

EXTENDED FLOW MATCHING : A METHOD OF CONDITIONAL GENERATION WITH GENERALIZED CONTINUITY EQUATION

Anonymous authors

Paper under double-blind review

ABSTRACT

Conditional generative modeling (CGM), which approximates the conditional probability distribution of data given a condition, holds significant promise for generating new data across diverse representations. While CGM is crucial for generating images, video, and text, its application to scientific computing, such as molecular generation and physical simulations, is also highly anticipated. A key challenge in applying CGM to scientific fields is the sparseness of available data conditions, which requires extrapolation beyond observed conditions. This paper proposes the Extended Flow Matching (EFM) framework to address this challenge. EFM achieves smooth transitions in distributions when departing from observed conditions, avoiding the unfavorable changes seen in existing flow matching (FM) methods. By introducing a flow with respect to the conditional axis, EFM ensures that the conditional distribution changes gradually with the condition. Specifically, we apply an extended Monge–Kantorovich theory to conditional generative models, creating a framework for learning matrix fields in a generalized continuity equation instead of vector fields. Furthermore, by combining the concept of Dirichlet energy on Wasserstein spaces with Multi-Marginal Optimal Transport (MMOT), we derive an algorithm called MMOT-EFM. This algorithm controls the rate of change of the generated conditional distribution. Our proposed method outperforms existing methods in molecular generation tasks where conditions are sparsely observed.

1 INTRODUCTION

Conditional generative modeling (CGM), which involves approximating a conditional probability distribution $p(x | c)$ of data x given condition c , holds great promise for generating new, previously non-existent data across a wide range of representations. Currently, CGM is pivotal in generating images, videos (Rombach et al., 2021; Saharia et al., 2022a;b; Voleti, 2023), and text (Li et al., 2022; Strudel et al., 2022; Gao et al., 2024), but it is also expected to be applied to scientific computing, such as molecular generation (Kang & Cho, 2019) and physical simulations (Huang et al., 2024; Gebhard et al., 2023).

One of the key challenges of applying CGM in scientific fields is the sparsity of available data conditions. This sparsity necessitates extrapolating beyond the observed conditions (Lee et al., 2023). An important example of scientific applications is molecular generation—imagine that you wish to discover a new molecule x_{desired} with a desired chemical property c_{desired} , for which no molecular data may be available. Here, we have only observed a limited number of properties c_{obs} , which may be very sparse and require difficult extrapolation. This sparsity issue is more apparent when the condition or property is multi-dimensional.

In contrast, recent deep generative models for CGM have been designed mainly for situations where the conditions are densely observed. Consider the example of methods (Ding et al., 2021; Zhao et al., 2024; Ding et al., 2024) based on Vicinal risk minimization (VRM) by Chapelle et al. (2000). In VRM, the observed conditions c_{obs} are augmented with Gaussian noise $w_c \sim \mathcal{N}(0, I)$, and the generative model is trained so that the unknown conditional distribution $p(x | c_{\text{obs}} + w_c)$ becomes close to the known distribution $p(x | c_{\text{obs}})$. Thus, if we can only observe two conditions c_{obs}^1 and

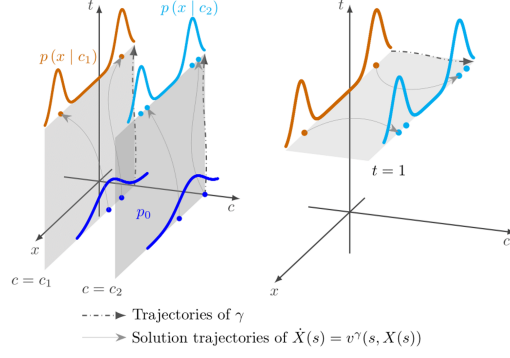
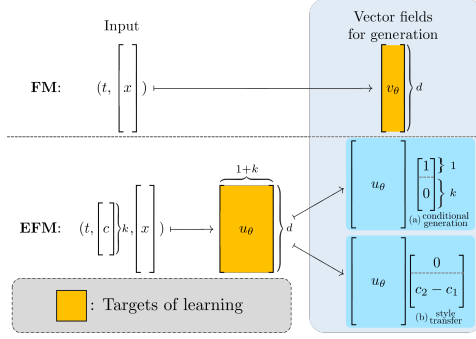


Figure 2: Visualization of the flow for (a) conditional generation along γ^{c_1} and γ^{c_2} (Algorithm 2), and (b) style transfer along $\gamma^{c_1 \rightarrow c_2}$ (Algorithm 3).

c_{obs}^2 , which are somewhat distant from each other, then we cannot introduce any inductive bias into the interpolated or extrapolated condition c_{desired} . As a result, the accuracy of the generation of data given c_{desired} would not improve. Indeed, Figure 4b will show another example where the quality of the generation at $c = c_{\text{desired}}$ deteriorates compared to $c = c_{\text{obs}}$ if no bias is introduced.

We expect that one of the hopes to overcome this difficulty is dynamical generative models, including diffusion models (Song et al., 2021; Ho et al., 2020) and, in particular, the simplest of these—Flow matching (FM) (Liu et al., 2023; Lipman et al., 2023; Albergo et al., 2023). FM itself is the method of generative modeling to approximate a probability distribution $p(x)$. In FM, two probability distributions are *gradually* deformed by flows induced by ordinary differential equations (ODEs). This deformation makes it possible to formulate the learning of the generative model as an estimation of the “vector field”, i.e., the way in which the ODE infinitesimally transformed the data. In particular, the methods based on FM stabilize the learning of vector fields, making it possible to generate a variety of data representations, including images (Esser et al., 2024), text (Hu et al., 2024), audio (Le et al., 2023), DNA (Stark et al., 2024), and molecules (Song et al., 2023; Miller et al., 2024).

This paper proposes the framework of *Extended Flow Matching (EFM)*, which realizes a smooth change of distributions for departure from the observed conditions to prevent unfavorable change observed in the existing FM methods. This framework introduces a flow with respect to the conditional axis c , and allows the conditional distribution $p(x | c)$ generated by EFM to also *gradually* change with respect to c . More specifically, we apply the extended Monge–Kantorovich theory introduced by Brenier (2003) to conditional generative models. This leads to a framework for learning *matrix fields* in a generalized continuity equation instead of vector fields in the continuity equations in FM.

Furthermore, by combining the concept of Dirichlet energy on Wasserstein spaces introduced by Lavanant (2019) with Multi-Marginal Optimal Transport (MMOT), we can derive an algorithm called *MMOT-EFM* that controls the rate of change of the generated conditional distribution. In addition, our proposed method is shown to outperform existing methods in the task of molecular generation in situations where conditions are sparsely observed.

NOTATION

Let us use \cdot to denote a placeholder, $\|\cdot\|$ to denote the Euclidean norm, and $0_k := (0, \dots, 0)^\top \in \mathbb{R}^k$ to denote the zero vector. We denote by $\mathcal{P}(M)$ the space of probability distributions on a metric space M , and denote by $\delta_x \in \mathcal{P}(M)$ the delta distribution supported on $x \in M$. For a distribution $\mu \in \mathcal{P}(M)$ on M and a vector-valued function f on M , we denote by $\mathbb{E}_{X \sim \mu}[f(X)]$ the expectation of a random variable $f(X)$, where $X \sim \mu$ is a random variable following μ .

We also denote $I := [0, 1]$ and $[m : n] := \{m, m + 1, \dots, n\}$ for $m, n \in \mathbb{N}$ such that $m < n$. For a function g on I , we write $\dot{g}(t)$ for the derivative $\frac{dg}{dt}(t)$ with respect to time $t \in I$. Further,

we let $D \subset \mathbb{R}^d$ be the data space. For any subscript ξ , we will denote by p_ξ the density of a probability distribution μ_ξ on $D \subset \mathbb{R}^d$, i.e., $\mu_\xi(dx) = p_\xi(x)dx$ in a measure-theoretic notation. In the following mathematical discussion, we will assume that any probability distribution has a density, but this assumption is superficial and is used only for simplicity of explanation.

2 PRELIMINARIES

In this section, we present Flow Matching (Lipman et al., 2023) and its variant, OT-CFM (Pooladian et al., 2023; Tong et al., 2023b), through the lens of Monge–Kantorovich theory to motivate the definition of EFM.

2.1 FLOW MATCHING (FM)

Continuity Equation: As a method of generative modeling, the goal of FM is to learn a map that transforms a source distribution to a target distribution in the form of $\mu: [0, 1] \rightarrow \mathcal{P}(D)$, where D is the space of dataset. Instead of learning μ directly, flow matching as a method learns a vector field $v: [0, 1] \times D \rightarrow \mathbb{R}^d$ such that the *continuity equation* (CE)

$$\partial_t p_t(x) + \operatorname{div}_x(p_t(x)v(t, x)) = 0 \quad (2.1)$$

holds with respect to the density p_t of μ_t , and we use this v for the sample generation.

Inference: $X_1 \sim \mu_1$ can be sampled by solving the ODE with $\dot{X}(t) = v(t, X(t))$, $X(0) \sim p_0$.

2.2 OT-CFM

OT-CFM, which has been proposed to use optimal transport for constructing the vector field, can be interpreted as a method of minimizing the Dirichlet energy, or the energy of transport for μ conditional to the boundary condition $\mu_0 = \mu_{\text{source}}, \mu_1 = \mu_{\text{target}}$. Specifically, we will show that a straight line in the construction of OT-CFM can be regarded as a minimizer of the Dirichlet energy.

Objective energy: Formerly, Dirichlet or the kinetic energy of the curve μ can be written as

$$\operatorname{Dir}(\mu) := \inf_{v: I \times D \rightarrow \mathbb{R}^d} \left\{ \frac{1}{2} \iint_{I \times D} \|v(t, x)\|^2 p_t(x) dx dt \mid \text{The pair } (p, v) \text{ satisfies (2.1)} \right\}. \quad (2.2)$$

Objective function: To derive the algorithm used in OT-CFM, let us present $\mu = \mu^Q$ from a random path as

$$\mu^Q: I \ni t \mapsto \mathbb{E}_{\psi \sim Q}[\mu_t^\psi] \in \mathcal{P}(D), \quad (2.3)$$

where $\mu_t^\psi = \delta_{\psi(t)}$ is the point-mass distribution at $\psi(t)$ and Q is a distribution over a space $H(I; D) := \{\psi: I \rightarrow D \mid \psi \text{ is differentiable}\}$ of paths that maps time t to an instance $x \in D$. The minimizer Q^* of the problem $\inf_Q \operatorname{Dir}(\mu^Q)$ subject to $\mu_0^Q = \mu_{\text{source}}, \mu_1^Q = \mu_{\text{target}}$ turns out to be concentrated on the set of “straight lines” $\psi(t \mid x_1, x_2) = tx_2 + (1-t)x_1$ between joint samples (x_1, x_2) from the target and the source. By Brenier (2003); Ambrosio et al. (2008), the function $D \times D \ni (x_1, x_2) \mapsto \psi(\cdot \mid x_1, x_2) \in H(I; D)$ allows a *parametrization* of Q with the optimal transport plan π with marginals μ_{source} and μ_{target} . This would allow us to write $\|\dot{\psi}(t \mid x_1, x_2)\|^2 = \|x_1 - x_2\|^2$ for the optimal Q^* . This would reduce the optimization with respect to Q to the classic optimal transport problem for the joint probability π with cost $c(x, y) = \|x - y\|^2$. In OT-CFM, this is approximated through batches. Following the same logic as in Kerrigan et al. (2024a), or our later theorem (Theorem 3.4), the vector field v which generates μ_{Q^*} via CE can be obtained as the minimizer of

$$\mathbb{E}_{\psi \sim Q^*, t \sim \operatorname{Unif}(I)} [\|v(t, \psi(t)) - \dot{\psi}(t)\|^2] = \mathbb{E}_{(x_1, x_2) \sim \pi^*, t \sim \operatorname{Unif}(I)} [\|v(t, \psi(t)) - \dot{\psi}(t \mid x_1, x_2)\|^2]. \quad (2.4)$$

This derives the learning of v through a neural network v_θ as shown in Algorithm 5. Indeed, Dirichlet energy that OT-CFM is aiming to minimize is a form of inductive bias regarding the continuity of the *generation* process with respect to time t .

In naive application of OT-CFM to conditional generation, $\psi(t)$ may be replaced with $\psi(t, c)$ given condition c . Note that the energy in this situation only considers $\|\partial_t \psi(t, c)\|^2$, which is different from our EFM discussed in Section 3.

3 THEORY OF EFM

In this section, we extend the standard FM theory to consider conditional probability with conditions c within a bounded domain $\Omega \subset \mathbb{R}^k$. Let $p_c(x) := p(x | c)$ be the unknown target conditional probability density, and let $p_{0,c}(x) := p_0(x | c)$ be a user-chosen tractable conditional density given $c = (c^i)_{i \in [1:k]} = (c^1, \dots, c^k) \in \Omega$, such as normal distributions with mean and variance parameterized by c . We will use the notation in the previous section, that is, we will denote by μ_c and $\mu_{0,c}$ the distribution of the probability density function p_c and $p_{0,c}$, respectively.

3.1 EXTENSION OF FM

We will present this subsection in parallel with Subsection 2.1.

Generalized Continuity Equation: We directly extend the interpretation of FM by extending the domain of ψ in (2.3) from I to $I \times \Omega$, where Ω is the space of conditions. For brevity, instead of using explicit $I \times \Omega$, we would like to use a general bounded domain Ξ in Euclidean space as an analog of Ω of the previous section and analogously set the goal of EFM to the learning of $\mu: \Xi \rightarrow \mathcal{P}(D)$. Now, just like FM, instead of learning μ directly, EFM aims to learn a *matrix* field $u: \Xi \times D \rightarrow \mathbb{R}^{d \times \dim \Xi}$ such that *generalized CE* (Brenier, 2003; Lavenant, 2019)

$$\nabla_\xi p_\xi(x) + \operatorname{div}_x(p_\xi(x)u(\xi, x)) = 0 \quad ((\xi, x) \in \Xi \times D) \quad (3.1)$$

holds for the density p_ξ of μ_ξ . Here, div is an extended divergence operator, see Appendix (A.1).

Inference: Inference based on the matrix field u is slightly more complicated than in FM, which provides a single vector field to integrate the ODE. Various tasks can be solved with solely the matrix field, including the typical cases of generation and transfer. For $\Xi = I \times \Omega$, the generation given condition c will be performed by transforming $\mu_{0,c} \rightarrow \mu_{1,c}$, and the transfer from c to c' by transforming $\mu_{1,c} \rightarrow \mu_{1,c'}$. Both are performed by integrating the matrix field along the path in $I \times \Omega$. More precisely, the following result justifies our use of the matrix field u in (3.1) to achieve the goal of conditional generative modeling:

Proposition 3.1 (GCE generates γ -induced CE). *Let $\mu: \Xi \rightarrow \mathcal{P}(D)$ and $u: \Xi \times D \rightarrow \mathbb{R}^{d \times \dim \Xi}$ be a probability path and a matrix field, respectively, that satisfy (3.1). Then, for any differentiable path $\gamma: I \rightarrow \Xi$, the γ -induced probability path $\mu^\gamma := \mu \circ \gamma$ and the γ -induced vector field $v^\gamma: I \times D \ni (s, x) \mapsto u(\gamma(s), x)\dot{\gamma}(s) \in \mathbb{R}^d$ satisfy the continuity equation, i.e., the density p^γ of μ^γ and v^γ satisfy $\partial_s p_s^\gamma(x) + \operatorname{div}_x(p_s^\gamma(x)v^\gamma(s, x)) = 0$.*

The rigorous version of Proposition 3.1 is given in Proposition A.2 in the Appendix. Proposition 3.1 shows that the flow on D corresponding to an arbitrary probability path on $\{\mu_\xi \in \mathcal{P}(D) \mid \xi \in \Xi\}$ can be constructed from the γ -induced vector field obtained from multiplying the matrix u to the vector $\dot{\gamma}$. Thus, once the matrix field u is obtained, the desired vector field v^γ is to be calibrated by choosing an appropriate γ that suits the purpose of choice. When the pair of p_ξ and u_ξ satisfies GCE (3.1), the designs of γ in the following two examples possess significant practical importance (See Figure 1 and Figure 2):

Example 3.2 (Conditional generation). When the goal is to sample from the unknown conditional distribution μ_{c_*} given condition $c_* \in \Omega$, we can choose $\gamma^{c_*}: I \rightarrow I \times \Omega$ such that $\gamma^{c_*}(1) = (1, c_*)$; typically, we can set $\gamma^{c_*}(s) = (s, c_*)$ for $s \in I$. Then, by virtue of Proposition 3.1 and the continuity equation (2.1), we only need to compute the flow ϕ by solving the ODE

$$\begin{cases} \dot{\phi}_s(x_0) = u(s, c_*, \phi_s(x_0)) \begin{bmatrix} 1 \\ 0_k \end{bmatrix} \quad (s \in I), \\ x_0 \sim \mu_{0,c_*}, \end{cases}$$

and obtain samples $\phi_1(x_0)$ from $\mu_{1,c_*} = \mu_{c_*}$. The trajectories in the front and rear plane of (a) in Figure 2 respectively represent the flows corresponding to this example with $c_* = c_1$ and $c_* = c_2$.

Example 3.3 (Style transfer). When the goal is to transform a sample generated from μ_{c_1} to a sample of another distribution μ_{c_2} given $c_2 \in \Omega$, we may choose $\gamma^{c_1 \rightarrow c_2}: I \rightarrow I \times \Omega$ satisfying $\gamma^{c_1 \rightarrow c_2}(0) = (1, c_1)$ and $\gamma^{c_1 \rightarrow c_2}(1) = (1, c_2)$. For example, we can set $\gamma^{c_1 \rightarrow c_2}(s) = (1, (1-s)c_1 + sc_2)$ for $s \in I$. In this case, we only need to solve the ODE

$$\begin{cases} \dot{\phi}_s(x_0) = u(1, \gamma^{c_1 \rightarrow c_2}(s), \phi_s(x_0)) \begin{bmatrix} 0 \\ c_2 - c_1 \end{bmatrix} (s \in I), \\ x_0 \sim \mu_{c_1}. \end{cases}$$

The solution trajectories in (b) in [Figure 2](#) represent the flows corresponding to this style transfer.

3.2 OBJECTIVE ENERGY AND MMOT-EFM

Now we extend the arguments in [Subsection 2.2](#) to EFM.

Objective energy: Just like in [Subsection 2.2](#), we use the representation of μ as (2.3) through a distribution Q over a space $H(\Xi; D)$ of differentiable maps ψ from Ξ to D . Now, the construction of EFM allows us to introduce inductive bias regarding a property of $\psi: \Xi \rightarrow D$ and hence how μ behaves with respect to ξ . In particular, if a given energy \mathcal{E} with respect to μ^ψ is convex, then by Jensen’s inequality we can bound $\mathcal{E}(\mu)$ from above by $\mathbb{E}_{\psi \sim Q}[\mathcal{E}(\mu^\psi)]$. Please also see [Propositions B.1](#) and [B.2](#) for more precise statements of these results.

In MMOT-EFM, we consider the case in which \mathcal{E} is the following generalization of the Dirichlet energy (2.2). According to [Lavenant \(2019\)](#), a generalization of Dirichlet energy of a function $\mu: \Xi \rightarrow \mathcal{P}(D)$ is given by

$$\text{Dir}(\mu) := \inf_{u: \Xi \times D \rightarrow \mathbb{R}^d} \left\{ \frac{1}{2} \iint_{\Xi \times D} \|u(\xi, x)\|^2 p_\xi(x) dx d\xi \mid \text{The pair } (p, u) \text{ satisfies (3.1)} \right\}, \quad (3.2)$$

where p_ξ is the density of μ_ξ . This energy is of great practical importance because it also measures how large μ changes with respect to ξ .

Objective function: Unfortunately, unlike in the case of OT, the energy-minimizing μ that can be written as $\mu = \mu^Q := \mathbb{E}_{\psi \sim Q}[\mu^\psi]$ is not necessarily achieved with Q concentrated on “straight paths”, or (flat) hyperplanes interpolating joint samples from $\{\mu_\xi\}$. Thus we choose to constrain the search of Q to a specific subspace \mathcal{F} of $H(\Xi; D)$, such as Reproducing Kernel Hilbert Space (RKHS). In this search, we also require Q to satisfy the boundary condition (BC) that

$$\mathbb{E}_{\psi \sim Q} [\delta_{\psi(\xi)}] = \mu_\xi (\xi \in A), \quad (3.3)$$

where $A \subset \Xi$ is a finite set for which $\mu_\xi (\xi \in A)$ is either known or observed. Instead of (3.3), suppose $\mathbf{x}_A := (x_\xi)_{\xi \in A}$ for $A \subset \Xi$ is a joint sample with $x_\xi \sim \mu_\xi$. Then, let $\phi: D^{|A|} \rightarrow \mathcal{F}$ be the function-valued mapping, which returns the function $\Xi \ni \xi \mapsto \phi(\xi \mid \mathbf{x}_A) \in D$ defined by the regression

$$\phi(\cdot \mid \mathbf{x}_A) \in \arg \min_{f \in \mathcal{F}} \sum_{\xi \in A} \|f(\xi) - x_\xi\|^2, \quad (3.4)$$

i.e., $\phi(\cdot \mid \mathbf{x}_A)$ satisfies $\sum_{\xi \in A} \|\phi(\xi \mid \mathbf{x}_A) - x_\xi\|^2 = \min_{f \in \mathcal{F}} \sum_{\xi \in A} \|f(\xi) - x_\xi\|^2$ for each $\mathbf{x}_A \in D^{|A|}$. For a joint distribution on π on $D^{|A|}$, the parametrization $Q \rightarrow \phi_\# \pi$ of random paths allows us to bound the energy from above in the following way:

$$\inf_Q \text{Dir}(\mu^Q) \leq \inf_Q \iint_{H(\Xi; D) \times \Xi} \|\nabla_\xi \psi(\xi)\|^2 Q(d\psi) d\xi \leq \inf_\pi \iint_{D^{|A|} \times \Xi} \|\nabla_\xi \phi(\xi \mid \mathbf{x}_A)\|^2 \pi(d\mathbf{x}_A) d\xi.$$

Now observe that the upper bound is the form of a marginal optimal transport problem about π with marginals μ_A and $c(\mathbf{x}_A) = \int_\Xi \|\nabla_\xi \phi(\xi \mid \mathbf{x}_A)\|^2 d\xi$, whose solution π^* can be approximated with batch as in the OT-CFM case. See [Table 1](#) for the parallel correspondence between MMOT-EFM and OT-CFM.

Table 1: Constructions of $\psi: [0, 1] \rightarrow D$ and $\bar{\psi}: \Omega \rightarrow D$ and π in OT-CFM and MMOT-EFM. Note that they agree when \mathcal{F} is a set of linear functions from Ω to D and when $\Omega = [0, 1] \subset \mathbb{R}$.

	OT-CFM	MMOT-EFM
Interpolator	$\psi(t \mid x, y) = tx + (1 - t)y$	$\bar{\psi}(\cdot \mid \mathbf{x} = (x_i)_i) \in \arg \min_{\phi \in \mathcal{F}} \sum_i \ \phi(c_i) - x_i\ ^2$
Cost	$\iint_{[0,1] \times D^2} \ \psi(t \mid x, y)\ ^2 dt \pi(dx, dy)$ $(\quad = \iint_{D^2} \ x - y\ ^2 \pi(dx, dy))$	$\iint_{\Omega \times D^{ \mathcal{C} }} \ \nabla_c \bar{\psi}(c \mid \mathbf{x})\ ^2 dc \pi(d\mathbf{x})$

Similarly to (2.4), Theorem 3.4 which is provided below allows us to train u which generates μ^{Q^*} via (3.1) as the minimizer of

$$\mathbb{E}_{\psi \sim Q^*, \xi \sim \text{Unif}(\Xi)} [\|u(\xi, \psi(\xi)) - \nabla_\xi \psi(\xi)\|^2] = \mathbb{E}_{\mathbf{x}_A \sim \pi^*, \xi \sim \text{Unif}(\Xi)} [\|u(\xi, \psi(\xi)) - \nabla_\xi \phi(\xi \mid \mathbf{x}_A)\|^2] \quad (3.5)$$

which we would use as the objective function of MMOT-EFM.

Theorem 3.4. Assume we have a random path $\psi \sim Q \in \mathcal{P}(H(\Xi; D))$ that satisfies (3.3) and let $\mu_{t,c} = \mathbb{E}_{\psi \sim Q} [\delta_{\psi(\xi)}]$ for $\xi \in \Xi$. For neural networks u_θ , set

$$\mathcal{L}'(\theta) = \int_{\Xi} \mathbb{E}_{\psi \sim Q} [\|u_\theta(\xi, \psi(\xi)) - \nabla_\xi \psi(\xi)\|^2] d\xi. \quad (3.6)$$

If there exists a matrix field $u: \Xi \times D \rightarrow \mathbb{R}^{d \times (1+k)}$ satisfying (3.1), then it follows that $\nabla_\theta \mathcal{L}(\theta) = \nabla_\theta \mathcal{L}'(\theta)$ for $\theta \in \mathbb{R}^p$. Here, we set $\mathcal{L}(\theta) := \int_{\Xi} \mathbb{E}_{x \sim \mu_\xi} [\|(u_\theta - u)(\xi, x)\|^2] d\xi$.

See Lemma A.4 in the Appendix.

4 TRAINING ALGORITHM

In this section, we leverage the EFM theory of Section 3 to construct an algorithm for learning u_θ in Proposition 3.1, which can be used for conditional generation tasks as well as for style transfer. We summarize the training algorithm in Algorithms 1 and 8.

Because EFM is a direct extension of FM, our algorithm roughly follows the same line of procedures as that of FM (Algorithm 5): (a) sampling data, (b) constructing the supervisory signal $\nabla \psi$, and (c) updating the network by averaged loss. However, in our algorithm, the domain of ψ is $I \times \Omega$ as opposed to just I . We developed our algorithm so that, when it is applied to the unconditional case, the trained model agrees with FM. Although the general EFM, as opposed to MMOT-EFM, does not necessarily need to parametrize Q with respect to joint distribution π , in this paper we focus on the procedure that uses the joint distribution π and ψ in the form of (3.4) and (3.5).

Step 1 Sampling from Datasets: Our objective begins from the sampling of ψ , whose Jacobian serves as the supervisory signal in the objective (3.5). In order to sample ψ , we construct Q from a joint distribution π defined over D^{2N_c} with marginals that are approximately $(\mu_{t,c})_{t \in \{0,1\}, c \in C_0}$. To this end, we begin by randomly choosing a subset $C_0 := \{c_i\}_{i=1}^{N_c}$ from C so that C_0 consists of close points. We then sample a batch $B_{0,c}$ from $\mu_{0,c}$ and $B_{1,c}$ from D_c for each $c \in C_0$. For the reason we describe at the end of this section, we chose $\mu_{0,c} = \text{Law}(R(c) + z)$ with z being a common Gaussian component, and $R: \Omega \rightarrow D$ is regressed from $\{(c_i, \text{Mean}[D_{c_i}])\}_i$ by a linear map. We choose this option because it theoretically aids us in reducing $\text{Dir}(\mu)$ (See Proposition B.2).

Step 2 Constructing the supervisory paths: Given the samples $B = (B_{t,c})_{t \in \{0,1\}, c \in C_0}$, we sample $(x_{t,c})_{c \in C_0, t \in \{0,1\}}$ from a joint distribution π over D^{2N_c} with support on B . In MMOT-EFM,

Algorithm 1 Algorithm of EFM

Input: Conditions $C \subset \Omega$, set of datasets $D_c \subset D$ ($c \in C$), network $u_\theta: I \times \Omega \times D \rightarrow \mathbb{R}^{d \times (1+k)}$, source distributions $p_0(\cdot | c)$ ($c \in C$)

Return: $\theta \in \mathbb{R}^p$

- 1: **for** each iteration **do**
- # Step 1: Sample
- 2: Sample C_0 from C , $B_{0,c}$ from $p_0(\cdot | c)$ and $B_{1,c}$ from D_c ($c \in C_0$). Put $B^0 := \{B_{0,c}\}_{c \in C_0}$, $B^1 := \{B_{1,c}\}_{c \in C_0}$
- # Step 2: Construct $\psi: I \times \Omega \rightarrow D$
- 3: Construct a transport plan π among B^0 and B^1 #Section 4
- 4: Sample $(x_{t,c})_{t,c} \sim \pi$
- 5: Define $\psi: I \times \Omega \rightarrow D$ s.t. (4.1)
- 6: Sample $t \sim \text{Unif}(I)$, $c \sim \text{Unif}(\text{ConvHull}(C_0))$
- 7: Compute

$$\psi_{t,c} := \psi(t, c)$$

$$\nabla \psi_{t,c} := \nabla_{t,c} \psi(t, c)$$
- 8: Update θ by $\nabla_\theta \|u_\theta(t, c, \psi_{t,c}) - \nabla \psi_{t,c}\|^2$
- 9: **end for**

as an internal step, we train the joint distribution π with $c(\mathbf{x}_A) = \int_{I \times \Omega} \|\nabla_{t,c} \phi(t, c | \mathbf{x}_A)\|^2 dt dc$ with ϕ solved analytically for (3.4) with $\Xi := I \times \Omega$, by e.g., Kernel Regression, Linear regression. When possible, the regression function may be chosen to reflect the prior knowledge of the metrics on Ω by extending the philosophy of Chen & Lipman (2024) to the space of conditions. In practice, however, the computational cost of MMOT scales exponentially with the number of marginals, so we optimize the joint distributions over $B_1 = (B_{1,c})_{c \in C_0}$ only and couple the analogous B_0 to B_1 via the usual optimal transport. Please see Appendix D.3 for a more detailed sampling procedure. Now, given a joint sample $(x_{t,c})_{c \in C_0, t \in \{0,1\}}$, we construct ψ as

$$\psi(t, c | x_{0,c}, \mathbf{x}_{C_0}) = (1-t)x_{0,c} + t\bar{\psi}(c | \mathbf{x}_{C_0}) \quad (4.1)$$

where $\bar{\psi}(c | \mathbf{x}_{C_0})$ is the solution of the kernel regression problem for the map $T: \mathbb{R}^k \ni c \mapsto x_{1,c} \in \mathbb{R}^d$ with any choice of kernel on \mathbb{R}^k . Note that this construction of ψ satisfies the boundary condition (3.3) with $A = \{0, 1\} \times C_0$, and generalizes the ψ used in OT-CFM.

Step 3 Learning the matrix fields: Thanks to the result of Theorem 3.4, we may train $u_\theta: I \times \Omega \rightarrow \mathbb{R}^{d \times (1+k)}$ via the loss function being the Monte Carlo approximation of (3.6).

5 INFERENCE METHOD

The sampling procedures for style transfer and conditional generation respectively follow Example 3.3 and Example 3.2. For the task of style transfer from c_0 to c_* , we use the flow along the path $\mu_{1,c_0} \rightarrow \mu_{1,c_*}$. For the task of conditional generation with target condition c_* , we use the flow along $\mu_{0,c_*} \rightarrow p_{\mu,c_*}$. See Algorithms 2 and 3 for the pseudo-codes. When generating a sample for $c^* \notin C$, the source distribution μ_{0,c^*} is constructed by $R(c^*) + \mathcal{N}(0, I)$ where R is given as in training.

6 RELATED WORKS

Guidance-based methods: Since the debut of Lipman et al. (2023), several studies have explored ways to formalize the application of flow-based models to conditional generation tasks. Some works (Dao et al., 2023; Zheng et al., 2023) take the approach of parametrizing the vector field v with the conditional value c together with the so-called guidance scale $\omega \in \mathbb{R}$ in the form of $v(t, c, x) = \omega v_t(x | \emptyset) + (1 - \omega)v_t(x | c)$, which is inspired by the classifier-free guidance scheme of Ho & Salimans (2022). Zheng et al. (2023) in particular has shown that if $v_t(x | c)$ in this expression well-approximates the conditional score $\nabla \log p(x | c)$, then with the appropriate choice of ω , $v_t(x, c)$ does correspond to the sequence of probability distributions beginning from the standard Gaussian

Algorithm 2 Generation using the matrix field u_θ

Input: Trained u_θ , source distribution $p_{0,0}$, target condition c_*
Return: A sample x_1 from $p(\cdot | c_*)$
 Sample z from source distribution $p_{0,0}$
 Solve the regression problem $R: c \mapsto \text{Mean}[D_c]$ on C
 Set $x_{0,c} = z + R(c)$
 Return $\text{ODEsolve}(x_{0,c}, u_\theta(\cdot, c, \cdot)) \big[\frac{1}{0_k} \big]$

Algorithm 3 Transfer using the matrix field u_θ

Input: Trained Network u_θ , source sample $x_0 \sim p_{1,c_1}$ with condition label c_1 , target condition c_2
Return: A sample x_2 from $p(\cdot | c_2)$
 Return
 $\text{ODEsolve}(x_0, u_\theta(1, \gamma^{c_1 \rightarrow c_2}(\cdot), \cdot) \big[\frac{0}{c_2 - c_1} \big])$
 # $\gamma^{c_1 \rightarrow c_2}$ is defined in Example 3.3

distribution and ending at the target distribution. The success of this scheme hinges on the quality of the approximation of the conditional score, and it is reported (Lipman et al., 2023) that in image applications, a guidance scale with a range from 1.2 to 1.3 yields competitive performance in terms of FID. Meanwhile, (Hu et al., 2023) takes the approach of creating a guidance vector by the average of $v_t(x_{c_{\text{targets}}}) - v_t(x_{c_{\text{others}}})$. Like the naive application of OT-CFM to a conditional generation that simply concatenates the conditional value to the input of the network modeling the vector field, however, these approaches do not allow the user to control the continuity of generated μ_c with respect to c , except through the black box architecture of the network modeling v .

Unlike these approaches, EFM constructs the flow of generation for an arbitrary condition $c \in \Omega$ through the matrix field $u: I \times \Omega \times D \rightarrow \mathbb{R}^{d \times (1+k)}$ which solves GCE, or the system of continuity equations defined over $I \times \Omega$, and one can introduce an inductive bias to the continuity of μ_c with respect to c through the design of the distribution Q of ψ used in the objective function. The Dirichlet energy that we use in the demonstration of EFM is akin to the control of the Lipschitz constant for ψ and hence μ , except that it also comes with the boundary condition to assure the generation of the conditional distributions used at the time of the training. Also, when u is trained with the random conditional paths with appropriate boundary conditions, our EFM theory in Section 3 guarantees that the flow ϕ^{γ^c} in Example 3.2 transforms the source distribution to the target conditional distribution whenever c is a condition used in the training.

Dynamical generative models (DGMs) for CGM: In addition to the VRM-based method mentioned in Section 1, there are two other methods: COT-FM (Kerrigan et al., 2024b) and Bayesian-FM (Chemseddine et al., 2024), both based on Conditional Optimal Transport (Hosseini et al., 2024). These methods rely on the relatively weak assumption that the map of conditional distributions $c \mapsto p(x | c)$ is measurable with respect to c . In contrast, the learning algorithm of EFM is designed under the assumption that $p(x|c)$ is piecewise continuous with respect to c . This distinction arises because the former addresses situations where high-dimensional conditions, such as inverse problems of PDEs, can be densely observed, while the latter addresses scenarios where relatively low-dimensional conditions, such as molecular generation, can be sparsely observed.

The authors in (Wildberger et al., 2023; Atanackovic et al., 2024) developed FM-based models to estimate the posterior distribution when the prior distribution $p(c)$ of conditions is known. In contrast, our approach tackles situations where the conditions can only be sparsely observed, and the prior distribution is unknown.

Energy principles in DGMs: We shall also mention the family of methods based on the Schrödinger bridge, which also aims to interpolate between an arbitrary pair of distribution (Tong et al., 2023a). This direction can be regarded as the problem of solving the continuity equation while minimizing the regularized energy of user’s choice (Koshizuka & Sato, 2022) in the generation process. (Kim et al., 2023) also uses Wasserstein Barycenter for distributional interpolation. Stochastic interpolants (Albergo et al., 2023) learns a model that is similar to generalized geodesics Appendix F and it aims to optimize the path in the space of conditions with respect to kinetic energy. This approach, however, neither simultaneously models the generation process along I nor formulates the energy with respect to $\mu: \Omega \rightarrow \mathcal{P}(D)$ itself.

7 EXPERIMENTS

We conducted the following experiments to investigate EFM in applications.

7.1 SYNTHETIC 2D POINT CLOUDS

We first demonstrate the performance of our method on a conditional distribution consisting of synthetic point clouds in a two-dimensional domain $D \subset \mathbb{R}^2$. Here, we consider the case where the space Ω of the condition is square, i.e., $\Omega = [0, 1]^2$, and train the model when only samples from the conditional distributions $p(\cdot | c)$ at the four corner points c of the square Ω can be observed, see Figure 5 in Appendix. We




Figure 3: Training distributions for Subsection 7.1

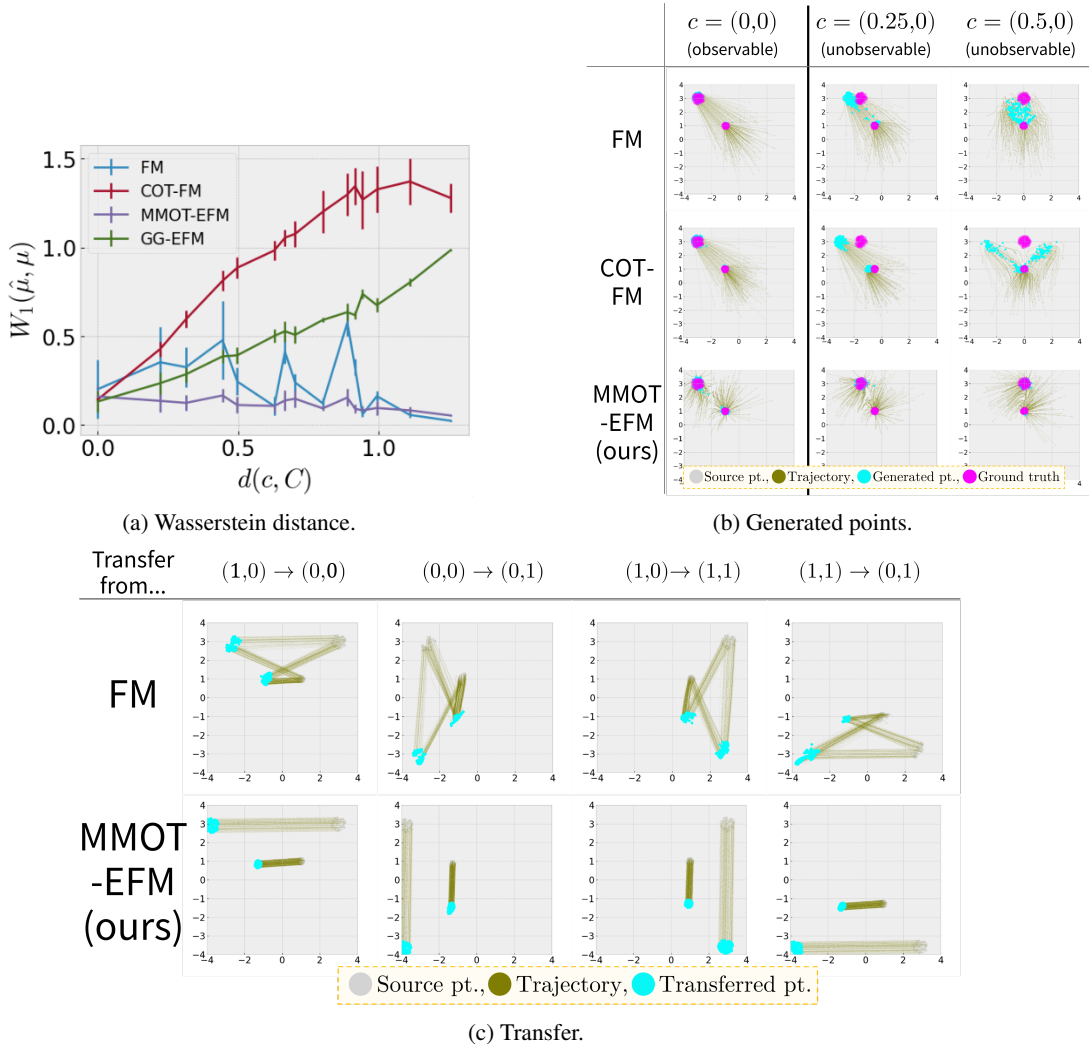


Figure 4: Results of Subsection 7.1

compared our method against COT-FM (Chemseddine et al., 2024; Kerrigan et al., 2024b), as well as OT-CFM and the EFM with the plan π which is constructed in the way of generalized geodesic, see Appendix F. See Figures 4b and 4c for the generation and transfer visualizations, and see Figure 4a for the error between GT and predicted distributions. Note that our method, MMOT-EFM, performs competitively with all its rivals in interpolation and generation tasks. Also note that the style transfer

with MMOT-EFM preserves the structure of the inner and outer clusters, just as mentioned in the introduction.

7.2 CONDITIONAL MOLECULAR GENERATION

In molecular design applications, multiple chemical properties are often required to be considered simultaneously. Most traditional molecular design methods combine all property requirements and their constraints into a single objective function. We applied MMOT-EFM to the conditional generation task for two simultaneous properties: (1) the number of rotatable bonds and (2) the number of hydrogen bond acceptors (HBA). The experimental setup is described in detail in [Appendix G](#).

We first trained a VAE model to encode molecular structures into a 32-dimensional latent space and then trained EFM to perform out-of-distribution conditional generation over this latent space.

We measure the MAE between the condition and actual value of the generated compounds. As shown in [Table 2](#), our method outperforms all baseline methods on the averaged MAE for out-of-distribution conditional generation.

Table 2: MMOT-EFM vs. baselines in conditional molecular generations.

Conditional Generation MAE	
FM	1.120 ± 0.142
COT-FM	0.966 ± 0.122
MMOT-EFM (ours)	0.918 ± 0.122

8 CONCLUSION

In this paper, we developed the theory of EFM, a direct extension of FM that learns the transformation of distributions along the conditional direction as well as along the direction of generation through the modeling of a matrix field instead of a vector field. EFM models how the distribution changes with respect to conditions in a more explicit form. We provide the mathematical theory of EFM together with MMOT-EFM, an extension of OT-CFM, with the aim of minimizing the average generation sensitivity with respect to continuous conditions and demonstrating its competitiveness. However, we shall note that our current algorithm is limited by the computational cost of MMOT, which grows exponentially with the number of conditional distributions to be used at each step of the algorithm ($|C_0|$). An advance in the efficient MMOT method or its alternative may significantly improve the scope of applications of EFM. The EFM theory is complementary to many powerful existing ideas, particularly through the design of ψ and Q , into which one may incorporate the structure of the space of conditions. Application to more complex datasets and incorporation of prior knowledge regarding the structure of Ω is an important future work. Finally, we note that our theory pertains to the generation of conditional distributions of unseen conditions and interpolation of distributions. We shall be aware that, without strong prior knowledge, the identification of unseen distribution is an ill-posed problem, and its solution also depends on the architectures and heuristics used therein, as well as the dataset used in the training.

REFERENCES

- Tagir Akhmetshin, Arkadii I. Lin, Daniyar Mazitov, Evgenii Ziaikin, Timur Madzhidov, and Alexandre Varnek. ZINC 250K data sets. 12 2021. doi: 10.6084/m9.figshare.17122427.v1. URL https://figshare.com/articles/dataset/ZINC_250K_data_sets/17122427.
- Michael S Albergo, Mark Goldstein, Nicholas M Boffi, Rajesh Ranganath, and Eric Vanden-Eijnden. Stochastic interpolants with data-dependent couplings. *arXiv preprint arXiv:2310.03725*, 2023.
- Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. *Gradient flows in metric spaces and in the space of probability measures*. Lectures in Mathematics ETH Zürich. Birkhäuser Verlag, Basel, 2 edition, 2008.

- Lazar Atanackovic, Xi Zhang, Brandon Amos, Mathieu Blanchette, Leo J Lee, Yoshua Bengio, Alexander Tong, and Kirill Neklyudov. Meta flow matching: Integrating vector fields on the wasserstein manifold. In *ICML 2024 Workshop on Geometry-grounded Representation Learning and Generative Modeling*, 2024. URL <https://openreview.net/forum?id=f9GsKvLdzs>.
- Yann Brenier. *Extended Monge–Kantorovich Theory*, pp. 91–121. Springer Berlin Heidelberg, Berlin, Heidelberg, 2003. ISBN 978-3-540-44857-0. doi: 10.1007/978-3-540-44857-0_4. URL https://doi.org/10.1007/978-3-540-44857-0_4.
- Olivier Chapelle, Jason Weston, Léon Bottou, and Vladimir Vapnik. Vicinal risk minimization. In *NIPS*, pp. 416–422. MIT Press, 2000.
- Jannis Chemseddine, Paul Hagemann, Christian Wald, and Gabriele Steidl. Conditional wasserstein distances with applications in bayesian ot flow matching, 2024.
- Ricky T. Q. Chen and Yaron Lipman. Flow matching on general geometries. In *ICLR*, 2024.
- Quan Dao, Hao Phung, Binh Nguyen, and Anh Tran. Flow matching in latent space. *arXiv preprint arXiv:2307.08698*, 2023.
- Xin Ding, Yongwei Wang, Zuheng Xu, William J. Welch, and Z. Jane Wang. Ccgan: Continuous conditional generative adversarial networks for image generation. In *ICLR*, 2021.
- Xin Ding, Yongwei Wang, Kao Zhang, and Z. Jane Wang. Ccdm: Continuous conditional diffusion models for image generation, 2024. URL <https://arxiv.org/abs/2405.03546>.
- Rick Durrett. *Probability: Theory and Examples*. Thomson, 2019.
- Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, Dustin Podell, Tim Dockhorn, Zion English, and Robin Rombach. Scaling rectified flow transformers for high-resolution image synthesis. In *ICML*. OpenReview.net, 2024.
- Jiaojiao Fan and David Alvarez-Melis. Generating synthetic datasets by interpolating along generalized geodesics. In *Uncertainty in Artificial Intelligence*, pp. 571–581. PMLR, 2023.
- Rémi Flamary, Nicolas Courty, Alexandre Gramfort, Mokhtar Z. Alaya, Aurélie Boisbunon, Stanislas Chambon, Laetitia Chapel, Adrien Corenflos, Kilian Fatras, Nemo Fournier, Léo Gautheron, Nathalie T.H. Gayraud, Hicham Janati, Alain Rakotomamonjy, Ievgen Redko, Antoine Rolet, Antony Schutz, Vivien Seguy, Danica J. Sutherland, Romain Tavenard, Alexander Tong, and Titouan Vayer. Pot: Python optimal transport. *Journal of Machine Learning Research*, 22(78): 1–8, 2021. URL <http://jmlr.org/papers/v22/20-451.html>.
- Zhujin Gao, Junliang Guo, Xu Tan, Yongxin Zhu, Fang Zhang, Jiang Bian, and Linli Xu. Empowering diffusion models on the embedding space for text generation, 2024. URL <https://arxiv.org/abs/2212.09412>.
- Timothy D Gebhard, Jonas Wildberger, Maximilian Dax, Daniel Angerhausen, Sascha P Quanz, and Bernhard Schölkopf. Inferring atmospheric properties of exoplanets with flow matching and neural importance sampling. *arXiv preprint arXiv:2312.08295*, 2023.
- Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020.
- Bamdad Hosseini, Alexander W. Hsu, and Amirhossein Taghvaei. Conditional optimal transport on function spaces, 2024.
- Vincent Tao Hu, David W Zhang, Meng Tang, Pascal Mettes, Deli Zhao, and Cees G. M. Snoek. Latent space editing in transformer-based flow matching. In *ICML Workshop on New Frontiers in Learning, Control, and Dynamical Systems*, 2023. URL <https://openreview.net/forum?id=Bi6E5rPtBa>.

- Vincent Tao Hu, Di Wu, Yuki Markus Asano, Pascal Mettes, Basura Fernando, Björn Ommer, and Cees Snoek. Flow matching for conditional text generation in a few sampling steps. In *EACL (2)*, pp. 380–392. Association for Computational Linguistics, 2024.
- Jiahe Huang, Guandao Yang, Zichen Wang, and Jeong Joon Park. DiffusionPDE: Generative PDE-solving under partial observation. In *ICML 2024 AI for Science Workshop*, 2024. URL <https://openreview.net/forum?id=8B9x6UW5pD>.
- Ryuichiro Ishitani, Toshiki Kataoka, and Kentaro Rikimaru. Molecular design method using a reversible tree representation of chemical compounds and deep reinforcement learning. *Journal of Chemical Information and Modeling*, 62(17):4032–4048, 2022.
- Noboru Isobe. A convergence result of a continuous model of deep learning via Łojasiewicz–Simon inequality. *arXiv preprint arXiv:2311.15365*, 2023.
- Wengong Jin, Regina Barzilay, and Tommi Jaakkola. Junction tree variational autoencoder for molecular graph generation. In *International conference on machine learning*, pp. 2323–2332. PMLR, 2018.
- Seokho Kang and Kyunghyun Cho. Conditional molecular design with deep generative models. *Journal of Chemical Information and Modeling*, 59(1):43–52, Jan 2019. ISSN 1549-9596. doi: 10.1021/acs.jcim.8b00263. URL <https://doi.org/10.1021/acs.jcim.8b00263>.
- Gavin Kerrigan, Giosue Miglierini, and Padhraic Smyth. Functional flow matching. In Sanjoy Dasgupta, Stephan Mandt, and Yingzhen Li (eds.), *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, pp. 3934–3942. PMLR, 02–04 May 2024a. URL <https://proceedings.mlr.press/v238/kerrigan24a.html>.
- Gavin Kerrigan, Giosue Miglierini, and Padhraic Smyth. Dynamic conditional optimal transport through simulation-free flows, 2024b.
- Young-geun Kim, Kyunghyun Lee, Youngwon Choi, Joong-Ho Won, and Myunghee Cho Paik. Wasserstein geodesic generator for conditional distributions. *arXiv preprint arXiv:2308.10145*, 2023.
- Takeshi Koshizuka and Issei Sato. Neural Lagrangian Schrödinger bridge: Diffusion modeling for population dynamics. In *The Eleventh International Conference on Learning Representations*, 2022.
- Hugo Lavenant. Harmonic mappings valued in the wasserstein space. *Journal of Functional Analysis*, 277(3):688–785, 2019. ISSN 0022-1236. doi: <https://doi.org/10.1016/j.jfa.2019.05.003>. URL <https://www.sciencedirect.com/science/article/pii/S0022123619301478>.
- Matthew Le, Apoorv Vyas, Bowen Shi, Brian Karrer, Leda Sari, Rashel Moritz, Mary Williamson, Vimal Manohar, Yossi Adi, Jay Mahadeokar, and Wei-Ning Hsu. Voicebox: Text-guided multilingual universal speech generation at scale. In *NeurIPS*, 2023.
- Seul Lee, Jaehyeong Jo, and Sung Ju Hwang. Exploring chemical space with score-based out-of-distribution generation. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 18872–18892. PMLR, 23–29 Jul 2023. URL <https://proceedings.mlr.press/v202/lee23f.html>.
- Xiang Lisa Li, John Thickstun, Ishaan Gulrajani, Percy Liang, and Tatsunori B. Hashimoto. Diffusion-lm improves controllable text generation. In *NeurIPS*, 2022.
- Tianyi Lin, Nhat Ho, Marco Cuturi, and Michael I. Jordan. On the complexity of approximating multimarginal optimal transport. *Journal of Machine Learning Research*, 23(65):1–43, 2022. URL <http://jmlr.org/papers/v23/19-843.html>.

- Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=PqvMRDCJT9t>.
- Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. In *ICLR*, 2023.
- Benjamin Kurt Miller, Ricky T. Q. Chen, Anuroop Sriram, and Brandon M Wood. Flowmm: Generating materials with riemannian flow matching, 2024. URL <https://arxiv.org/abs/2406.04713>.
- Caroline Moosmüller and Alexander Cloninger. Linear optimal transport embedding: Provable wasserstein classification for certain rigid transformations and perturbations. *arXiv preprint arXiv:2008.09165*, 2020.
- Zoe Piran, Michal Klein, James Thornton, and Marco Cuturi. Contrasting multiple representations with the multi-marginal matching gap. In *International conference on machine learning*, 2024.
- Aram-Alexandre Pooladian, Heli Ben-Hamu, Carles Domingo-Enrich, Brandon Amos, Yaron Lipman, and Ricky T. Q. Chen. Multisample flow matching: Straightening flows with minibatch couplings. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 28100–28127. PMLR, 23–29 Jul 2023. URL <https://proceedings.mlr.press/v202/pooladian23a.html>.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. *CoRR*, abs/2112.10752, 2021.
- Chitwan Saharia, William Chan, Huiwen Chang, Chris A. Lee, Jonathan Ho, Tim Salimans, David J. Fleet, and Mohammad Norouzi. Palette: Image-to-image diffusion models. In *SIGGRAPH (Conference Paper Track)*, pp. 15:1–15:10. ACM, 2022a.
- Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L. Denton, Seyed Kamyar Seyed Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, Jonathan Ho, David J. Fleet, and Mohammad Norouzi. Photorealistic text-to-image diffusion models with deep language understanding. In *NeurIPS*, 2022b.
- Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *ICLR*. OpenReview.net, 2021.
- Yuxuan Song, Jingjing Gong, Minkai Xu, Ziyao Cao, Yanyan Lan, Stefano Ermon, Hao Zhou, and Wei-Ying Ma. Equivariant flow matching with hybrid probability transport for 3d molecule generation. In *NeurIPS*, 2023.
- Hannes Stark, Bowen Jing, Chenyu Wang, Gabriele Corso, Bonnie Berger, Regina Barzilay, and Tommi Jaakkola. Dirichlet flow matching with applications to dna sequence design, 2024. URL <https://arxiv.org/abs/2402.05841>.
- Robin Strudel, Corentin Tallec, Florent Althé, Yilun Du, Yaroslav Ganin, Arthur Mensch, Will Grathwohl, Nikolay Savinov, Sander Dieleman, Laurent Sifre, and Rémi Leblond. Self-conditioned embedding diffusion for text generation, 2022. URL <https://arxiv.org/abs/2211.04236>.
- Alexander Tong, Nikolay Malkin, Kilian Fatras, Lazar Atanackovic, Yanlei Zhang, Guillaume Hugué, Guy Wolf, and Yoshua Bengio. Simulation-free schrödinger bridges via score and flow matching. *arXiv preprint 2307.03672*, 2023a.
- Alexander Tong, Nikolay Malkin, Guillaume Hugué, Yanlei Zhang, Jarrid Rector-Brooks, Kilian Fatras, Guy Wolf, and Yoshua Bengio. Improving and generalizing flow-based generative models with minibatch optimal transport. *arXiv preprint 2302.00482*, 2023b.
- Vikram Voleti. Conditional generative modeling for images, 3d animations, and video, 2023. URL <https://arxiv.org/abs/2310.13157>.

- Johannes von Lindheim. Approximative algorithms for multi-marginal optimal transport and free-support wasserstein barycenters, 2022. URL <https://arxiv.org/abs/2202.00954>.
- Jonas Wildberger, Maximilian Dax, Simon Buchholz, Stephen R. Green, Jakob H. Macke, and Bernhard Schölkopf. Flow matching for scalable simulation-based inference. In *NeurIPS*, 2023.
- Jason Yim, Andrew Campbell, Emile Mathieu, Andrew YK Foong, Michael Gastegger, José Jiménez-Luna, Sarah Lewis, Victor Garcia Satorras, Bastiaan S Veeling, Frank Noé, et al. Improved motif-scaffolding with SE(3) flow matching. *arXiv preprint arXiv:2401.04082*, 2024.
- Yanxuan Zhao, Peng Zhang, Guopeng Sun, Zhigong Yang, Jianqiang Chen, and Yueqing Wang. Ccdpm: A continuous conditional diffusion probabilistic model for inverse design. In *AAAI*, pp. 17033–17041. AAAI Press, 2024.
- 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.

A MATHEMATICAL DESCRIPTION OF EXTENDED FLOW MATCHING THEORY

Our aim is to sample from the unknown conditional distribution $\Omega \ni c \mapsto p(\bullet \mid c) \in \mathcal{P}(D)$. We extend the flow matching technique developed in (Lipman et al., 2023) for this aim. The technique evolves unconditional probability distributions $\mu_t \in \mathcal{P}(D)$, $t \in [0, 1]$ from a source distribution μ_0 (such as Gaussian $\mathcal{N}(\cdot)$) to a target distribution $\mu_1 \approx p^{\text{data}}$ by means of a continuity equation. We then introduce a generalized continuity equation that evolves conditional distributions $\mu_{t,c}$, $t \in [0, 1]$, $c \in \Omega$ from source distributions μ_0 to the target distributions $\mu_{t=1,c} \approx p^{\text{data}}(\bullet \mid c)$.

To realize this evolution, this section gives an example of how to construct a (at least approximate) solution of the generalized continuity equation and a design of the source distributions $\mu_{t=0,c}$, $c \in \Omega$.

A.1 NOTATIONS

- $\langle \bullet, \bullet \rangle$ is the standard inner product and $|\bullet| := \sqrt{\langle \bullet, \bullet \rangle}$.
- $D \ni x = (x^1, \dots, x^q)$; data space
- $t \in [0, 1]$; generation time
- $c \in \Omega \subset \mathbb{R}^p$; conditions in a bounded domain Ω .
- $\xi = (\xi^0, \xi^1, \dots, \xi^p) := (t, c) \in \tilde{\Omega} := [0, 1] \times \Omega$.
- $x \in D \subset \mathbb{R}^q$; data in a compact subset D
- For a matrix-valued function $u: \Xi \times D \rightarrow \mathbb{R}^{d \times \dim \Xi}$, let $u_{i,j}$ denote its (i, j) -th coordinate, where $i \in [d]$, $j \in [\dim \Xi]$. We then define

$$\operatorname{div}_x u: \Xi \times D \rightarrow \mathbb{R}^{\dim \Xi} \quad \text{as} \quad \operatorname{div}_x u(\xi, x) := \left(\sum_{i=1}^d \partial_i u_{i,0}(\xi, x), \dots, \sum_{i=1}^d \partial_i u_{i,\dim \Xi}(\xi, x) \right)^\top. \quad (\text{A.1})$$

- For $\varphi \in C^1(\tilde{\Omega} \times D; \mathbb{R}^{p+1})$,

$$\nabla_x \varphi := \begin{pmatrix} \partial_{x^1} \varphi^0 & \dots & \partial_{x^1} \varphi^p \\ \vdots & \ddots & \vdots \\ \partial_{x^q} \varphi^0 & \dots & \partial_{x^q} \varphi^p \end{pmatrix} \in \mathbb{R}^{q \times (p+1)}.$$

- $\mathcal{P}(X)$; the space of Borel probability measures on a space X , endowed with the narrow topology
- $\mathcal{P}_2(X)$; the L^2 -Wasserstein space
- $\delta_x \in \mathcal{P}_2(X)$; the delta measure supported at $x \in X$
- $\mu_\bullet: \tilde{\Omega} \ni \xi \mapsto \mu_\xi \in \mathcal{P}(D)$ conditional probability distribution
- $L^2(\Omega; X)$; the Lebesgue space valued in a metric space X , see (Lavenant, 2019, Definition 3.1)
- $H^1(\Omega; X)$; the Sobolev space valued in a metric space X , see (Lavenant, 2019, Definition 3.18). In particular, we set $\Gamma := H^1(\tilde{\Omega}; D)$
- $\operatorname{Dir}(\mu)$ is the Dirichlet energy of $\mu \in L^2(\Omega; \mathcal{P}(D))$, see (Lavenant, 2019, Definition 3.5).
- $\operatorname{Unif}(S)$ is the uniform distribution on a subset S of a Euclidean space with unit mass.
- $Q \in \mathcal{P}(\Psi)$. We will denote by ψ the sample from a probability distribution Q .
- $\sigma(X)$ denotes the σ -algebra of a random variable

Following the notation in (Durrett, 2019), we also use the notation $x \sim p$ to designate that x is sampled from the distribution p .

A.2 GENERALIZED CONTINUITY EQUATION

According to (Lavenant, 2019, Definition 3.4), we introduce a distributional solution of a generalized continuity equation formally given as

$$\nabla_\xi \mu(\xi, x) + \operatorname{div}_x(\mu(\xi, x)v(\xi, x)) = 0. \quad (\text{A.2})$$

The rigorous sense of (A.2) is stated in the following.

Definition A.1 (A distributional solution of the generalized continuity equation). A pair (μ, v) of a Borel mapping $\mu: \tilde{\Omega} \rightarrow \mathcal{P}(D)$ valued in probability measures and a Borel matrix field $v: \tilde{\Omega} \times D \rightarrow \mathbb{R}^{q \times (p+1)}$ is a *solution of the continuity equation* if it holds that

$$\int_{\tilde{\Omega}} \int_{\mathbb{R}^q} |v(\xi, x)|^2 d\mu_\xi(x) d\xi < +\infty,$$

and

$$\int_{\tilde{\Omega}} \int_{\mathbb{R}^q} (\operatorname{div}_\xi \varphi(\xi, x) + \langle \nabla_x \varphi(\xi, x), v(\xi, x) \rangle) d\mu_\xi(x) d\xi = 0,$$

for all $\varphi \in C_c^\infty(\tilde{\Omega} \times \mathbb{R}^q; \mathbb{R}^{p+1})$.

If a solution (μ, v) of the continuity equation is smooth, a path γ on $\tilde{\Omega}$ induces a path on $\mathcal{P}(D)$:

Proposition A.2 (Lifting conditional paths to probability paths). *Let (μ, v) be a solution of the continuity equation and $\gamma: [0, 1] \ni s \mapsto \gamma(s) \in \tilde{\Omega}$ be a continuously differentiable curve in $\tilde{\Omega}$. Set $\mu^\gamma := \mu_{\gamma(\bullet)}: [0, 1] \rightarrow \mathcal{P}(D)$ and $v^\gamma(s, x) := v(\gamma(s), x)\dot{\gamma}(s) \in \mathbb{R}^q$ for $(s, x) \in [0, 1] \times \mathbb{R}^q$.*

Suppose that $\operatorname{Dir}(\mu) < +\infty$ and there exists a probability density $\rho \in C^\infty(\tilde{\Omega}; L^\infty(D))$ of μ with respect to the Lebesgue measure.

Then, (μ^γ, v^γ) satisfies the continuity equation in the sense of distributions, i.e.,

$$\int_0^1 \int_{\mathbb{R}^q} (\partial_s \zeta(s, x) + \langle \nabla_x \zeta(s, x), v^\gamma(s, x) \rangle) d\mu_s^\gamma(x) ds = 0,$$

for all $\zeta \in C_c^\infty([0, 1] \times \mathbb{R}^q)$.

Proof. By (Lavenant, 2019, Proposition 3.16), there exists a unique $\varphi(\xi, \bullet) \in H^1(D; \mathbb{R}^{p+1})$ for every $\xi \in \tilde{\Omega}$ satisfying

$$\nabla_\xi \rho(\xi, x) + \operatorname{div}_x(\rho(\xi, x)\nabla_x \varphi(\xi, x)) = 0, \quad x \in \overset{\circ}{D},$$

and $v = \nabla_x \varphi$ on $\operatorname{supp} \mu$, where $\overset{\circ}{X}$ is the interior of a subset X . Thus, we have

$$\begin{aligned} \partial_s \rho(\gamma(s)) + \operatorname{div}_x(\rho(\gamma(s), x)v^\gamma(s, x)) &= (\nabla_\xi \rho(\gamma(s), x) + \operatorname{div}_x(\rho(\gamma(s), x)v(\gamma(s), x)))\dot{\gamma}(s) \\ &= (\nabla_\xi \rho(\gamma(s), x) + \operatorname{div}_x(\rho(\gamma(s), x)\nabla_x \varphi(\gamma(s), x)))\dot{\gamma}(s) \\ &= 0. \end{aligned}$$

■

Remark A.3. The smoothness assumption of Proposition A.2 recommends us to use some smooth probability measures as source distributions $\mu_{t=1,c}$, $c \in \Omega$.

According to Proposition A.2 and the well-known fact (see (Ambrosio et al., 2008, Proposition 8.1.8)), if we want a sample under a certain condition $c \in \Omega$, we can flow samples from a source distribution according to the family $(v^\gamma(s, \bullet))_{s \in [0, 1]}$ of vector fields determined from a path γ satisfying $\gamma(1) = (1, c)$.

A.3 PRINCIPLED MASS ALIGNMENT

A straightforward generalization of (Kerrigan et al., 2024a, Theorem 1 and Theorem 3) yields the following principle in flow marching theory.

Lemma A.4 (Principled mass alignment lemma). *Let \mathcal{F} be a separable (complete) metric space and P be a Borel probability measure on \mathcal{F} . Let (μ^f, v^f) be a solution of the continuity equation, in the sense of Definition A.1, for each $f \in \mathcal{F}$. Set the marginal distribution as*

$$\bar{\mu} := \int_{\mathcal{F}} \mu^f \, dP(f).$$

Assume that

$$\int_{\mathcal{F}} \int_{\tilde{\Omega}} \int_{\mathbb{R}^q} |v^f(\xi, x)|^2 \, d\mu_{\xi}^f(x) \, d\xi \, dP(f) < +\infty,$$

and μ_{ξ}^f is absolutely continuous with respect to $\bar{\mu}_{\xi}$ for P -a.e. f and a.e. $\xi \in \tilde{\Omega}$. Then, $(\bar{\mu}, \bar{v})$ is also a solution, where

$$\bar{v}(\xi, x) = \int_{\mathcal{F}} v^f(\xi, x) \frac{d\mu_{\xi}^f(x)}{d\bar{\mu}_{\xi}(x)} \, dP(f),$$

for $(\xi, x) \in \tilde{\Omega} \times D$. Moreover, for another matrix field u satisfying

$$\int_{\tilde{\Omega}} \int_{\mathbb{R}^q} |u(\xi, x)|^2 \, d\bar{\mu}_{\xi}(x) \, d\xi < +\infty,$$

we have

$$\int_{\tilde{\Omega}} \int_{\mathbb{R}^q} \langle \bar{v}(\xi, x), u(\xi, x) \rangle \, d\bar{\mu}_{\xi}(x) \, d\xi = \int_{\mathcal{F}} \int_{\tilde{\Omega}} \int_{\mathbb{R}^q} \langle v^f(\xi, x), u(\xi, x) \rangle \, d\mu_{\xi}^f(x) \, d\xi \, dP(f). \quad (\text{A.3})$$

Lemma A.4 leads to Theorem 3.4 as follows: first, in Lemma A.4, identify (\bar{v}, u) with (u, u_{θ}) in Theorem 3.4. then we see from (A.3) that

- $\int_{\Xi} \mathbb{E}_{x \sim \mu_{\xi}} [\langle u(\xi, x), u_{\theta}(\xi, x) \rangle] \, d\xi$ and
- $\int_{\Xi} \mathbb{E}_{\psi \sim Q, x \sim \mu_{\xi}^{\psi}} [\langle v^{\psi}(\xi, x), u_{\theta}(\xi, x) \rangle] \, d\xi$ are equal,

where v^{ψ} is a matrix field such that $v^{\psi}(\xi, \psi(\xi)) = \nabla_{\xi} \psi(\xi)$ with $\xi \in \Xi$. Also, because $\mu_{\xi}^{\psi} = \delta_{\psi(\xi)}$ is a delta distribution concentrated on $\psi(\xi)$, these are both equal to $\int_{\Xi} \mathbb{E}_{\psi \sim Q} [\langle \nabla_{\xi} \psi(\xi), u_{\theta}(\psi(\xi)) \rangle] \, d\xi$, as well. If we use this identity to the expansion of the square norm in (3.6), then the Theorem 3.4 follows from the same logic as (Kerrigan et al., 2024a, Theorem 3).

A.4 LIFTING DATA-VALUED FUNCTION TO PROBABILITY-MEASURE-VALUED FUNCTION

In order to construct a solution of the generalized continuity equation, we start to consider a particle-based solution of the continuity equation.

According to (Brenier, 2003, Subsection 3.1) and (Lavenant, 2019, Section 5), we can easily construct a solution of the continuity equation from a given function $\psi \in H^1(\tilde{\Omega}; D)$.

Lemma A.5. *Let $\psi \in H^1(\tilde{\Omega}; D)$ be a function satisfying*

$$\int_{\tilde{\Omega}} |\nabla_{\xi} \psi(\xi)|^2 \, d\xi < +\infty.$$

Set $\mu_{\bullet}^{\psi} := \delta_{\psi(\bullet)} \in H^1(\tilde{\Omega}; \mathcal{P}(D))$. Assume that there exists a matrix field satisfying

$$v^{\psi}(\xi, \psi(\xi)) = \nabla_{\xi} \psi(\xi), \quad (\text{A.4})$$

for $\xi \in \tilde{\Omega}$. Then, (μ^{ψ}, v^{ψ}) is a solution of the continuity equation.

Combining [Lemmas A.4](#) and [A.5](#), we can construct another solution of the continuity equation.

Corollary A.6 (The paths make the solution.). *Let $Q \in \mathcal{P}(H^1(\tilde{\Omega}; D))$ be a Borel probability measure, and (μ^ψ, v^ψ) be a solution defined in [Lemma A.5](#) Q -a.e. $\psi \in H^1(\tilde{\Omega}; D)$ and*

$$\mu^Q := \int_{H^1(\tilde{\Omega}; D)} \mu^\psi dQ(\psi)$$

is their marginal distribution. Assume that

$$\int_{H^1(\tilde{\Omega}; D)} \int_{\tilde{\Omega}} \int_{\mathbb{R}^q} |v^\psi(\xi, x)|^2 d\mu_\xi^\psi(x) d\xi dQ(\psi) < +\infty,$$

and $\mu^\psi \ll \mu^Q$. Then, (μ^Q, v^Q) is also a solution of the continuity equation, where

$$v^Q = \int_{H^1(\tilde{\Omega}; D)} v^\psi(\xi, x) \frac{d\mu_\xi^\psi(x)}{d\mu_\xi}(x) dQ(\psi).$$

B TECHNICAL PROOFS

The following claim follows immediately from the convexity of the Dirichlet energy as shown in [Lavenant \(2019, Proposition 3.13\)](#) and from Jensen's inequality:

Proposition B.1 (Straightness is controlled by ψ). *Let $\mu_{t,c} = \mathbb{E}_{\psi \sim Q} [\delta_{\psi(t,c)}]$ $((t, c) \in I \times \Omega)$ with $\eta \in \mathcal{P}(D)$. Then, the Dirichlet energy of $\mu: I \times \Omega \rightarrow \mathcal{P}(D)$ is bounded as*

$$\text{Dir}_{I \times \Omega}(\mu) \leq \iint_{I \times \Omega} \mathbb{E}_{\psi \sim Q} \|\nabla_{t,c} \psi(t, c)\|^2 dt dc.$$

Proposition B.2. *Let $\mu \in H^1(\tilde{\Omega}; \mathcal{P}(D))$ be a [smooth](#) solution of the continuity equation, and $v: \tilde{\Omega} \times \mathbb{R}^q \rightarrow \mathbb{R}^{q \times (p+1)}$ is the matrix field associated with μ . Assume that $v \in C^1(\tilde{\Omega} \times \mathbb{R}^q; \mathbb{R}^{q \times (p+1)})$ and the derivatives $\partial_c v$, $\partial_x v$ of v is bounded on $\tilde{\Omega} \times \mathbb{R}^q$. Then, there exists a constant $C > 0$ depend on p, q such that*

$$\text{Dir}(\mu(1, \bullet)) \leq C \exp\left(\|\partial_x v\|_{L^\infty(\tilde{\Omega} \times \mathbb{R}^q; \mathcal{B}(\mathbb{R}^q \times \tilde{\Omega}; \mathbb{R}^q))}\right) (\text{Dir}(\mu(0, \bullet)) + \|\partial_c v\|_\infty).$$

Here, $\|f\|_\infty = \sup_{(\xi, x) \in \tilde{\Omega} \times \mathbb{R}^q} |f(\xi, x)|$ for a finite-dimensional valued continuous function f on $\tilde{\Omega} \times \mathbb{R}^q$.

The proof of [Proposition B.2](#) is similar to ([Isobe, 2023, Proposition 5.4](#)).

Proof. By virtue of ([Lavenant, 2019, Proposition 3.21](#)), we have to estimate

$$\text{Dir}(\mu(1, \bullet)) = \lim_{\varepsilon \rightarrow 0} \frac{C_p}{\varepsilon^{p+2}} \iint_{\Omega^2} W_2^2(\mu(1, c^1), \mu(1, c^2)) dc^1 dc^2.$$

The integrand of the above is decomposed as

$$\begin{aligned} W_2(\mu(1, c^1), \mu(1, c^2)) &= W_2\left(\Phi_{\#}^{1, c^1} \mu(0, c^1), \Phi_{\#}^{1, c^2} \mu(0, c^2)\right) \\ &\leq W_2\left(\Phi_{\#}^{1, c^1} \mu(0, c^1), \Phi_{\#}^{1, c^2} \mu(0, c^1)\right) + W_2\left(\Phi_{\#}^{1, c^2} \mu(0, c^1), \Phi_{\#}^{1, c^2} \mu(0, c^2)\right). \end{aligned} \tag{B.1}$$

Here $\Phi^{t,c}: \mathbb{R}^q \rightarrow \mathbb{R}^q$ is a flow mapping satisfying

$$\Phi^{t,c}(x) = x + \int_0^t v(s, c, \Phi^{s,c}(x)) \begin{pmatrix} 1 \\ 0 \end{pmatrix} ds.$$

The first term of (B.1) is bounded as

$$W_2\left(\Phi_{\#}^{1,c^1}\mu(0,c^1), \Phi_{\#}^{1,c^2}\mu(0,c^1)\right)^2 \leq \int_{\mathbb{R}^q} \left|\Phi^{t,c^1}(x) - \Phi^{t,c^2}(x)\right|^2 d\mu_{0,c^1}(x).$$

Then, the integrand is also bounded by

$$\begin{aligned} \left|\Phi^{t,c^1}(x) - \Phi^{t,c^2}(x)\right| &\leq \int_0^t \left\|v(s, c^1, \Phi^{s,c^1}(x)) - v(s, c^2, \Phi^{s,c^2}(x))\right\|_{\text{op}} ds \\ &\leq |c^1 - c^2| \|\partial_c v\|_{\infty} \\ &\quad + \int_0^t \|\partial_x v\|_{\infty} \left|\Phi^{t,c^1}(x) - \Phi^{t,c^2}(x)\right| ds. \end{aligned}$$

Thus, the Gronwall inequality yields

$$\left|\Phi^{t,c^1}(x) - \Phi^{t,c^2}(x)\right| \leq |c^1 - c^2| \|\partial_c v\|_{L^{\infty}(\tilde{\Omega} \times \mathbb{R}^q; \mathcal{B}(\Omega \times \tilde{\Omega}; \mathbb{R}^q))} \exp\left(\|\partial_x v\|_{L^{\infty}(\tilde{\Omega} \times \mathbb{R}^q; \mathcal{B}(\mathbb{R}^q \times \tilde{\Omega}; \mathbb{R}^q))}\right). \quad (\text{B.2})$$

By a similar argument, the second term of (B.1) is also bounded as

$$W_2\left(\Phi_{\#}^{1,c^2}\mu(0,c^1), \Phi_{\#}^{1,c^2}\mu(0,c^2)\right) \leq W_2(\mu(0,c^1), \mu(0,c^2)) \exp\left(\|\partial_x v\|_{L^{\infty}(\tilde{\Omega} \times \mathbb{R}^q; \mathcal{B}(\mathbb{R}^q \times \tilde{\Omega}; \mathbb{R}^q))}\right). \quad (\text{B.3})$$

Combining (B.2) and (B.3) completes the proof. \blacksquare

C PSEUDO-CODES

Algorithm 4 Algorithm of OT-CFM

Input: Neural Network $v_{\theta}: I \times D \rightarrow \mathbb{R}^d$, the source distribution μ_0 , the dataset $D_* \subset D$ from a target distribution μ .

Return: $\theta \in \mathbb{R}^p$

1: **for** each iteration **do**

 # Step 1: Sample from datasets

2: Sample a batch B^0 from μ_0

3: Sample a batch B^1 from D_*

 # Step 2: Construct $\psi: I \rightarrow D$

4: Construct an optimal transport plan π between B^0 and B^1

5: Jointly sample $(x_0, x_1) \sim \pi$

6: Sample $t \sim \text{Unif}(I)$

7: Compute

$$\begin{aligned} \psi_t &:= \psi(t \mid x_0, x_1) \\ &= (1-t)x_0 + tx_1 \\ \dot{\psi}_t &:= \dot{\psi}(t \mid x_0, x_1) \\ &= x_1 - x_0 \end{aligned}$$

8: Update θ by the gradient of $\|v_{\theta}(t, \psi_t) - \dot{\psi}_t\|^2$

9: **end for**

D SAMPLING OF $\bar{\psi}$ IN (4.1) IN SECTION 4 FOR MMOT-EFM

In this section, we follow the notation in Section 4 and describe in more detail the construction of $\bar{\psi}(c \mid \mathbf{x}_{C_0})$ in (4.1), which is

$$\psi(t, c \mid x_{0,c}, \mathbf{x}_{C_0}) = (1-t)x_{0,c} + t\bar{\psi}(c \mid \mathbf{x}_{C_0})$$

Algorithm 5 Flow Matching (Training)

Input: Neural Network $v_\theta: I \times D \rightarrow \mathbb{R}^d$, the source distribution μ_0 , the dataset $D_* \subset D$ from a target distribution μ .

Return: $\theta \in \mathbb{R}^p$

```

1: for each iteration do
  # Step 1: Sampling from datasets
2: Sample batches  $B^0 = \{x_0^i\}_{i=1}^N$  from source  $p_0$ 
3: Sample batches  $B^1 = \{x_1^j\}_{j=1}^N$  from dataset  $D_*$ 
  # Step 2: Constructing a supervisory path  $\psi$ 
4: Construct an optimal transport plan  $\pi \in \mathbb{R}^{N \times N}$  between  $B^0$  and  $B^1$ 
5: Jointly sample  $(x_0, x_1) \in B^0 \times B^1$  from  $\pi$ 
6: Sample  $t \in I$ 
7: Compute
  (A)  $\psi_t := \psi(t \mid x_0, x_1) = (1-t)x_0 + tx_1$ 
  (B)  $\nabla \psi_t := \nabla_t \psi(t \mid x_0, x_1) = x_1 - x_0$ 
  # Step 3: Learning vector fields
8: Update  $\theta$  by the gradient of  $\|v_\theta(t, \psi_t) - \nabla \psi_t\|^2$ 
9: end for

```

Algorithm 6 ODEsolve for generation

Input: Initial data $x_0 \in D$, vector fields $v: I \times D \rightarrow \mathbb{R}^d$

Return: Terminal value $\phi_1^v(x_0)$ of the solution of ODE $\dot{\phi}_t^v(x_0) = v(t, \phi_t^v(x_0))$

1: Compute $\phi_1(x_0)$ via a discretization of the ODE in t

Algorithm 7 Extended Flow Matching (Training)

Input: Condition set $C \subset \Omega \subset \mathbb{R}^k$, set of datasets $D_c \subset D \subset \mathbb{R}^d$ for each $c \in C$, network $u_\theta: I \times \Omega \times D \rightarrow \mathbb{R}^{d \times (1+k)}$, source distributions $p_0(\cdot \mid c)$ ($c \in C$)

Return: $\theta \in \mathbb{R}^p$

```

1: for each iteration do
  # Step 1: Sampling from datasets
2: Sample  $C_0 = \{c_i\}_{i=1}^{N_c} \subset C$ 
3: Sample a batch  $B_{0,c}$  from  $p_0(x \mid c)$  for each  $c \in C_0$ 
4: Sample a batch  $B_{1,c}$  from  $D_c$  for each  $c \in C_0$ 
5: Put  $B^0 := \{B_{0,c}\}_{c \in C_0}$  and  $B^1 := \{B_{1,c}\}_{c \in C_0}$ 
  # Step 2: Constructing supervisory paths  $\{\psi_j\}_{j=1}^N$ 
6: Construct a transport plan  $\pi$  among  $B^0$  and  $B^1$  # see Section 4
7: Sample  $\{(x_{t,c}^j)_{(t,c) \in \{0,1\} \times C_0}\}_{j=1}^N \subset D^{2N_c}$  from  $\pi$ 
8: For all  $j \in [1 : N]$ , define  $\psi_j: I \times \Omega \rightarrow D$  that regresses  $(x_{t,c}^j)_{(t,c) \in \{0,1\} \times C_0}$  on  $\{0,1\} \times C_0$  # see Equation (4.1)
9: Sample  $\{t_k\}_{k=1}^{N_t} \subset I$ 
10: Sample  $\{c'_l\}_{l=1}^{N'_c} \subset \text{ConvHull}(C_0)$ 
11: For all  $j \in [1 : N]$ ,  $k \in [1 : N_t]$ ,  $l \in [1 : N'_c]$ , compute
  (A)  $\psi_{j,k,l} := \psi_j(t_k, c'_l)$ 
  (B)  $\nabla \psi_{j,k,l} := \nabla_{t,c} \psi_j(t_k, c'_l)$ 
  # Step 3: Learning matrix fields
12: Compute the loss

```

$$L(\theta) = \frac{1}{NN_tN'_c} \sum_{j,k,l} \|u_\theta(t_k, c'_l, \psi_{j,k,l}) - \nabla \psi_{j,k,l}\|^2$$

13: Update θ by the gradient of $L(\theta)$

14: **end for**

and the corresponding joint distribution of $x_{C_0} := \{x_i\}_{c_i \in C_0}$ on $D^{2|C_0|}$ we used in step 2 of the training algorithm. In the final part of this section, we also elaborate how we couple $x_{0,c}$ with x_{C_0} .

As we describe in the main manuscript, we introduce our EFM as a direct extension of FM as a method to transform one distribution to another through a learned vector field. In particular, we present in this paper an implementation of EFM which extends OT-CFM [Tong et al. \(2023b\)](#), which aims to train FM as an approximate optimal transport between two distributions (source μ_0 and target μ_1). To formalize this extension, we need to describe OT as a minimization of Dirichlet Energy.

D.1 OT-CFM AS APPROXIMATE DIRICHLET ENERGY MINIMIZATION

As is principally described in [Lavenant \(2019\)](#), OT emerges as a coupling of the source μ_0 and the target μ_1 constructed from the constant-speed geodesic (with respect to Wasserstein distance) between μ_0 and μ_1 , which can be realized by minimizing the Dirichlet energy

$$\text{Dir}(\mu) = \inf_{v: I \times D \rightarrow \mathbb{R}^d} \left\{ \int_{[0,1] \times D} \frac{1}{2} \|v(t, x)\|^2 \mu_t(dx) dt \mid \partial_t \mu_t(x) + \text{div}_x(\mu_t(x)v(t, x)) = 0 \right\} \quad (\text{D.1})$$

over all set of $\mu: [0, 1] \rightarrow \mathcal{P}(D)$ satisfying $\mu(0) = \mu_0$, $\mu(1) = \mu_1$. It is well known that in the standard Euclidean metric space, the minimal energy is achieved by μ corresponding to $v(t, x)$ that is the derivative of a straight-line of form $\psi^T(t|x) = tT(x) + (1-t)x$ where $T: D \rightarrow D$, and more particularly as the minimum of

$$\int_{D \times D} \frac{1}{2} \|x - y\|^2 \pi(dx, dy) = \int_D \frac{1}{2} \|\partial_t \psi^T(t|x)\|^2 (I \times T)_{\#} \mu_0(dx) \quad (\text{D.2})$$

over all $\pi \in \mathcal{P}(D \times D)$ with marginal distribution μ_0 and μ_1 or equivalently over all T with $T_{\#} \mu_0 = \mu_1$. In OT-CFM, this π (or T) is approximated by the discrete optimal transport solution over a pair of batches B_0, B_1 sampled respectively from source and target distributions. Note that, in this view, $(I \times T)_{\#} \mu_0$ induces a distribution Q on the path $[0, 1] \rightarrow D$ generating $\psi^T(t|x)$ with randomness derived from x .

Theorem 3.1 of [Yim et al. \(2024\)](#) guarantees that the (batch)sample-averaged version of μ and the (batch)sample-averaged version of v satisfies the continuity equation, thereby yielding the approximation of the dirichlet energy minimizing flow map.

D.2 MMOT-EFM AS APPROXIMATE DIRICHLET ENERGY MINIMIZATION

To mimic this construction in multi-marginal setting of EFM, we aim to approximate the solution to the minimization of

$$\text{Dir}(\mu) = \inf_{v: \Omega \times D \rightarrow \mathbb{R}^{d \times k}} \left\{ \int_{\Omega \times D} \frac{1}{2} \|v(c, x)\|^2 \mu_{\xi}(dx) dc \mid \partial_c \mu_{\xi}(x) + \text{div}_x(\mu(c, x)v(c, x)) = 0 \right\} \quad (\text{D.3})$$

over all set of $\mu: \Omega \rightarrow \mathcal{P}(D)$ satisfying $\mu(c_i) = \mu_i$ for all $c_i \in C_0$. Note that when $\Omega = [0, 1]$, this minimization problem (i.e. Dirichlet Problem) agrees with that of the OT problem on which the method of FM is established.

Now, in a similar philosophy as FM, we would aim to approximate this Dirichlet energy through multi-marginal optimal transport [Piran et al. \(2024\)](#) over discrete samples. Now, under *sufficient* regularity condition (Prop 5.6 [Lavenant \(2019\)](#)), we can similarly argue that there exists some probability Q on the space $\mathcal{F} = H^1(\Omega, D)$ of a map from “condition” to “data” satisfying

$$\text{Dir}(\mu) = \int_{\Omega \times \mathcal{F}} \|\partial_c \psi(c)\|^2 Q(d\psi) dc \quad (\text{D.4})$$

and our goal winds down to finding the energy-minimizing distribution Q . In this endeavor, we implicitly find Q by specifying a particular space of functions \mathcal{F} and generating $\psi: \Omega \rightarrow D$ from

a set of $\{(c_i, x_i)\}_{c_i \in C_0}$ of "condition value" and "observation" for jointly sampled $\{x_i\}_i$ as the regression

$$\bar{\psi}(\cdot|\{x_i\}_i) = \arg \min_{\psi \in \mathcal{F}} \sum_{c_i \in C_0} \|\psi(c_i) - x_i\|^2 \quad (\text{D.5})$$

and minimize the energy with respect to the joint distribution π on $D^{|C|}$ from which to sample $\{x_i\}_i$. That is, we aim to minimize

$$\int \|\nabla_c \bar{\psi}(c|\{x_i\}_i)\|^2 \pi(\{dx_i\}_i) dc \quad (\text{D.6})$$

with respect to π . This, indeed, is in the format of MMOT problem, where $c(\{x_i\}_i) := \|\partial_c \psi(c|\{x_i\}_i)\|^2$. \mathcal{F} can be chosen for example, as an RKHS or a space of linear function, so that the regression can be solved analytically with respect to c .

Just as is done in OT-CFM, we approximate this π with the joint distribution over finite tuple of batches $\{B_i\}_i$ with each B_i sampled from μ_i corresponding to condition c_i . This approximation is indeed the very π that we adopt in MMOT version of our EFM in step 2.

Now, by the virtue of Theorem of principle-mass-alignment A.6, we can argue that the (batch)sample-averaged distributions μ^ψ and the (batch)sample-averaged $v^\psi = \partial_c \psi$ solve the *generalized* continuity equation, thereby yielding the approximation of the Dirichlet energy minimizing map $\mu : \Omega \rightarrow \mathcal{P}(D)$.

Note that the above constructions of $\psi \sim Q$ is in complete parallel with that of OT-CFM. See Table 3 for the correspondences. We also note that this argument can be extended to $\tilde{\Omega} = [0, 1] \times \Omega$ in place

Table 3: OT-CFM vs MMOT-EFM

Framework	OT-CFM	MMOT-EFM
μ	$[0, 1] \rightarrow \mathcal{P}(D)$	$\Omega \rightarrow \mathcal{P}(D)$
ψ	$[0, 1] \rightarrow D$	$\Omega \rightarrow D$
v	$\partial_t \psi$	$\partial_c \bar{\psi}$
(μ, v) relation	Continuity	Generalized Continuity
Boundaries	$\{\mu_0, \mu_1\}$	$\{\mu_i\}_{c_i \in C_0}$
Approximation	OT	MMOT

of Ω . However, because of the computational cost of MMOT, we construct our generative model from (4.1), which combines ψ and the OT-CFM construction. In the next section, we elaborate on the construction of the approximation of π in (D.6) from which to sample $\bar{\psi}$ in (4.1)

D.3 APPROXIMATING MMOT

In general, MMOT is computationally heavy, and even with the advanced methods like the multi-marginal Sinkhorn method developed in (Lin et al., 2022), the computational cost scales as $|B|^{|C|}$, where $|B|$ is the batch size and $|C|$ is the number of conditions to be simultaneously considered. To reduce this cost, we took the approach of approximating MMOT through clustering. More particularly, when a batch from B_i is sampled each from μ_i for condition c_i , we applied K -means nearest neighborhood clustering (KNN) to B_i , yielding sub-batches $\{U_{ik}\}_{c_i \in C_0, k \in [1:K]}$ with mean values $\{m_{ik}\}_{c_i \in C_0, k \in [1:K]}$, where $\cup_{k \in [1:K]} U_{ik} = B_i$. Let $M_i = \{m_{ik}\}_{k \in [1:K]}$ be the set of cluster-means for batch i . Instead of conducting MMOT directly on batch B_i , we conduct the MMOT on $\{M_i\}_i$, whose cost will be on the order of $K^{|C|}$. Applying argmax operations on the result of MMOT from methods like the Sinkhorn method, we can obtain the deterministic coupling $\pi_m = (\times_i T_i)_{\#} \text{Unif}(M_0)$ where $\text{Unif}(M_0)$ is the uniform distribution on M_0 . After sampling $m_{0k^*} \sim \text{Unif}(M_0)$, we couple $U_{iT_i(k^*)}$ with a method of user's choice, where $T_i(k^*)$ is an *abuse of notation* satisfying

$$m_{iT_i(k^*)} = T_i(m_{0k^*}).$$

In our implementation of MMOT-EFM, we coupled $\{U_{iT_i(k^*)}\}_i$ with generalized-geodesic coupling as is used in Fan & Alvarez-Melis (2023), with center distribution being the standard Gaussian with

mean being the average of $\{U_{iT_i(k^*)}\}_i$. Although we provide a brief description of generalized-geodesic in reference F, we would like to refer to Ambrosio et al. (2008) for a more thorough study. Below, we summarize the sampling procedure of $\{x_i\}_{c_i \in C_0}$ in $\psi(\cdot | \{x_i\}_{c_i \in C_0})$ of MMOT-EFM.

Algorithm 8 MMOT sampling with Cluster

Input: Set of batches $\{B_i\}_i$ with each B_i sampled from $p(\cdot | c_i)$
Return: Joint sample $\{x_i\}_i$ from $\{B_i\}_i$

Step 1: Cluster MMOT setup

- 1: Cluster each B_i as $\cup_{k \in [1:K]} U_{ik} = B_i$ with $\text{mean}(U_{ik}) = m_{ik}$
- 2: Set $M_i = \{m_{ik}\}_{k \in [1:K]}$
- 3: Use MMOT to produce coupling on $\{M_i\}_i$ via $\{T_i\}_i \# \text{Unif}(M_0)$

Step 2: Sampling

- 4: Sample m_{0k^*} from $\text{Unif}(M_0)$
- 5: Compute $m_{iT_i(k^*)} := T_i(m_{0k^*})$
- 6: Jointly sample from $\{U_{iT_i(k^*)}\}$ with the method of user's choice, preferably with deterministic coupling, such as another round of MMOT or generalized-geodesic.

D.4 COUPLING OF $\{x_{0,c_i}\}_{c_i \in C_0}$ AND $\{x_i\}_{c_i \in C_0}$

Ideally, it is more closely aligned with the theory of Dirichlet energy to include the source distributions $\{\mu(0, c_i)\}_i$ into the set of distributions to be coupled in the MMOT, and enact the argument in Appendix D.2 with $\tilde{\Omega} = [0, 1] \times \Omega$ in place of Ω . As mentioned in the previous section, however, the cost of empirical MMOT scales exponentially with the number of distributions to couples. We, therefore, took an alternative coupling strategy as a computational compromise.

First, recall from the step 1 of Section 4 that $\{x_{0,c_i}\}_{c_i \in C_0}$ are already coupled with common standard Gaussian sample in the form of $\mu_{0,c} = \text{Mean}[D_c] + \mathcal{N}(0, I)$. To couple $\{x_{0,c_i}\}_{c_i \in C_0}$ with $\{x_i\}_{c_i \in C_0}$ which are deterministically coupled through the routine of Section D.3 as $\{x_i\}_{c_i \in C_0} = \{\mathcal{T}_i(x_0)\}_{c_i \in C_0}$ with x_0 sampled from $p(\cdot | c_0)$, we may simply couple x_{0,c_0} with x_0 and this will automatically induce the deterministic coupling of $\{x_{0,c_i}\}_{c_i \in C_0}$ and $\{x_i\}_{c_i \in C_0}$. In particular, if B_{0,c_0} is a batch of samples from $p_0(\cdot | c_0)$ and B_{1,c_0} is a batch of samples from D_{c_0} in the step1 of the training, we may couple B_{0,c_0} with B_{1,c_0} with optimal transport with the methods of user's choice, such as those provided in Flamary et al. (2021).

D.5 APPROXIMATION OF MMOT BY THE ORDINAL OT

According to von Lindheim (2022),

E APPROXIMATION OF MMOT BY WASSERSTEIN BARYCENTER

F A REMARK ON GENERALIZED GEODESIC COUPLING(GGC) AND THE SAMPLING OF $\bar{\psi}$ IN (4.1) IN SECTION 4 FOR GGC-EFM

As we have mentioned in Section 3.1, EFM can be defined with any distribution $Q \in \mathcal{P}(\Psi)$ on the space of functions $\Psi := \{\psi: I \times \Omega \rightarrow D \mid \psi \text{ is differentiable}\}$ satisfying the boundary conditions (3.3). We also present still another construction of $\bar{\psi}$ derived from different coupling.

F.1 GENERALIZED GEODESIC COUPLING

Generalized geodesic of $\{\mu_i\}$ with base $\nu \in \mathcal{P}(D)$, also known in the name of linear optimal transport Moosmüller & Cloninger (2020) in mathematical literatures, was introduced by Ambrosio et al. (2008) as

$$\rho_a := \left(\sum_{i=1}^m a_i T_i \right) \# \nu, \quad a \in \Delta_{m-1} \quad (\text{F.1})$$

Algorithm 9 Reference algorithm

```

1: Input: Discrete measures  $\mu^i = \sum_{j=1}^{n_i} \mu_j^i \delta(x_j^i)$ ,  $i = 1, \dots, N$ , with  $x_1^1 < \dots < x_{n_1}^1$  if  $d = 1$ 
2: for  $i = 2, \dots, N$  do
3:   Compute
      
$$\pi^i \in \arg \min_{\pi \in \Pi(\mu^1, \mu^i)} \langle c, \pi \rangle = \sum_j \pi_j^i(x_{1,i,j}, x_{i,j}) \quad \text{s.t.} \quad \#\text{supp}(\pi) \leq n_1 + n_i - 1 \quad (3.9)$$

4: end for
5: Initialization:  $\tilde{\pi} = 0$ 
6: for  $k = 1, \dots, n_1$  do
7:   while  $x_k^1 \in P_1(\text{supp}(\pi^i))$  for  $i = 2, \dots, N$  do
8:     for  $i = 2, \dots, N$  do
9:        $j_i \leftarrow \min\{j : x_{1,i,j} = x_k^1\}$ 
10:    end for
11:     $h \leftarrow \min_j \pi_{j_i}^i$ 
12:     $\tilde{\pi} \leftarrow \tilde{\pi} + h\delta(x_k^1, x_{2,j_2}, \dots, x_{N,j_N})$ 
13:    for  $i = 2, \dots, N$  do
14:       $\pi^i \leftarrow \pi^i - h\delta(x_k^1, x_{i,j_i})$ 
15:    end for
16:  end while
17: end for
18: Output:  $\tilde{\pi}$ 

```

where T_i is the optimal map from ν to μ_i and Δ_{m-1} is the set of all $\{a_i\}_{i=1}^m$ with $\sum_i a_i = 1$. This is indeed one of the generalizations to the McCann’s interpolation used in OT between μ_0 and μ_1 through the expression

$$\rho_t := ((1-t)\text{Id} + tT)\#\mu_0, \quad t \in [0, 1]$$

which runs along the geodesic in $\mathcal{P}(D)$ with respect to Wasserstein distance. Note that ρ_a in Generalized Geodesic provides not only provides deterministic coupling of $\{\mu_i\}$ through $\rho_{e_i} = T_{i\#}\nu = \mu_i$, it also interpolates unknown distributions for any $a \in \Delta_{m-1}$. We would refer to the deterministic coupling in the form of $T_{i\#}\nu = \mu_i$ as GGC-coupling.

F.2 GGC SAMPLING OF $\bar{\psi}$

In analogy to the sampling procedure of $\bar{\psi}(\cdot|\{x_i\}_i)$ in MMOT-EFM with MMOT-coupled $\{x_i\}_i$, we may sample $\bar{\psi}(\cdot|\{x_i\}_i)$ with $\{x_i\}_i$ that is jointly sampled with GGC-coupling. We emphasize that $\bar{\psi}$ constructed in such a way does not necessarily minimize an explicit objective as Dirichlet energy and this might result in EFM with a somewhat erratic style transfer. For more empirical investigations, please see the main manuscript.

G EXPERIMENT DETAILS FOR CONDITIONAL MOLECULAR GENERATION

G.1 METRICS

To evaluate our conditional generation, we use the pre-trained VAE model to encode EFM-generated latent vectors into molecular structures and compute the Mean Absolute Error(MAE) between the generated molecule’s property values and the conditioning property values. MAEs are calculated separately for interpolation and extrapolation. All MAEs are first calculated for each property and then averaged for both properties.

G.2 DATASET AND BASELINES

We first trained a Site-information-encoded Junction Tree Variational Autoencoder (SJT-VAE) model, a variant implementation of the Junction Tree Variational Autoencoder (JT-VAE) (Jin et al.,

2018). SJT-VAE was initially designed to eliminate the arbitrariness of JT-VAE and enable applications such as RJT-RL (Ishitani et al., 2022). We chose SJT-VAE over JT-VAE due to its superior reconstruction accuracy and faster training times. However, we expect that similar results could be reproduced with the original JT-VAE implementation.

Our SJT-VAE model was trained on the ZINC-250k dataset (Akhmetshin et al., 2021). A random subset of 80,000 molecules was labeled with the number of HBAs and the number of rotatable bonds, with all labels computed using RDKit. These 80,000 molecules were then binned into a 2D matrix based on their property values. From this matrix, we selected a region with concentrated data: molecules with 2 and 4 rotatable bonds and 3 and 5 HBAs, forming 4 bins with property sets (2, 3), (2, 5), (4, 3), and (4, 5). To balance the dataset, we up-sampled or capped the number of training examples to 5,000 per bin.

To evaluate out-of-distribution conditional generation, we generated molecules with property sets not included in the training set, specifically (3, 4), (2, 4), (4, 4), (3, 3), and (5, 5). For property sets where only one property is out-of-distribution, we calculated the MAE based solely on the out-of-distribution property.




Figure 5: Training set rotatable bonds and HBAs label distribution

All flow matching-based models, including MMOTEFM and baselines, are trained with a batch size of 250 and the learning rate of $1e^{-4}$ for 160,000 iterations. Training on a single Nvidia V-100 GPU with evaluation every 5000 iteration took around 4 hours.

COMPUTATIONAL RESOURCES

All models were trained on a single Nvidia V100-16G GPU, and 100 epochs were completed within 4 hours. Training for the MMOT-EFM model is performed on a single Nvidia V100-16G GPU within 2.5 hours. The results of MMOT-EFM for synthetic experiments were yielded from a model trained over 100000 iterations in 5 hours.

ADDITIONAL FIGURES




Figure 6: Conditional generation of the synthetic dataset by FM, organized in the grid for two axes of conditions.




Figure 7: Conditional generation of the synthetic dataset by MMOT-EFM, organized in the grid for two axes of conditions. The figures in the bottom row are the result of style transfer.




Figure 8: Conditional generation of synthetic dataset by Bayesian(COT)-FM with $\beta = 10^2$, organized in grid for two axis of conditions.