DEEP GENERATIVE WASSERSTEIN GRADIENT FLOWS

Anonymous authors
Paper under double-blind review

ABSTRACT

Deep generative modeling is a rapidly-advancing field with a wealth of modeling choices developed in the past decade. Amongst them, Wasserstein gradient flows (WGF) are a powerful and theoretically rich class of methods. However, their applications to high-dimensional distributions remain relatively underexplored. In this paper, we present Deep Generative Wasserstein Gradient Flows (DGGF), which constructs a WGF minimizing the entropy-regularized f-divergence between two distributions. We demonstrate how to train a deep density ratio estimator that is required for the WGF and apply it to the task of generative modeling. Experiments demonstrate that DGGF is able to synthesize high-fidelity images of resolutions up to 128×128, directly in data space. We demonstrate that DGGF has an interpretable diagnostic of sample quality by naturally estimating the KL divergence throughout the gradient flow. Finally, we show DGGF's modularity by composition with external density ratio estimators for conditional generation, as well as for unpaired image-to-image translation without modifications to the underlying framework.

1 INTRODUCTION

Gradient flow methods are a powerful and general class of techniques with diverse applications ranging from physics (Carrillo et al., 2019; Adams et al., 2011) and sampling (Bernton, 2018) to neural network optimization (Chizat & Bach, 2018) and reinforcement learning (Richemond & Maginnis, 2017; Zhang et al., 2018). In particular, Wasserstein gradient flow (WGF) methods are a popular specialization that model the gradient dynamics on the space of probability measures with respect to the Wasserstein metric; these methods aim to construct the *optimal* path between two probability measures — a source distribution $q(\mathbf{x})$ and a target distribution $p(\mathbf{x})$ — where the notion of optimality refers to the path of steepest descent in Wasserstein space.

The freedom in choosing $q(\mathbf{x})$ and $p(\mathbf{x})$ when constructing the WGF makes the framework a natural fit for a variety of generative modeling tasks. For data synthesis, we choose $q(\mathbf{x})$ to be a simple distribution easy to draw samples from (e.g., Gaussian), and $p(\mathbf{x})$ to be a complex distribution which we would like to learn (e.g., the distribution of natural images). The WGF then constructs the optimal path from the simple distribution to synthesize data resembling that from the complex distribution. Furthermore, we could choose both $p(\mathbf{x})$ and $q(\mathbf{x})$ to be distributions from different domains of the same modality (e.g., images from separate domains). The WGF then naturally performs domain translation.

However, despite this fit and the wealth of theoretical work established over the past decades (Ambrosio et al., 2005; Santambrogio, 2017), applications of WGFs to generative modeling of high-dimensional distributions remain under-explored and limited. A key difficulty is that the 2-Wasserstein distance and divergence functionals are generally intractable. Existing works rely on complex optimization schemes with constraints that contribute to model complexity, such as approximations of the 2-Wasserstein distance with input convex neural networks (Mokrov et al., 2021), dual variational optimization schemes with the Fenchel conjugate (Fan et al., 2021) or adopting a particle simulation approach, but amortizing sample generation to auxiliary generators (Gao et al., 2019; 2022).

In this work, we take a step towards resolving the shortcomings of WGF methods for deep generative modeling. We propose Deep Generative Wasserstein Gradient Flows (DGGF), which is formulated using the gradient flow of entropy-regularized f-divergences (Fig. 1). As this formulation involves

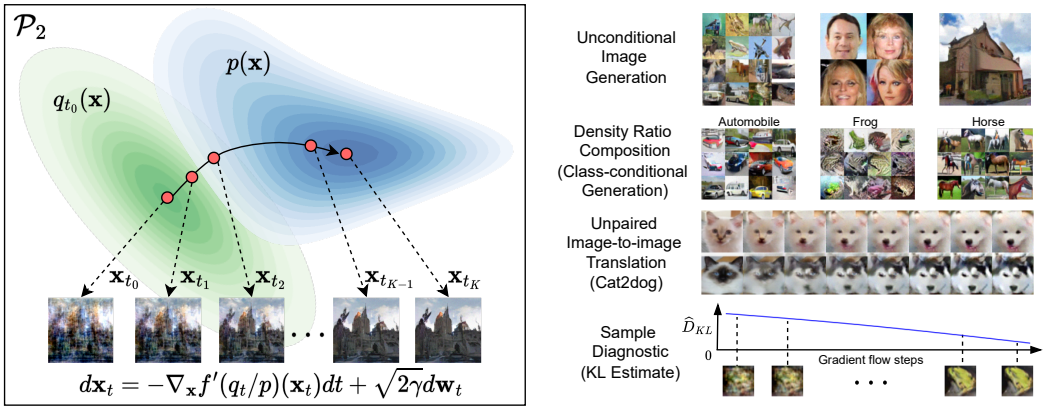

Figure 1: Left: illustration of the generative gradient flow process using DGGF. The evolution of the gradient flow is governed by the SDE shown in the figure. We visualize intermediate samples of the LSUN Church dataset. Right: visualization of application domains of DGGF. At its core, DGGF is able to perform high-fidelity unconditional image generation. The unconditional model can be used for class-conditional generation via density ratio composition with external pretrained classifiers. Additionally, DGGF is able to perform unpaired image-to-image translation with no modifications needed to the framework. Finally, DGGF possesses an innate sample diagnostic by estimating the KL divergence over the flow, which approaches 0 as samples converge to the target distribution.

density ratio estimation, we introduce a novel algorithm for training deep density ratio estimators and show experimentally for the first time that gradient flow methods can scale to image dimensions as high as 128×128. Our gradient flow is formulated entirely in the data space, with no need for additional generator networks. The density ratio formulation allows DGGF to be composed with external density ratio estimators, which we show allows us to utilize pretrained external classifiers for class-conditional generation. In addition, we demonstrate that DGGF can be viewed as estimating the KL divergence of samples over the flow, providing it with an innate diagnostic to evaluate sample quality that also enhances model interpretability. We also show a simple technique of leveraging data-dependent priors to boost generative performance. Finally, by leveraging the freedom of choosing the source and target distributions, we show DGGF can be applied to unpaired image-to-image translation with no modifications to the framework.

2 Background

In the following, we give a brief overview of gradient flows and density ratio estimation. For a more comprehensive introduction to gradient flows, please refer to Santambrogio (2017). A thorough overview of density ratio estimation can be found in Sugiyama et al. (2012).

Wasserstein Gradient Flows. To motivate the concept of gradient flows, we consider Euclidean space equipped with the familiar L_2 distance metric $(\mathcal{X}, \|\cdot\|_2)$. Given a function $F : \mathcal{X} \to \mathbb{R}$, the curve $\{\mathbf{x}(t)\}_{t\in\mathbb{R}^+}$ that follows the direction of steepest descent is called the *gradient flow* of F:

$$\mathbf{x}'(t) = -\nabla F(\mathbf{x}(t)). \tag{1}$$

In generative modeling, we are interested in sampling from the *probability distribution* of a given dataset. Hence, instead of Euclidean space, we consider the space of *probability measures* with finite second moments equipped with the 2-Wasserstein metric $(\mathcal{P}_2(\Omega), \mathcal{W}_2)$. Given a functional $\mathcal{F} : \mathcal{P}_2(\Omega) \to \mathbb{R}$ in the 2-Wasserstein space, the gradient flow of $\mathcal{F}$ is the steepest descent curve of $\mathcal{F}$. We call such curves Wasserstein gradient flows (WGF).

Density Ratio Estimation via Bregman Divergence. Let $q(\mathbf{x})$ and $p(\mathbf{x})$ be two distributions over $\mathcal{X} \in \mathbb{R}^d$ where we have access to i.i.d samples $\mathbf{x}_q \sim q(\mathbf{x})$ and $\mathbf{x}_p \sim p(\mathbf{x})$. The goal of density ratio estimation (DRE) is to estimate the true density ratio $r^*(\mathbf{x}) = \frac{q(\mathbf{x})}{p(\mathbf{x})}$ based on samples $\mathbf{x}_q$ and $\mathbf{x}_p$.

We will focus on density ratio fitting under the Bregman divergence (BD), which is a framework that unifies many existing DRE techniques (Sugiyama et al., 2012). Let $g : \mathbb{R}^+ \to \mathbb{R}$ be a twice continuously differentiable convex function with a bounded derivative. The BD seeks to quantify the discrepancy between the estimated density ratio r_θ and the true density ratio r^*:

$$BD_g(r^* || r_\theta) = \mathbb{E}_{p(\mathbf{x})}[g(r^*(\mathbf{x})) - g(r_\theta(\mathbf{x})) + \partial g(r_\theta(\mathbf{x}))r_\theta(\mathbf{x})] - \mathbb{E}_{q(\mathbf{x})}[\partial g(r_\theta(\mathbf{x}))]. \tag{2}$$

As we only have access to samples, we estimate the expectations in Eq. 2 using Monte Carlo estimates:

$$BD_g(r_\theta) = \frac{1}{N}\sum_{n=1}^{N}[\partial g(r_\theta(\mathbf{x}_p^{(n)}))r_\theta(\mathbf{x}_p^{(n)}) - g(r_\theta(\mathbf{x}_p^{(n)}))] - \frac{1}{N}\sum_{n=1}^{N}[\partial g(r_\theta(\mathbf{x}_q^{(n)}))], \tag{3}$$

where we drop the term $\mathbb{E}_p[g(r^*(\mathbf{x}))]$ as it does not depend on the model r_θ during optimization. The minimizer of Eq. 3, which we denote θ^*, satifies $r_{\theta^*}(x) = r^*(\mathbf{x}) = q(\mathbf{x})/p(\mathbf{x})$. For ease of notation, we will use the hatted symbol $\widehat{\mathbb{E}}_p$ to refer to Monte Carlo estimates.

3 Generative Modeling with Wasserstein Gradient Flows

This section describes our key contribution: Deep Generative Wasserstein Gradient Flows (DGGF), in which we show how to train and leverage a deep density ratio estimator for high-fidelity generative modeling. Let $q(\mathbf{x})$ and $p(\mathbf{x})$ be two distributions over $\mathcal{X} \in \mathbb{R}^d$. Assume that $p(\mathbf{x})$ is the target distribution that we wish to learn. We choose $q(\mathbf{x})$ to be a known distribution that we can sample from, such as a uniform or Gaussian distribution. Our goal is to construct a WGF starting from $q(\mathbf{x})$, such that trajectories along the flow decrease some notion of distance between $q(\mathbf{x})$ and $p(\mathbf{x})$. This will allow us to flow samples from $q(\mathbf{x})$ to $p(\mathbf{x})$; in our case, when $q(\mathbf{x})$ is a simple prior and $p(\mathbf{x})$ is the distribution of natural images, we can perform generative modeling of images.

Formally, the functional $\mathcal{F} : \mathcal{P}_2(\Omega) \to \mathbb{R}$ encodes the aforementioned notion of distance between $q(\mathbf{x})$ and $p(\mathbf{x})$. Following Ansari et al. (2021), we choose $\mathcal{F}$ to be from the family of entropy-regularized f-divergences defined as

$$\mathcal{F}_p^f(q) = \int p(\mathbf{x}) f(q(\mathbf{x})/p(\mathbf{x}))d\mathbf{x} + \gamma \int q(\mathbf{x}) \log q(\mathbf{x}) d\mathbf{x}, \tag{4}$$

where $f : \mathbb{R}^+ \to \mathbb{R}$ is a twice-differentiable convex function with $f(1) = 0$. We can understand the first term, the f-divergence, as measuring the discrepancy between $q(\mathbf{x})$ and $p(\mathbf{x})$. Popular f-divergences include the Kullback-Leibler (KL), Pearson-χ^2 divergence and Jensen-Shannon (JS) divergence. The first term of Eq. 4 thus ensures that the "distance" between $q(\mathbf{x})$ and $p(\mathbf{x})$ decreases along the gradient flow, while the second (differential entropy) term improves expressiveness and prevents the collapse of the gradient flow onto the data points (Ansari et al., 2021).

From Distributions to Particles. In this study, we are interested in functional optimization of the form $\min_{q \in \mathcal{P}(\Omega)} \mathcal{F}_p^f(q)$. One way to formulate this optimization is to construct a gradient flow in Wasserstein space. The gradient flow of the functional $\mathcal{F}_q^f(p)$ in Wasserstein space is the curve of measures $\{q_t\}_{t \in \mathbb{R}^+}$ which solves the following Fokker-Planck equation (FPE) (Ambrosio et al., 2005; Villani, 2009):

$$\partial_t q_t(\mathbf{x}) = \operatorname{div}(q_t(\mathbf{x}) \nabla_{\mathbf{x}} f'(q_t(\mathbf{x})/p(\mathbf{x})) + \gamma \nabla_{\mathbf{x}}^2 q_t(\mathbf{x}), \tag{5}$$

where div and $\nabla_{\mathbf{x}}^2$ denote the divergence and Lapacian operators respectively, and f' denotes the first derivative of f. Eq. 5 is a non-linear partial differential equation, which makes solving for the distribution q_t challenging. Instead of attempting to obtain the WGF at the population density level by solving for q_t, we can utilize the connection of FPEs with stochastic differential equations (SDE) and *simulate* the equivalent particle system described by the following SDE:

$$d\mathbf{x}_t = -\nabla_{\mathbf{x}} f'(q_t(\mathbf{x}_t)/p(\mathbf{x}_t))dt + \sqrt{2\gamma} d\mathbf{w}_t, \tag{6}$$

where $d\mathbf{w}_t$ denotes the standard Wiener process. Eq. 6 describes the stochastic evolution of a particle $\mathbf{x}_t$; in image modeling terms, $\mathbf{x}_t$ represents an image sample as it is evolved through the SDE. Eq. 6 and Eq. 5 are equivalent in the sense that the marginal distribution q_t of particles that

evolve under Eq. 6 satisfies the FPE of Eq. 5. In other words, we are able to obtain samples from q_t along the gradient flow of $\mathcal{F}_p^f(q)$ by first drawing samples $\mathbf{x}_{t_0} \sim q_{t_0}$ and then simulating the SDE in Eq. 6. Empirically, we simulate the discretized version of Eq. 6 using the Euler-Maruyama method

$$\mathbf{x}_{t_{k+1}} = \mathbf{x}_{t_k} - \eta \nabla_{\mathbf{x}} f'(q_{t_k}(\mathbf{x}_{t_k})/p(\mathbf{x}_{t_k})) + \sqrt{2\gamma\eta}\boldsymbol{\xi}_{t_k}, \tag{7}$$

where $\boldsymbol{\xi}_{t_k} \sim \mathcal{N}(0, I)$, η can be interpreted as the step size and the time interval $[0, T]$ is discretized into equal intervals $\{(t_{k-1}, t_k)\}_{k\in[1:K]}$ such that $t_0 < t_1 < \cdots < t_K$.

3.1 Gradient Flow via Density Ratio Estimation

Simulating the gradient flow in Eq. 7 requires an estimate of the density ratio $q_t(\mathbf{x})/p(\mathbf{x})$, which is typically unavailable. In this section, we demonstrate how to train a parametric density ratio estimator $r_\theta(\mathbf{x}_t) = q_t(\mathbf{x}_t)/p(\mathbf{x}_t)$ that estimates the required density ratio throughout the gradient flow, and apply it to the task of generative modeling. By using $r_\theta(\mathbf{x}_t)$ as the density ratio estimate in Eq. 7, we obtain the following form for the discretized gradient flow that is amenable to numerical simulation

$$\mathbf{x}_{t_{k+1}} = \mathbf{x}_{t_k} - \eta \nabla_{\mathbf{x}} f'(r_\theta(\mathbf{x}_{t_k})) + \nu\boldsymbol{\xi}_{t_k}, \tag{8}$$

where we combine the constants of the Gaussian noise term into a single hyperparameter ν.

We will use the Bregman divergence objective of Eq. 3 to train the deep density ratio estimator, parameterized by a neural network. The DRE objective takes the form

$$\mathcal{L}(\theta) = \widehat{\mathbb{E}}_p[\partial f(r(\mathbf{x}))r(\mathbf{x}) - f(r_\theta(\mathbf{x}))] - \widehat{\mathbb{E}}_{q_t}[\partial f(r(\mathbf{x}))]. \tag{9}$$

The expectation over $p(\mathbf{x})$ in the first term can be estimated using samples drawn from the dataset, $\mathbf{x} \sim p(\mathbf{x})$. The second term requires us to draw samples from the marginal distribution of the gradient flow $\mathbf{x}_t \sim q_t(\mathbf{x}_t)$, which we will approximate by numerically simulating Eq. 8.

Concretely, at each training step, we first draw samples from the the prior distribution $\mathbf{x} \sim q_{t_0}(\mathbf{x})$ and simulate Eq. 8 for K steps and use the resultant samples to estimate the expectation over $q_t(\mathbf{x})$. The simulation of Eq. 8 requires a time-dependent estimate of the density ratio; however, we train the density ratio estimator r_θ without explicit dependence on time. As the model improves progressively through training, samples from simulating Eq. 8 improve correspondingly. We hypothesize that minimizing Eq. 9 over the course of training with these samples of varying quality allows the network to implicitly associate sample quality with time. Thus, a single time-independent estimator is able to approximate the density ratio $q_t(\mathbf{x})/p(\mathbf{x})$ along the gradient flow. We provide a pseudocode of the training algorithm in Algorithm 1 in the appendix.

Our training scheme bears resemblance to contrastive divergence, which is used in Energy-based models (EBM). At each training step, DGGF learns to refine its estimate of the density ratio $r_\theta(\mathbf{x})$ by looking at positive samples from the data distribution and negative samples drawn from its implicit distribution by simulating Eq. 8. However, despite the similarity, DGGF is fundamentally distinct from EBMs: DGGF is formulated as the path of steepest descent in Wasserstein space, while EBMs are derived from maximum likelihood. We further elaborate on the salient distinctions from EBMs and other models in Sec. 4. Once DGGF is trained, generating new samples at test time is simply a matter of running Eq. 8 directly in data space, without additional generator networks.

Choices for f-divergences and Bregman divergence. DGGF allows for flexibility in the choice of f-divergence in the gradient flow, as well as the choice of g in the Bregman divergence objective. We consolidate a list of common f-divergences and their first derivatives that we study in this paper in Table 2 in the appendix.

Certain forms for g pair naturally with specific f-divergences for the gradient flow, as they simplify the calculations needed. In our experiments, we utilize two forms of g: the first is when $g(t) = \frac{1}{2}(t-1)^2$, which corresponds to the Least-Squares Importance Fitting (LSIF) objective:

$$\mathcal{L}_{LSIF}(\theta) = \frac{1}{2}\widehat{\mathbb{E}}_p[r_\theta(\mathbf{x})]^2 - \widehat{\mathbb{E}}_{q_t}[r_\theta(\mathbf{x})], \tag{10}$$

and the second is $g(t) = t\log t - (1+t)\log(1+t)$, which corresponds to the Logistic Regression (LR) objective:

$$\mathcal{L}_{LR}(\theta) = -\widehat{\mathbb{E}}_p\left[\log\frac{1}{1+r_\theta(\mathbf{x})}\right] - \widehat{\mathbb{E}}_{q_t}\left[\log\frac{r_\theta(\mathbf{x})}{1+r_\theta(\mathbf{x})}\right]. \tag{11}$$

Due to numerical compatibility issues that we discuss in Appendix B, we pair the LSIF objective with the Pearson-χ^2 divergence and the LR objective with the KL, JS and logD divergences. We abbreviate them as LSIF-χ^2, LR-KL, LR-JS and LR-logD respectively and study these pairings in our experiments in Sec. 5.

3.2 Density Chasm and Data-Dependent Priors

In principle, the gradient flow formulation does not place any restrictions on the prior distribution $q_{t_0}(\mathbf{x})$. In the $t \to \infty$ limit, the gradient flow is guaranteed to converge to $p(\mathbf{x})$ (see Appendix D for a proof). However, empirically, we observe that when implementing the flow with a finite number of steps, the choice of prior distribution can significantly affect model performance.

This observation may be attributed to the density chasm problem (Rhodes et al., 2020). Consider two densities $q(\mathbf{x})$ and $p(\mathbf{x})$ which differ greatly, for example as measured by their KL divergence. A binary classifier that is trained on distinguishing the two distributions, which is equivalent to density ratio estimation using the "density ratio trick", can obtain near perfect accuracy while learning a relatively poor estimate of the density ratio (Ansari et al., 2021). A simple example would be two narrow Gaussians with means that are far apart. A classifier can learn a trivial boundary with near perfect accuracy, such as a straight line between the two modes, without having to estimate the density ratio between the two distributions accurately.

We found that using common priors for $q_{t_0}(\mathbf{x})$ such as a uniform distribution led to poorer sample quality due to the large chasm between the simple distribution and the complex multimodal data distribution. Inspired by generation from seed distributions with robust classifiers (Santurkar et al., 2019), we leverage a *data-dependent* prior by fitting a simple multivariate Gaussian to the training dataset, and sampling $\mathbf{x}_0$ from the data-dependent prior

$$q_{t_0}(\mathbf{x}) = \mathcal{N}(\boldsymbol{\mu}_D, \boldsymbol{\Sigma}_D), \quad \text{where } \boldsymbol{\mu}_D = \mathbb{E}_D[\mathbf{x}], \quad \boldsymbol{\Sigma}_D = \mathbb{E}_D[(\mathbf{x} - \boldsymbol{\mu}_D]^T(\mathbf{x} - \boldsymbol{\mu}_D)] \tag{12}$$

where D represents the training dataset. Samples from this prior are poor but this approach is sufficient to cross the density chasm; we visualize some samples from the prior in Fig. 8, and experimental results support this approach (Sec. 5.2).

4 Related Works

A well-known method to simulate WGFs at the population density level is using the Jordan, Kinderlehrer, and Otto (JKO) scheme (Jordan et al., 1998), which approximates the dynamics of the Fokker-Planck equation through an iterative time-discretization update. The JKO scheme requires computation of the 2-Wasserstein distance and a free energy functional. The former is intractable due to the minimization over the space of probability measures. Several works have been proposed to circumvent this problem: Mokrov et al. (2021) leverages Brenier's theorem and convex neural networks to approximately estimate $\mathcal{W}_2$, while Fan et al. (2021) leverages the Fenchel conjugate to evaluate the free energy functional such as f-divergences. Our work avoids the JKO scheme entirely by adopting a particle-based approach, where we use the Euler-Maruyama method to simulate the gradient flow, parameterized by neural networks. In this way, we avoid the need to estimate $\mathcal{W}_2$, as well as the dual optimization needed in the variational formulation of Fan et al. (2021).

More closely related to DGGF are other particle-based WGF approaches, notably DG$f$low (Ansari et al., 2021), VGrow (Gao et al., 2019) and Euler Particle Transport (EPT) (Gao et al., 2022). DG$f$low tackles the problem of sample refinement by constructing the WGF using pretrained GAN discriminators as density ratio estimators. DGGF can be seen as a general case of DG$f$low, where we recover DG$f$low if we fix the prior q_{t_0} to be the implicit distribution defined by the GAN generator. When q_{t_0} is chosen as a simple prior, DGGF can perform generation from scratch, as opposed to purely refinement. VGrow and EPT solve the WGF problem via application of iterative residual maps parameterized by deep density ratio estimators. Similarly to DGGF, these estimators are trained using the Bregman divergence and applied to the task of unconditional image generation. There are several differences that distinguish DGGF from VGrow and EPT. As VGrow can be seen as a specific instance of EPT with the logistic regression objective, we focus on EPT while noting that our discussion applies to both methods. EPT utilizes a formalism that performs generation on a predetermined batch of particles concurrently with the training of the density ratio estimator. When

formulated directly in data space, EPT is unable to generate new samples at test time without retraining from scratch. In the same vein, EPT also has poor generative performance when formulated in the data space[1]. It was necessary to amortize the cost of generation to an auxiliary generator, and formulate the WGF in latent space by assuming that data lies on a lower-dimensional manifold. By flowing the latent codes and matching the generator's transformation of the codes to data, the auxiliary generator can perform generation at test time. We show from our results in Sec. 5 that such assumptions are unnecessary for particle-based WGF methods and generation can be done entirely in the data space, without the complexity of training an additional generator.

Our model also shares similarities with EBMs (LeCun et al., 2006) and score-based models (SGM) (Song & Ermon, 2019). Both models perform generation via Langevin dynamics, which require the data score $\nabla_{\mathbf{x}} \log p(\mathbf{x})$. In the former, a neural network is trained to estimate the energy of a Boltzmann distribution, while the latter is trained to predict the score directly. SGMs are trained with a denoising autoencoder objective, while EBMs are trained using contrastive divergence. Fundamentally, the key difference is that DGGF is formulated by finding the path of greatest descent in Wasserstein space, while both EGMs and SGMs can be derived from the perspective of maximum likelihood. As shown in prior works (Jordan et al., 1998; Liu et al., 2019), Langevin dynamics can be derived from Wasserstein gradient flows when f is chosen to be the KL divergence, which we reproduce in Appendix E. The density ratio formulation of DGGF allows us the flexibility of composition with other density ratio estimators, as well as the freedom of choosing the source and target distributions of the gradient flow, which opens up the possibility of applications such as composition with pretrained classifiers (Sec. 5.3) and image-to-image translation (Sec. 5.4). These are not directly applicable in the EBM and SGM formulation, and adaptations of both to such domains require significant theoretical and experimental modifications to the generative framework (Zhao & Chen, 2021; Sasaki et al., 2021; Song et al., 2020).

5 Experiments

In this section, we present empirical results of the generative performance of DGGF on various common image datasets across different resolutions. Our goal is to demonstrate that DGGF is able to generate diverse, high-quality images across various f-divergences and Bregman divergence objectives, as well as generalizes to other tasks such as class-conditional generation and image-to-image translation. We will use DGGF-DDP to abbreviate experiments with the data-dependent prior and DGGF-UP for ablation experiments with the uniform prior.

Table 1: CIFAR10 and CelebA scores.

Model	FID ↓
CIFAR10 32 × 32	
EPT (no outer loop) (Gao et al., 2022)	46.63
IGEBM (Du & Mordatch, 2019)	40.58
SNGAN (Miyato et al., 2018)	21.7
PixelCNN (Van Oord et al., 2016)	65.93
NVAE (Vahdat & Kautz, 2020)	51.67
NCSN (Song & Ermon, 2019)	25.32
DGGF-DDP (LSIF-χ^2)	28.12
DGGF-DDP (LR-KL)	28.80
DGGF-DDP (LR-JS)	29.92
DGGF-DDP (LR-logD)	30.72
DGGF-UP (LSIF-χ^2)	35.11
DGGF-UP (LR-KL)	39.90
CelebA 64 × 64	
NCSN (Song & Ermon, 2019)	26.89
NVAE (Vahdat & Kautz, 2020)	14.74
EBM-SR (Nijkamp et al., 2019)	23.02[2]
DGGF-DDP (LSIF-χ^2)	22.42
DGGF-DDP (LR-KL)	22.88

5.1 Setup

We test DGGF with unconditional generation on CIFAR10, CelebA and LSUN Church datasets, class-conditional generation on CIFAR10 and image-to-image translation on the Cat2dog dataset. All pixel values are normalized to the range [-1, 1]. For CIFAR10, we keep the resolution of the images at 32×32, while for CelebA and LSUN Church we resize them to 64×64 and 128×128, respectively. We use modified ResNet architectures for all experiments in this study. See Appendix G for more details.

[1] Based on experiments of EPT in Sec. 5 and on unpublished scores cited by the authors; please see https://openreview.net/forum?id=awMgJJ9H-0q.

[2] Score as reported in Gao et al. (2020).

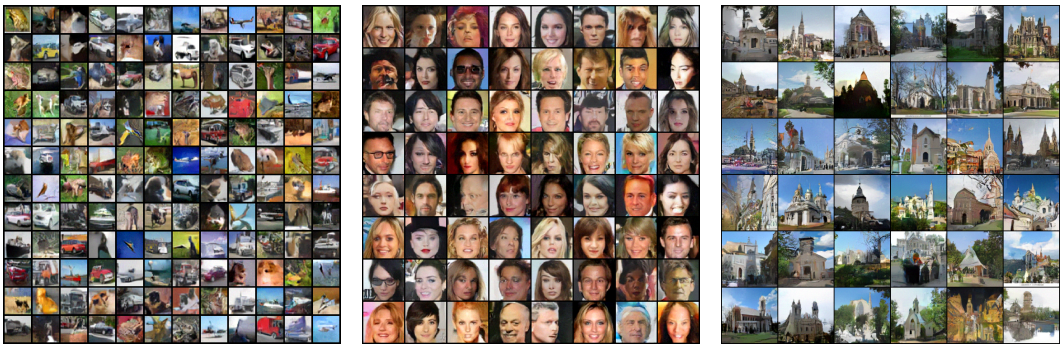

Figure 2: Samples from DGGF-DDP on CIFAR10 32^2, CelebA 64^2 and LSUN Church 128^2 using LSIF-χ^2. More results using various BD objectives and f-divergences can be found in Appendix.

Figure 3: Illustration of the gradient flow process for LSUN Church, starting from a sample from the data-dependent prior on the leftmost column.

5.2 Image Generation

In Fig. 2, we show uncurated samples of DGGF-DDP on different combinations of g and f-divergences. Visually, our model is able to produce high-quality samples on a variety of datasets up to resolutions of 128×128, surpassing existing gradient flow techniques (Gao et al., 2019; 2022). More samples with other f-divergences can be found in Appendix I. In Table 1, we show the FID scores of DGGF-DDP in comparison with relevant baselines utilizing different generative approaches. On CIFAR10, our model performs comparably with SNGAN and the score-based NCSN, while outperforming the baseline EBM and autoregressive method PixelCNN. Most notably, our method strongly outperforms EPT without an auxiliary generator (no outer loop), which is the gradient flow baseline utilizing a density ratio method. For CelebA, our model is outperformed by the state-of-the-art variational autoencoder NVAE, but outperforms NCSN and the short-run EBM approach.

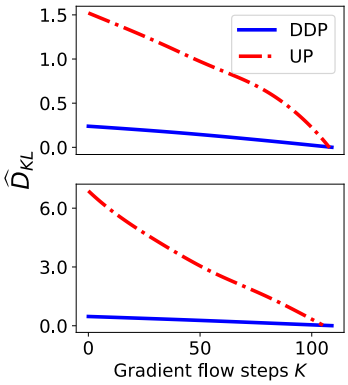

Figure 4: KL estimate $\widehat{D}_{KL}$ of the gradient flow for CIFAR10 generation for both DDP and UP for LSIF-χ^2 (upper) and LR-KL (lower).

To provide intuition for the gradient flow process, we provide intermediate samples for the LSUN Church dataset in Fig. 3, which visualizes how samples drawn from the data-dependent prior is evolved to a high quality sample. In this scenario, the prior contains low level features such as the color of the sky and rough silhouette of the building. The gradient flow retains these coarse features, but generates the higher frequency details necessary to create a realistic image. This matches the intuition that the gradient flow is the steepest descent curve, as the shortest path would mean changing as little of the image as possible to reach the target distribution. We also visualize samples obtained by interpolating in the prior space for CelebA in Fig. 5. Despite the use of a relatively complex prior, the model is able to smoothly interpolate in the latent space, indicating the model has learnt a semantically relevant latent representation that is characteristic of

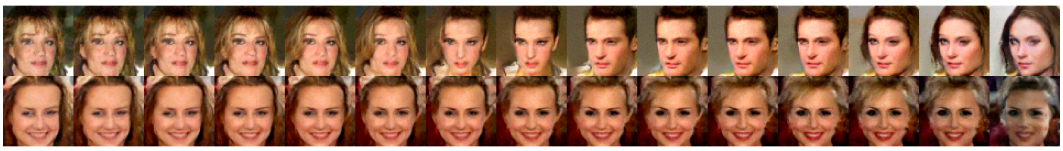

Figure 5: Interpolation results between leftmost and rightmost samples with CelebA.

a valid generative model. Finally, to verify that our model has not merely memorized the dataset, particularly due to concerns that the prior is fitted from data, we show nearest neighbor samples of the generated images in the training set in Appendix H. We can clearly see that samples produced by DGGF-DDP are distinct from the closest samples in the training set, which tells us that DGGF-DDP is capable of generating new and diverse samples beyond the data it was trained on.

Model Estimate of KL over Flow As our model outputs the density ratio $r_\theta(\mathbf{x}_{t_k}) = q_{t_k}(\mathbf{x}_{t_k})/p(\mathbf{x}_{t_k})$ throughout the gradient flow, our model can be interpreted as estimating the KL divergence $D_{KL}(q_{t_k}||p) = \int q_{t_k}(\mathbf{x}_{t_k}) \log(q_{t_k}/p) d\mathbf{x}_{t_k} = \mathbb{E}_{\mathbf{x}_{t_k}}[\log r_\theta(\mathbf{x}_{t_k})]$ where the expectation is taken over the batch of samples being evolved. We denote this estimate as $\widehat{D}_{KL}$. We show that this estimate is valid in Fig. 4, which plots $\widehat{D}_{KL}$ over the gradient flow for both DDP and UP (the UP results serve as ablations in the next paragraph). Focusing on the DDP results, we observe that $\widehat{D}_{KL}$ decreases monotonically over the flow, which agrees with the notion that $q_{t_k}(\mathbf{x}_{t_k})$ approaches $p(\mathbf{x}_{t_k})$, as seen in Fig. 3. Intuitively, we wish to terminate the flow when $\widehat{D}_{KL} = 0$, which indicates convergence to $p(\mathbf{x})$. Interestingly, this result is achieved at approximately $K = 110$ steps, despite setting $K = 100$ during training. This motivates us to use $K = 110$ when evaluating our DDP experiments, which we found to improve FID scores slightly. As such, the observation that decreasing $\widehat{D}_{KL}$ corresponds to improving sample quality, and its natural interpretation as the KL estimate, provides DGGF with an interpretable diagnostic of its own sample quality that comes innately with the density ratio formulation.

Ablations with Uniform Prior We motivate the use of the data-dependent prior by conducting ablation experiments with q_{t_0} being a uniform prior, $\mathbf{x}_{t_0} \sim U[-1, 1]$. All hyperparameters are kept identical to DDP experiments to isolate the effects of the choice of prior distribution. We include qualitative samples of DGGF-UP in Fig. 18 in the appendix. The quantitative results can be seen from the FID scores in Table 1. Visually, DGGF-UP produces diverse and appealing samples even with a uniform prior. However, when comparing quantitative scores we observe that the use of DDP improves results significantly. Support for the density chasm hypothesis can be found by comparing the $\widehat{D}_{KL}$ of DDP and UP in Fig. 4. For both LSIF-χ^2 and LR-KL, DGGF-UP has a significantly larger $\widehat{D}_{KL}$ at the start of the flow as compared to DGGF-DDP. This corresponds to our intuition, as the uniform prior is 'farther' from the data distribution as compared to the data-dependent prior. As a result, the model has to push the particles over a larger density chasm, leading to poorer performance given a fixed computational budget. Interestingly, DGGF-UP reaches $\widehat{D}_{KL} = 0$ at a smaller number of steps than DGGF-DDP, although it is still larger than the $K = 100$ at training time.

5.3 Conditional Generation with Robust Classifiers

The density ratio framework allows us to compose different density ratio estimators together, therefore allowing us to construct WGFs between distributions different from those in training. To illustrate this, consider a multiclass classifier which predicts the probability that a given image belongs to any one of N classes. We show in Appendix F that we can express such classifiers as a density ratio $p(y = n|\mathbf{x}) = N^{-1}p(\mathbf{x}|y = n)/p(\mathbf{x})$. We can thus obtain a conditional density ratio estimator $r_\theta(\mathbf{x}_t|y = n) = q_t(\mathbf{x}_t)/p(\mathbf{x}_t|y = n)$ by composing our unconditional estimator $r_\theta(\mathbf{x}_t)$ with the classifier output (see Appendix F):

$$r_\theta(\mathbf{x}_t|y = n) = \frac{1}{N} r_\theta(\mathbf{x}_t) p(y = n|\mathbf{x}_t)^{-1}. \tag{13}$$

When $r_\theta(\mathbf{x}_t|y = n)$ is used in simulating the WGF in Eq. 8, we obtain a class-conditional generative model. This is conceptually similar to the idea proposed in Song et al. (2020), where an uncondi-

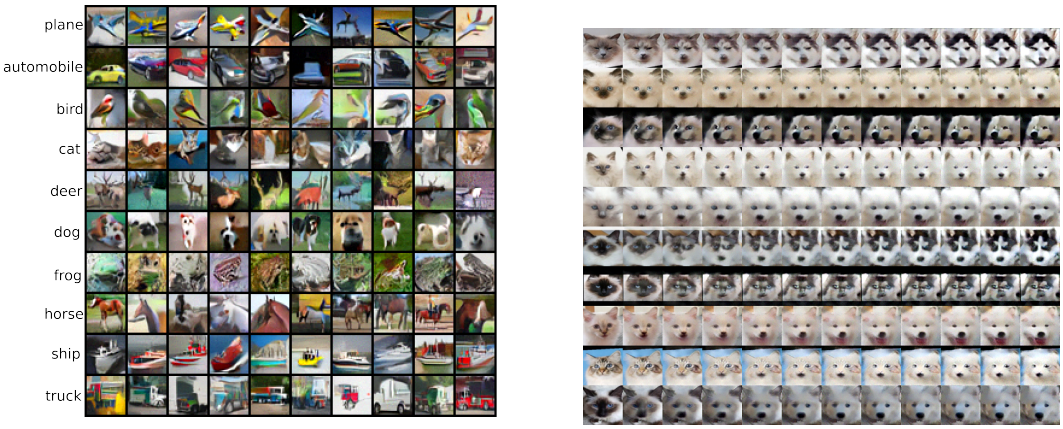

(a) (b)

Figure 6: (a) Class-conditional samples from composition with a robust classifier. (b) Image-to-image translation process from cat to dog images using DGGF.

tional score model $\nabla_{\mathbf{x}} \log p_t(\mathbf{x}(t))$ is composed with a time-dependent classifier $\nabla_{\mathbf{x}} \log p_t(y|\mathbf{x}(t))$ to form a class-conditional model. However, whereas Song et al. (2020) requires the separate training of a time-dependent classifier, our formulation allows us to use off-the-shelf pretrained classifiers with *no further retraining*. Inspired by earlier work on image synthesis with robust classifiers (Santurkar et al., 2019), we found that using a pretrained adversarially-robust classifier was necessary in obtaining useful gradients for the gradient flow. We show our results in Fig. 6a, where each row represents conditional samples of each class in the CIFAR10 dataset.

5.4 Unpaired Image-to-image Translation

Our framework can also be directly applied to unpaired image-to-image-translation (I2I). We simply fix the prior distribution $q_{t_0}(\mathbf{x})$ to a source domain and $p(\mathbf{x})$ to a target domain. We then train the model in exactly the same manner as unconditional generation (see Algorithm 1).

We test our I2I model on the Cat2dog dataset (Lee et al., 2018). From Fig. 6b, DGGF is able to smoothly translate images of cats to dogs while maintaining relevant semantic features of the image. For example, the background colors of the image and the pose of the cat are unchanged—a cat that is facing a certain direction is translated to a dog that faces the same direction. We also observe that the facial tones of the cat is preserved—a cat with light fur is translated to a dog with light fur. CycleGAN (Zhu et al., 2017) achieves better FID scores than DGGF (Table 4 in the appendix) but like many I2I methods (Lee et al., 2018; Choi et al., 2020; Zhao & Chen, 2021; Nie et al., 2021), CycleGAN incorporates specific inductive biases, such as dual generators and discriminators together with the cycle-consistency loss. Incorporating such inductive biases into the gradient flow process can improve the translation and would make for interesting future work.

6 Conclusion

In this paper, we proposed DGGF, a method to simulate the Wasserstein gradient flow between two distributions that minimizes the entropy-regularized f-divergence. As constructing such a flow requires an estimate of the density ratio, we showed how to leverage the Bregman divergence to train a deep density ratio estimator that is able to synthesize diverse images of high quality. We showed that the modularity of the density ratio-based WGF framework allows for composition with external density ratio estimators, as well as direct application to tasks such as unpaired image-to-image translation. Given the flexibility of choosing the source and target distribution, future work can investigate different choices of the two distributions that could correspond to entirely new applications. Another possible avenue is investigating how we can incorporate task-specific inductive biases into the gradient flow process, which should allow the WGF to perform better on the problem at hand.

REFERENCES

Stefan Adams, Nicolas Dirr, Mark A Peletier, and Johannes Zimmer. From a large-deviations principle to the wasserstein gradient flow: a new micro-macro passage. *Communications in Mathematical Physics*, 307(3):791–815, 2011.

Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. *Gradient flows: in metric spaces and in the space of probability measures*. Springer Science & Business Media, 2005.

Abdul Fatir Ansari, Ming Liang Ang, and Harold Soh. Refining deep generative models via discriminator gradient flow. In *ICLR*, 2021.

Espen Bernton. Langevin monte carlo and jko splitting. In *Conference On Learning Theory*, pp. 1777–1798. PMLR, 2018.

José Antonio Carrillo, Katy Craig, and Francesco S Patacchini. A blob method for diffusion. *Calculus of Variations and Partial Differential Equations*, 58(2):1–53, 2019.

Lenaic Chizat and Francis Bach. On the global convergence of gradient descent for over-parameterized models using optimal transport. *Advances in neural information processing systems*, 31, 2018.

Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. Stargan v2: Diverse image synthesis for multiple domains. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 8188–8197, 2020.

Yilun Du and Igor Mordatch. Implicit generation and generalization in energy-based models. *arXiv preprint arXiv:1903.08689*, 2019.

Logan Engstrom, Andrew Ilyas, Shibani Santurkar, and Dimitris Tsipras. Robustness (python library), 2019. URL https://github.com/MadryLab/robustness.

Jiaojiao Fan, Amirhossein Taghvaei, and Yongxin Chen. Variational wasserstein gradient flow. *arXiv preprint arXiv:2112.02424*, 2021.

Ruiqi Gao, Yang Song, Ben Poole, Ying Nian Wu, and Diederik P Kingma. Learning energy-based models by diffusion recovery likelihood. *arXiv preprint arXiv:2012.08125*, 2020.

Yuan Gao, Yuling Jiao, Yang Wang, Yao Wang, Can Yang, and Shunkang Zhang. Deep generative learning via variational gradient flow. In *International Conference on Machine Learning*, pp. 2093–2101. PMLR, 2019.

Yuan Gao, Jian Huang, Yuling Jiao, Jin Liu, Xiliang Lu, and Zhijian Yang. Deep generative learning via euler particle transport. In *Mathematical and Scientific Machine Learning*, pp. 336–368. PMLR, 2022.

Richard Jordan, David Kinderlehrer, and Felix Otto. The variational formulation of the fokker–planck equation. *SIAM journal on mathematical analysis*, 29(1):1–17, 1998.

Yann LeCun, Sumit Chopra, Raia Hadsell, M Ranzato, and F Huang. A tutorial on energy-based learning. *Predicting structured data*, 1(0), 2006.

Hsin-Ying Lee, Hung-Yu Tseng, Jia-Bin Huang, Maneesh Singh, and Ming-Hsuan Yang. Diverse image-to-image translation via disentangled representations. In *Proceedings of the European conference on computer vision (ECCV)*, pp. 35–51, 2018.

Chang Liu, Jingwei Zhuo, and Jun Zhu. Understanding mcmc dynamics as flows on the wasserstein space. In *International Conference on Machine Learning*, pp. 4093–4103. PMLR, 2019.

Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. *arXiv preprint arXiv:1802.05957*, 2018.

Petr Mokrov, Alexander Korotin, Lingxiao Li, Aude Genevay, Justin M Solomon, and Evgeny Burnaev. Large-scale wasserstein gradient flows. *Advances in Neural Information Processing Systems*, 34:15243–15256, 2021.

Weili Nie, Arash Vahdat, and Anima Anandkumar. Controllable and compositional generation with latent-space energy-based models. *Advances in Neural Information Processing Systems*, 34:13497–13510, 2021.

Erik Nijkamp, Mitch Hill, Song-Chun Zhu, and Ying Nian Wu. Learning non-convergent non-persistent short-run mcmc toward energy-based model. *Advances in Neural Information Processing Systems*, 32, 2019.

Benjamin Rhodes, Kai Xu, and Michael U Gutmann. Telescoping density-ratio estimation. *Advances in neural information processing systems*, 33:4905–4916, 2020.

Pierre H Richemond and Brendan Maginnis. On wasserstein reinforcement learning and the fokker-planck equation. *arXiv preprint arXiv:1712.07185*, 2017.

Filippo Santambrogio. {Euclidean, metric, and Wasserstein} gradient flows: an overview. *Bulletin of Mathematical Sciences*, 7(1):87–154, 2017.

Shibani Santurkar, Andrew Ilyas, Dimitris Tsipras, Logan Engstrom, Brandon Tran, and Aleksander Madry. Image synthesis with a single (robust) classifier. *Advances in Neural Information Processing Systems*, 32, 2019.

Hiroshi Sasaki, Chris G Willcocks, and Toby P Breckon. Unit-ddpm: Unpaired image translation with denoising diffusion probabilistic models. *arXiv preprint arXiv:2104.05358*, 2021.

Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in Neural Information Processing Systems*, 32, 2019.

Yang Song and Stefano Ermon. Improved techniques for training score-based generative models. *Advances in neural information processing systems*, 33:12438–12448, 2020.

Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020.

Masashi Sugiyama, Taiji Suzuki, and Takafumi Kanamori. *Density ratio estimation in machine learning*. Cambridge University Press, 2012.

Arash Vahdat and Jan Kautz. Nvae: A deep hierarchical variational autoencoder. *Advances in Neural Information Processing Systems*, 33:19667–19679, 2020.

Aaron Van Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. In *International conference on machine learning*, pp. 1747–1756. PMLR, 2016.

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

Max Welling and Yee W Teh. Bayesian learning via stochastic gradient langevin dynamics. In *Proceedings of the 28th international conference on machine learning (ICML-11)*, pp. 681–688. Citeseer, 2011.

Ruiyi Zhang, Changyou Chen, Chunyuan Li, and Lawrence Carin. Policy optimization as wasserstein gradient flows. In *International Conference on Machine Learning*, pp. 5737–5746. PMLR, 2018.

Yang Zhao and Changyou Chen. Unpaired image-to-image translation via latent energy transport. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 16418–16427, 2021.

Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pp. 2223–2232, 2017.