffusion model discussed in Section 4, we propose learning the score function $\mathbf{s}_t^{\boldsymbol{\theta}}(\mathbf{a}, \mathbf{x})$ or the velocity field $\mathbf{v}_t^{\boldsymbol{\theta}}(\mathbf{a}, \mathbf{x})$ that generates the optimal policy $\pi(\cdot|\mathbf{x})$ using the Q-weighted diffusion loss $\mathcal{L}_{\text{QD}}$ or Q-weighted flow matching loss $\mathcal{L}_{\text{QF}}$, respectively, defined as:

$$\mathcal{L}_{\text{QD}}(\boldsymbol{\theta}) = \mathbb{E}_{t,(\mathbf{x},\mathbf{a}),\mathbf{a}_t} \left[\frac{\exp(\beta Q(\mathbf{x}, \mathbf{a}))}{\mathbb{E}_{\widetilde{\mathbf{a}}\sim\mu(\cdot|\mathbf{x})}[\exp(\beta Q(\mathbf{x}, \widetilde{\mathbf{a}}))]} \left\| \mathbf{s}_t^{\boldsymbol{\theta}}(\mathbf{a}_t; \mathbf{x}) - \nabla_{\mathbf{a}_t} \log p_{t0}(\mathbf{a}_t|\mathbf{a}_0) \right\|_2^2 \right], \tag{5.2}$$

$$\mathcal{L}_{\text{QF}}(\boldsymbol{\theta}) = \mathbb{E}_{t,(\mathbf{x},\mathbf{a}),\mathbf{a}_t} \left[\frac{\exp(\beta Q(\mathbf{x}, \mathbf{a}))}{\mathbb{E}_{\widetilde{\mathbf{a}}\sim\mu(\cdot|\mathbf{x})}[\exp(\beta Q(\mathbf{x}, \widetilde{\mathbf{a}}))]} \left\| \mathbf{v}_t^{\boldsymbol{\theta}}(\mathbf{a}_t; \mathbf{x}) - \mathbf{u}_{t0}(\mathbf{a}_t|\mathbf{a}) \right\|_2^2 \right], \tag{5.3}$$

where the expectation is taken over $t \sim \lambda(t)$, $(\mathbf{x}, \mathbf{a})$ is sampled from the offline RL dataset, and $\mathbf{a}_t \sim p_{t0}(\cdot|\mathbf{a})$. Two components are essential for training either (5.2) or (5.3): First, the behavioral policy $\mu(\cdot|\mathbf{x})$ can be trained via standard diffusion or flow matching models. Second, any Q function derived from offline RL algorithms can be used as $Q(\mathbf{x}, \mathbf{a})$ in (5.2) or (5.3).

5.1 PROPOSED ALGORITHM

We present the algorithm sketch for the Q-weighted diffusion process in Algorithm 2, and defer the detailed implementation to Appendix D. From a high-level overview, the algorithm first trains the score model to match the behavioral policy $\mu(\mathbf{a}, \mathbf{x})$ for K_1 episodes, then trains the Q function for K_2 episodes. The algorithm then performs the *Q-weighted iterative policy optimization* (QIPO) as follows: First, in Line 10, using the current score function $\mathbf{s}^{\boldsymbol{\theta}}$, the algorithm samples several support actions $\mathbf{a}_{ij}$ to estimate the expectation $\mathbb{E}_{\widetilde{\mathbf{a}}\sim\mu(\cdot|\mathbf{x})}$ in (5.2). In Line 15, the algorithm optimizes $\boldsymbol{\theta}$ with respect to the empirical estimation of $\mathcal{L}_{\text{QD}}(\boldsymbol{\theta})$. Therefore, assuming the score function $\mathbf{s}_t^{\boldsymbol{\theta}}(\cdot|\mathbf{x})$ generates the behavior distribution $\mu(\cdot|\mathbf{x})$ after the warm-up for K_1 episodes, using the sampled support action $\mathbf{a}_{ij}$, the optimal score function for Line 15 corresponds to the policy $\pi_1 = \mu(\cdot|\mathbf{x}) \exp(\beta Q^{\boldsymbol{\psi}}(\cdot|\mathbf{x}))$. When the number of episodes $k = K_{\text{renew}} + 1$, the algorithm revisits Line 10 with the new score function and regenerates the support action set with the new policy

Algorithm 2 Q-weighted iterative policy optimization for offline RL (diffusion)

Input: Score function $\mathbf{s}_t^{\boldsymbol{\theta}}(\cdot)$, schedule (μ_t, σ_t), guidance scale β, batch size B, weight on time $\lambda(t)$
Input: Offline RL dataset $\mathcal{D} = \{(\mathbf{x}, \mathbf{a}, \mathbf{x}', r)\}$, number of epochs K_1, K_2 and K_3, Q function $Q^{\boldsymbol{\psi}}$
Input: Support action set size M, support action set renew frequency K_{renew}
1: **for** diffusion model warm-up epoch $k \in [K_1]$ **do**
2: Train $\mathbf{s}_t^{\boldsymbol{\theta}}$ for each batch $\{(\mathbf{x}_i, \mathbf{a}_i, \mathbf{x}_i', r_i)\}_i \subset \mathcal{D}$ following standard diffusion model.
3: **end for**
4: **for** Q-learning warm-up epoch $k \in [K_2]$ **do**
5: Train $Q^{\boldsymbol{\psi}}$ for each batch $\{(\mathbf{x}_i, \mathbf{a}_i, \mathbf{x}_i', r_i)\}_i \subset \mathcal{D}$
6: **end for**
7: **for** policy improvement step $k \in [K_3]$ **do**
8: **for** batch $\{(\mathbf{x}_i, \mathbf{a}_i, \mathbf{x}_i', r_i)\}_i \subset \mathcal{D}$ **do**
9: **if** $k \bmod K_{\text{renew}} = 1$ **then**
10: Sample support action set $\mathbf{a}_{ij}$ using score function $\mathbf{s}^{\boldsymbol{\theta}}(\cdot|\mathbf{x}_i)$ for all $i \in [B], j \in [M]$
11: Denote $\mathbf{a}_{i0} = \mathbf{a}_i$, sample $t_{ij} \sim U(0,1)$ and $\mathbf{a}_{ij,t_{ij}} \sim \mathcal{N}(\mu_{t_{ij}} \mathbf{a}_{ij}, \sigma_{t_{ij}}^2 \mathbf{I})$ for all $j \in \overline{[M]}$
12: Calculate guidance $g_{ij} = \text{softmax}_j(\beta Q^{\boldsymbol{\psi}}(\mathbf{x}_i, \mathbf{a}_{ij})) = \frac{\exp(\beta Q^{\boldsymbol{\psi}}(\mathbf{x}_i, \mathbf{a}_{ij}))}{\sum_{j=0}^{M} \exp(\beta Q^{\boldsymbol{\psi}}(\mathbf{x}_i, \mathbf{a}_{ij}))}$
13: **end if**
14: Calculate loss $\mathcal{L}_{\text{QD}}(\boldsymbol{\theta}) = \sum_{i=1,j=0}^{B,M} \lambda(t_{ij}) g_{ij} \left\| \mathbf{s}_{t_{ij}}^{\boldsymbol{\theta}}(\mathbf{a}_{ij,t_{ij}}; \mathbf{x}_i) - \nabla_{\mathbf{a}_t} \log p_{t0}(\mathbf{a}_{ij,t_{ij}} | \mathbf{a}_{ij}) \right\|_2^2$
15: Update $\boldsymbol{\theta}$ using the gradient of $\mathcal{L}_{\text{QD}}(\boldsymbol{\theta})$
16: **end for**
17: **end for**

π_1. Thus the target score function for Line 15 to optimize is $\pi_2(\cdot|\mathbf{x}) \propto \pi_1(\cdot|\mathbf{x}) \exp(\beta Q^{\boldsymbol{\psi}}(\cdot|\mathbf{x})) \propto \mu(\cdot|\mathbf{x}) \exp(2\beta Q^{\boldsymbol{\psi}}(\cdot|\mathbf{x}))$. As a result, denoting $l = (k-1) \bmod K_{\text{renew}}$, the policy π_l generated by the score function $\mathbf{s}_t^{\boldsymbol{\theta}_k}$ is:

$$\pi_{l+1}(\mathbf{a}|\mathbf{x}) \propto \pi_l(\mathbf{a}|\mathbf{x}) \exp(\beta Q^{\boldsymbol{\psi}}(\mathbf{a}, \mathbf{x})) \propto \cdots \propto \mu(\mathbf{a}|\mathbf{x}) \exp((l+1)\beta Q^{\boldsymbol{\psi}}(\mathbf{a}, \mathbf{x})), \tag{5.4}$$

which will implicitly increase the guidance scale β.

Similar weighted policy optimization approaches have been applied in Kang et al. (2024); Ding et al. (2024). However, QIPO builds the relationship between the KL-regularized policy optimization so that QIPO can iteratively improve the policy as described in (5.4). Compared with directly setting a large guidance scale β, QIPO makes the support action set $\widetilde{a}$ to be more concentrated in the space with higher Q values. As a result, QIPO learns a more robust Q-weighted score function $\mathbf{s}_t^{\boldsymbol{\theta}}$ compared to one-step Q-weighted diffusion with a larger β. Second, QGPO (Lu et al., 2023) introduces a scaling factor s and composes the score function as $\nabla_{\mathbf{a}_t} \log \pi_t(\mathbf{a}_t|\mathbf{x}) = \nabla_{\mathbf{a}_t} \log \mu_t(\mathbf{a}_t|\mathbf{x}) + s\nabla_{\mathbf{a}_t} Q_t(\mathbf{x}, \mathbf{a}_t)$, where Q_t is the intermediate Q function, similar to the $\mathcal{E}_t$ in Section 4. However, as we discussed in Section 4.3 about the comparison of the CFG, since

$$sQ_t(\mathbf{a}_t, \mathbf{x}) = -s \log \mathbb{E}_{p_{0t}(\mathbf{a}|\mathbf{a}_t)}[\exp(\beta Q(\mathbf{a}, \mathbf{x})] \neq \log \mathbb{E}_{p_{0t}(\mathbf{a}|\mathbf{a}_t)}[\exp(s\beta Q(\mathbf{a}, \mathbf{x})],$$

using a guidance scale $s > 1$ does not guarantee generating a policy strictly regularized by the behavioral policy $\mu(\mathbf{a}|\mathbf{x})$. In contrast, as (5.4) suggests, our approach strictly follows the formulation $\pi(\mathbf{a}|\mathbf{x}) \propto \mu(\mathbf{a}|\mathbf{x}) \exp(s\beta Q^{\boldsymbol{\psi}}(\mathbf{a}|\mathbf{x}))$ regularized by the behavioral policy.

5.2 Experiment Results

We evaluate the performance of QIPO with flow matching and diffusion model on the D4RL tasks (Fu et al., 2020) in this subsection.

Experiment configurations We implement the flow matching model **QIPO-OT** using the optimal-transport conditional velocity fields (Lipman et al., 2022) and the diffusion model **QIPO-Diff** using VP-SDE (Song et al., 2020). We use the same network structure as QGPO for a fair comparison of the efficiency with QGPO. We defer the detailed experiment configurations in Appendix E.1.

We compare our results with other state-of-the-art benchmarks, including Diffusion-QL (Wang et al., 2022), QGPO (Lu et al., 2023), IDQL (Hansen-Estruch et al., 2023), SRPO (Chen et al., 2023) and Guided Flows (Zheng et al., 2023) and present the results in Table 2. As the experiment results

Table 2: Evaluation numbers of D4RL benchmarks (normalized as suggested by Fu et al. (2020)). We report mean ± standard deviation of algorithm performance across 8 random seeds. The highest performance is **boldfaced highlighted.** The performance within 5% of the maximum absolute value in every individual task are highlighted.

Dataset	Environment	SfBC	QGPO	IDQL	SRPO	Guided Flows	QIPO-Diff (ours)	QIPO-OT (ours)
Medium-Expert	HalfCheetah	92.6	93.5	95.9	92.2	**97**	94.14 ± 0.48	94.45 ± 0.49
Medium-Expert	Hopper	108.6	108.0	108.6	100.1	105	$\mathbf{112.12 \pm 0.42}$	108.02 ± 5.19
Medium-Expert	Walker2d	109.8	110.7	112.7	**114.0**	94	110.14 ± 0.51	110.87 ± 1.04
Medium	HalfCheetah	45.9	54.1	51.0	**60.4**	49	48.19 ± 0.20	54.16 ± 1.27
Medium	Hopper	57.1	**98.0**	65.4	95.5	84	89.53 ± 9.96	94.05 ± 13.27
Medium	Walker2d	77.9	86.0	82.5	84.4	77	84.99 ± 0.46	$\mathbf{87.61 \pm 1.46}$
Medium-Replay	HalfCheetah	37.1	47.6	45.9	**51.4**	42	45.27 ± 0.42	48.04 ± 0.79
Medium-Replay	Hopper	86.2	96.9	92.1	101.2	89	101.23 ± 0.47	$\mathbf{101.25 \pm 2.18}$
Medium-Replay	Walker2d	65.1	84.4	85.1	84.6	78	90.08 ± 4.53	78.57 ± 26.09
Average (Locomotion)		75.6	86.6	82.1	87.1	79.4	86.2	86.3
Default	AntMaze-umaze	92.0	96.4	94.0	97.1	-	$\mathbf{97.5 \pm 0.53}$	93.62 ± 7.05
Diverse	AntMaze-umaze	**85.3**	74.4	80.2	82.1	-	73.88 ± 6.42	76.12 ± 9.93
Play	AntMaze-medium	81.3	83.6	**84.5**	80.7	-	82.75 ± 3.24	80.00 ± 13.66
Diverse	AntMaze-medium	82.0	83.8	84.8	75.0	-	$\mathbf{86.00 \pm 8.65}$	$\mathbf{86.42 \pm 5.44}$
Play	AntMaze-large	59.3	66.6	63.5	53.6	-	$\mathbf{73.25 \pm 10.90}$	55.5 ± 29.39
Diverse	AntMaze-large	45.5	64.8	**67.9**	53.6	-	40.5 ± 20.40	32.13 ± 23.16
Average (AntMaze)		74.2	78.3	**79.1**	73.6	-	77.3	71.96

suggests, QIPO-Diff and QIPO-OT consistently outperform the baselines in various tasks. We defer more baseline algorithms for comparison to Table 4 in Appendix E.

Among these benchmark algorithms, we would like to highlight the comparison between *Guided Flows* (Zheng et al., 2023) and *Q-Guided Policy Optimization* (QGPO, Lu et al. 2023). Firstly, compared with Guided Flows (Zheng et al., 2023), QIPO-OT enjoys higher performance across many different tasks. This improved performance is due to the fact that energy-based guidance will provide more accurate guidance compared with classifier-free guidance, as discussed in Section 4.3.

Secondly, compared with QGPO (Lu et al., 2023), QIPO-Diff directly learns the energy-guided score function without estimating the intermediate energy function. As a result, QIPO does not require the back-propagation to calculate the gradient of the intermediate energy function $\nabla_{\mathbf{x}}\mathcal{E}_t(\mathbf{x})$ and therefore enjoys a faster sampling speed compared with QGPO when using the same score network, as presented in Table 3. In addition, since the QIPO guarantees the strict formulation $\pi(\mathbf{a}|\mathbf{x}) \propto \mu(\mathbf{a}|\mathbf{x})\exp(s\beta Q^{\psi}(\mathbf{a}|\mathbf{x}))$ as shown in (5.4), QIPO enjoys better performance compared with QGPO on various tasks. We defer more detailed discussions on the advantage of QIPO-OT to Appendix E.2.

Table 3: Comparison of the running time for action generation between QGPO (Lu et al., 2023) and QIPO, averaged over 1500 runs. The percentage reduction in time compared to QGPO is also reported.

Method	Time (ms)
QGPO (Lu et al., 2023)	75.05
QIPO-OT (ours)	**27.26** (-63.68%)
QIPO-Diff (ours)	55.86 (-25.57%)

Ablation Study. We conduct ablation study on changing the support action set M, policy renew period K_{renew} and the guidance scale β. We defer the detailed ablation study result in Appendix E.3.

6 Conclusion and Future Work

In this paper, we explored the energy guidance in both flow matching and diffusion models. We introduced Energy-weighted Flow Matching (EFM) and Energy-weighted Diffusion (ED) by incorporating the energy guidance directly into these generative models without relying on auxiliary models or post-processing steps. We applied the proposed methods in offline RL and introduced Q-weighted Iterative Policy Optimization (QIPO), which enjoys competitive empirical performance on the D4RL benchmark. To the best of our knowledge, this work is the first to present an energy-guided flow matching model and the first algorithm to *directly* learn an energy-guided diffusion model. While the current QIPO focuses on offline RL without interacting with the environment, we leave the extension to online RL, where guidance from the Q-function can be updated through online interactions

ETHICS STATEMENT

This paper focuses on the methodology of generative models and strictly adheres to ethical guidelines and standards. As with other deep generative modeling techniques, energy-weighted diffusion and flow matching model have the potential to generate harmful content, such as “deepfakes”, and may reflect or amplify unwanted social biases present in the training data. However, these broader societal impacts are not directly relevant to the specific contributions of our work. Therefore, we do not believe any unique ethical concerns or negative social impacts arise from this paper.

REPRODUCIBILITY STATEMENT

We provide detailed descriptions of our experimental setups, dataset construction processes, and code implementation in the supplementary materials to ensure the reproducibility of the Energy-weighted Diffusion and QIPO methods. The full experimental configurations are presented in Appendix E.1. To further support research in the community, we will release the model checkpoints following the de-anonymization process.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their helpful comments. Partial of this work was done while WZ was a PhD student at UCLA. WZ and QG are supported in part by the NSF grants DMS-2323113, CPS-2312094, IIS-2403400 and the research fund from UCLA-Amazon Science Hub. WZ is also supported by UCLA dissertation year fellowship. Part of this work is supported by the Google Cloud Research Credits program with the award GCP376319164. The views and conclusions contained in this paper are those of the authors and should not be interpreted as representing any funding agencies.

REFERENCES

Anurag Ajay, Yilun Du, Abhi Gupta, Joshua B Tenenbaum, Tommi S Jaakkola, and Pulkit Agrawal. Is conditional generative modeling all you need for decision making? In *The Eleventh International Conference on Learning Representations*, 2022.

Jacob Austin, Daniel D Johnson, Jonathan Ho, Daniel Tarlow, and Rianne Van Den Berg. Structured denoising diffusion models in discrete state-spaces. *Advances in Neural Information Processing Systems*, 34:17981–17993, 2021.

Huayu Chen, Cheng Lu, Chengyang Ying, Hang Su, and Jun Zhu. Offline reinforcement learning via high-fidelity generative behavior modeling. In *The Eleventh International Conference on Learning Representations*, 2022.

Huayu Chen, Cheng Lu, Zhengyi Wang, Hang Su, and Jun Zhu. Score regularized policy optimization through diffusion behavior. In *The Twelfth International Conference on Learning Representations*, 2023.

Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018.

Julian Cremer, Tuan Le, Frank Noé, Djork-Arné Clevert, and Kristof T Schütt. Pilot: Equivariant diffusion for pocket conditioned de novo ligand generation with multi-objective guidance via importance sampling. *arXiv preprint arXiv:2405.14925*, 2024.

Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021.

Shutong Ding, Ke Hu, Zhenhao Zhang, Kan Ren, Weinan Zhang, Jingyi Yu, Jingya Wang, and Ye Shi. Diffusion-based reinforcement learning via q-weighted variational policy optimization. *arXiv preprint arXiv:2405.16173*, 2024.

Zeyu Fang and Tian Lan. Learning from random demonstrations: Offline reinforcement learning with importance-sampled diffusion models. *arXiv preprint arXiv:2405.19878*, 2024.

Justin Fu, Aviral Kumar, Ofir Nachum, George Tucker, and Sergey Levine. D4rl: Datasets for deep data-driven reinforcement learning, 2020.

Philippe Hansen-Estruch, Ilya Kostrikov, Michael Janner, Jakub Grudzien Kuba, and Sergey Levine. Idql: Implicit q-learning as an actor-critic method with diffusion policies. *arXiv preprint arXiv:2304.10573*, 2023.

Longxiang He, Linrui Zhang, Junbo Tan, and Xueqian Wang. Diffcps: Diffusion model based constrained policy search for offline reinforcement learning. *arXiv preprint arXiv:2310.05333*, 2023.

Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*, 2021.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.

Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. *Advances in Neural Information Processing Systems*, 35:8633–8646, 2022.

Emiel Hoogeboom, Vıctor Garcia Satorras, Clément Vignac, and Max Welling. Equivariant diffusion for molecule generation in 3d. In *International conference on machine learning*, pp. 8867–8887. PMLR, 2022.

Matthew Thomas Jackson, Michael Tryfan Matthews, Cong Lu, Benjamin Ellis, Shimon Whiteson, and Jakob Foerster. Policy-guided diffusion. *arXiv preprint arXiv:2404.06356*, 2024.

Michael Janner, Yilun Du, Joshua Tenenbaum, and Sergey Levine. Planning with diffusion for flexible behavior synthesis. In *International Conference on Machine Learning*, pp. 9902–9915. PMLR, 2022.

Bingyi Kang, Xiao Ma, Chao Du, Tianyu Pang, and Shuicheng Yan. Efficient diffusion policies for offline reinforcement learning. *Advances in Neural Information Processing Systems*, 36, 2024.

Jaewoo Lee, Sujin Yun, Taeyoung Yun, and Jinkyoo Park. Gta: Generative trajectory augmentation with guidance for offline reinforcement learning. *arXiv preprint arXiv:2405.16907*, 2024.

TP Lillicrap. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.

Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *The Eleventh International Conference on Learning Representations*, 2022.

Qiang Liu. Rectified flow: A marginal preserving approach to optimal transport. *arXiv preprint arXiv:2209.14577*, 2022.

Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *Advances in Neural Information Processing Systems*, 35:5775–5787, 2022a.

Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver++: Fast solver for guided sampling of diffusion probabilistic models. *arXiv preprint arXiv:2211.01095*, 2022b.

Cheng Lu, Huayu Chen, Jianfei Chen, Hang Su, Chongxuan Li, and Jun Zhu. Contrastive energy prediction for exact energy-guided diffusion sampling in offline reinforcement learning. In *International Conference on Machine Learning*, pp. 22825–22855. PMLR, 2023.

Cong Lu, Philip Ball, Yee Whye Teh, and Jack Parker-Holder. Synthetic experience replay. *Advances in Neural Information Processing Systems*, 36, 2024.

Xue Bin Peng, Aviral Kumar, Grace Zhang, and Sergey Levine. Advantage-weighted regression: Simple and scalable off-policy reinforcement learning. *arXiv preprint arXiv:1910.00177*, 2019.

Jan Peters and Stefan Schaal. Reinforcement learning by reward-weighted regression for operational space control. In *Proceedings of the 24th international conference on Machine learning*, pp. 745–750, 2007.

Jan Peters, Katharina Mulling, and Yasemin Altun. Relative entropy policy search. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 24, pp. 1607–1612, 2010.

Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image synthesis. In *The Twelfth International Conference on Learning Representations*, 2023.

Vıctor Garcia Satorras, Emiel Hoogeboom, and Max Welling. E (n) equivariant graph neural networks. In *International conference on machine learning*, pp. 9323–9332. PMLR, 2021.

Marcin Sendera, Minsu Kim, Sarthak Mittal, Pablo Lemos, Luca Scimeca, Jarrid Rector-Brooks, Alexandre Adam, Yoshua Bengio, and Nikolay Malkin. Improved off-policy training of diffusion samplers. In *The Thirty-Eighth Annual Conference on Neural Information Processing Systems*, pp. 1–30. ACM, 2024.

Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2020.

Cédric Villani et al. *Optimal transport: old and new*, volume 338. Springer, 2009.

Yan Wang, Lihao Wang, Yuning Shen, Yiqun Wang, Huizhuo Yuan, Yue Wu, and Quanquan Gu. Protein conformation generation via force-guided se (3) diffusion models. *arXiv preprint arXiv:2403.14088*, 2024.

Zhendong Wang, Jonathan J Hunt, and Mingyuan Zhou. Diffusion policies as an expressive policy class for offline reinforcement learning. In *The Eleventh International Conference on Learning Representations*, 2022.

Qinqing Zheng, Matt Le, Neta Shaul, Yaron Lipman, Aditya Grover, and Ricky TQ Chen. Guided flows for generative modeling and decision making. *arXiv preprint arXiv:2311.13443*, 2023.