



OPEN Taming hyperparameter tuning in continuous normalizing flows using the JKO scheme

Alexander Vidal^{1✉}, Samy Wu Fung², Luis Tenorio¹, Stanley Osher³ & Levon Nurbekyan³

A normalizing flow (NF) is a mapping that transforms a chosen probability distribution to a normal distribution. Such flows are a common technique used for data generation and density estimation in machine learning and data science. The density estimate obtained with a NF requires a change of variables formula that involves the computation of the Jacobian determinant of the NF transformation. In order to tractably compute this determinant, continuous normalizing flows (CNF) estimate the mapping and its Jacobian determinant using a neural ODE. Optimal transport (OT) theory has been successfully used to assist in finding CNFs by formulating them as OT problems with a soft penalty for enforcing the standard normal distribution as a target measure. A drawback of OT-based CNFs is the addition of a hyperparameter, α , that controls the strength of the soft penalty and requires significant tuning. We present JKO-Flow, an algorithm to solve OT-based CNF without the need of tuning α . This is achieved by integrating the OT CNF framework into a Wasserstein gradient flow framework, also known as the JKO scheme. Instead of tuning α , we repeatedly solve the optimization problem for a fixed α effectively performing a JKO update with a time-step α . Hence we obtain a “divide and conquer” algorithm by repeatedly solving simpler problems instead of solving a potentially harder problem with large α .

A normalizing flow (NF) is a type of generative modeling technique that has shown great promise in applications arising in physics^{1–3} as a general framework to construct probability densities for continuous random variables in high-dimensional spaces^{4–6}. An NF provides a C^1 -diffeomorphism f (i.e., a normalizing transformation) that transforms the density ρ_0 of an initial distribution P_0 to the density ρ_1 of the standard multivariate normal distribution P_1 —hence the term “normalizing”. Given such mapping f , the density ρ_0 can be recovered from the Gaussian density via the change of variables formula,

$$\log \rho_0(x) = \log \rho_1(f(x)) + \log |\det J_f(x)|, \quad (1)$$

where $J_f \in \mathbb{R}^{d \times d}$ is the Jacobian of f . Moreover, one can obtain samples with density ρ_0 by pushing forward Gaussian samples via f^{-1} .

Remark 1 Throughout the paper we slightly abuse notation, using the same notation for probability distributions and their density functions. Additionally, given a probability distribution P_0 on \mathbb{R}^d and a measurable mapping $f: \mathbb{R}^d \rightarrow \mathbb{R}^d$, we define the pushforward distribution of P_0 through f as $(f\#P_0)(B) = P_0(f^{-1}(B))$ for all Borel measurable $B \subseteq \mathbb{R}^d$ ^{7,8}.

There are two classes of normalizing flows: finite and continuous. A finite flow is defined as a composition of a finite number of C^1 -diffeomorphisms: $f = f_1 \circ f_2 \circ \dots \circ f_n$. To make finite flows computationally tractable, each f_i is chosen to have some regularity properties such as a Jacobian with a tractable determinant; for example, J_{f_i} may have a triangular structure^{9–11}.

On the other hand, continuous normalizing flows (CNFs) estimate f using a neural ODE of the form¹²:

$$\partial_t z(x, t) = v_\theta(z(x, t), t), \quad z(x, 0) = x, \quad 0 \leq t \leq T, \quad (2)$$

¹Department of Applied Mathematics and Statistics, Colorado School of Mines, Golden, USA. ²Department of Applied Mathematics and Statistics, Department of Computer Science, Colorado School of Mines, Golden, USA. ³Department of Mathematics, University of California, Los Angeles, USA. ✉email: vidal@mines.edu

where θ are the parameters of the neural ODE. In this case, f is defined as $f(x) = z(x, T)$ (for simplicity, we remove the dependence of z on θ).

One of the main advantages of CNFs is that we can tractably estimate the log-determinant of the Jacobian using Jacobi's identity, which is commonly used in fluid mechanics (see, e.g.,⁸, p. 114):

$$\partial_t \log |\det \nabla_x z(x, t)| = \nabla_z \cdot v_\theta(z(x, t), t) = \text{trace}(\nabla_z v_\theta(z(x, t), t)). \quad (3)$$

This is computationally appealing as one can replace the expensive determinant calculation by a more tractable trace computation of $\nabla_z v_\theta(z(x, t), t)$. Importantly, no restrictions on $\nabla_z v_\theta(z(x, t), t)$ (e.g., diagonal or triangular structure) are needed; thus, these Jacobians are also referred to as “free-form Jacobians”¹³.

The goal in training a CNF is to find parameters, θ , such that $f = z(\cdot, T)$ leads to a good approximation of ρ_1 or, assuming f is invertible, the pushforward of ρ_1 through f^{-1} is a good approximation of ρ_0 ^{5,6,10,13}. Indeed, let $\hat{\rho}_0$ be this pushforward density obtained with a CNF f ; that is, $\hat{\rho}_0 = f^{-1\#}\rho_1$. We then minimize the Kullback-Leibler (KL) divergence from $\hat{\rho}_0$ to ρ_0 given by

$$\min_{\theta} \mathbb{E}_{x \sim \rho_0} \log(\rho_0(x)/\hat{\rho}_0(x)) = \min_{\theta} \mathbb{E}_{x \sim \rho_0} [\log \rho_0(x) - \log \rho_1(z(x, T)) - \ell(x, T)],$$

where $\ell(x, T) = \log |\det \nabla z(x, T)|$. Dropping the θ -independent term $\log \rho_0$ and using Eqs. (2) and (3), this previous optimization problem reduces to the minimization problem

$$\min_{\theta} \mathbb{E}_{x \sim \rho_0} C(x, T), \quad C(x, T) := -\log \rho_1(z(x, T)) - \ell(x, T) \quad (4)$$

subject to ODE constraints

$$\partial_t \begin{bmatrix} z(x, t) \\ \ell(x, t) \end{bmatrix} = \begin{bmatrix} v_\theta(z(x, t), t) \\ \text{trace}(\nabla_z v_\theta(z(x, t), t)) \end{bmatrix}, \quad \begin{bmatrix} z(x, 0) \\ \ell(x, 0) \end{bmatrix} = \begin{bmatrix} x \\ 0 \end{bmatrix}. \quad (5)$$

The ODE Eq. (5) might be stiff for certain values of θ , leading to extremely long computation times. Indeed, the dependence of v on θ is highly nonlinear and might generate vector fields that lead to highly oscillatory trajectories with complex geometry.

Some recent work leverages optimal transport theory to find the CNF^{14,15}. In particular, a kinetic energy regularization term (among others) is added to the loss to “encourage straight trajectories” $z(x, t)$. That is, the flow is trained by solving the following minimization instead of Eq. (4):

$$\min_{\theta} \mathbb{E}_{x \sim \rho_0} \int_0^T \frac{1}{2} \|v_\theta(z(x, t), t)\|^2 dt + \alpha C(x, T) \quad (6)$$

subject to Eq. (5). The key insight in^{14,15} is that Eq. (4) is an example of a degenerate OT problem with a soft terminal penalty and without a transportation cost. The first term in the objective function in Eq. (6) given by the time integral is the transportation cost, whereas α is a hyperparameter that balances the soft penalty and the transportation cost. Including this cost makes the problem well-posed by forcing the solution to be unique¹⁶. Additionally, it enforces straight trajectories so that Eq. (5) is not stiff. Indeed^{14,15} empirically demonstrate that including optimal transport theory leads to faster and more stable training of CNFs. Intuitively, we minimize the KL divergence *and* the arclength of the trajectories.

Although including optimal transport theory into CNFs has been very successful^{14,15,17,18}, there are two key challenges that render them difficult to train. First, estimating the log-determinant in Eq. (4) via the trace in Eq. (5) is still computationally taxing and commonly used methods rely on stochastic approximations^{13,14}, which add extra error. Second, including the kinetic energy regularization requires tuning of the hyperparameter α . Indeed, if α is chosen too small in Eq. (6), then the kinetic regularization term dominates the training process, and the optimal solution consists of not moving, i.e., $f(x) = x$. On the other hand, if α is chosen too large, we return to the original setting where the problem is ill-posed, i.e., there are infinitely many solutions. Finally, finding an “optimal” α is problem dependent and requires tuning on a case-by-case basis.

Our contribution. We present JKO-Flow, an optimal transport-based algorithm for training CNFs without the need to tune the hyperparameter α in Eq. (6). Our approach also leverages fast numerical methods for exact trace estimation from the recently developed optimal transport flow (OT-Flow)^{15,19}.

The key idea is to integrate the OT-Flow approach into a Wasserstein gradient flow framework, also known as the Jordan, Kinderlehrer, and Otto (JKO) scheme²⁰. Rather than tuning the hyperparameter α (commonly done using a grid search), the idea is to simply pick any α and solve a sequence of “easier” OT problems that gradually approach the target distribution. Each solve is precisely a gradient descent in the space of distributions, a Wasserstein gradient descent, and the scheme provably converges to the desired distribution for all $\alpha > 0$ ²¹. Our experiments show that our proposed approach is effective in generating higher quality samples (and density estimates) and also allows us to reduce the number of parameters required to estimate the desired flow.

Our strategy is reminiscent of debiasing techniques commonly used in inverse problems. Indeed, the transportation cost that serves as a regularizer in Eq. (6) introduces a bias—the smaller α the more bias is introduced (see, e.g.,²²), so good choices of α tend to be larger. One way to remove the bias and avoid the need to tune the regularization parameter is to perform a sequence of Bregman iterations^{23,24}, also known as nonlinear proximal steps. Hence our approach reduces to debiasing via Bregman or proximal steps in the Wasserstein space. In the context of CNF training, Bregman iterations are advantageous due to the flexibility of the choice for α . Indeed,

the resulting loss function is non-convex and its optimization tends to get harder for large α . Thus, instead of solving one harder problem we solve several “easier” problems.

Optimal transport background and connections to CNFs

Denote by $\mathcal{P}_2(\mathbb{R}^d)$ the space of Borel probability measures on \mathbb{R}^d with finite second-order moments, and let $\rho_0, \rho_1 \in \mathcal{P}_2(\mathbb{R}^d)$. The quadratic optimal transportation (OT) problem (which also defines the Wasserstein metric W_2) is then formulated as

$$W_2^2(\rho_0, \rho_1) = \inf_{\pi \in \Gamma(\rho_0, \rho_1)} \int_{\mathbb{R}^{2d}} \|x - y\|^2 d\pi(x, y), \tag{7}$$

where $\Gamma(\rho_0, \rho_1)$ is the set of probability measures $\pi \in \mathcal{P}(\mathbb{R}^{2d})$ with fixed x and y -marginal distributions ρ_0 and ρ_1 , respectively. Hence the cost of transporting a unit mass from x to y is $\|x - y\|^2$, and one attempts to transport ρ_0 to ρ_1 as cheaply as possible. In Eq. (7), π represents a *transportation plan*, and $\pi(x, y)$ is the mass being transported from x to y . One can prove that $(\mathcal{P}_2(\mathbb{R}^d), W_2)$ is a complete separable metric space⁸. OT has recently become a very active research area in PDE, geometry, functional inequalities, economics, data science and elsewhere partly due to equipping the space of probability measures with a (Riemannian) metric^{8,16,25,26}.

As observed in prior works, there are many similarities between OT and NFs^{14,15,18,27}. This connection becomes more transparent when considering the dynamic formulation of Eq. (7). More precisely, the Benamou-Brenier formulation of the OT problem is given by²⁸:

$$\begin{aligned} \frac{T}{2} W_2^2(\rho_0, \rho_1) = \inf_{v, \rho} & \int_0^T \int_{\mathbb{R}^d} \frac{1}{2} \|v(x, t)\|_2^2 \rho(x, t) dx dt \\ \text{s.t.} & \quad \partial_t \rho(x, t) + \nabla \cdot (\rho(x, t)v(x, t)) = 0 \\ & \quad \rho(x, 0) = \rho_0(x), \quad \rho(x, T) = \rho_1(x). \end{aligned} \tag{8}$$

Hence, the OT problem can be formulated as a problem of flowing ρ_0 to ρ_1 with a velocity field v that achieves minimal kinetic energy. The optimal velocity field v has several appealing properties. First, particles induced by the optimal flow v travel in straight lines. Second, particles travel with constant speed. Moreover, under suitable conditions on ρ_0 and ρ_1 , the optimal velocity field is unique⁸.

Given a velocity field v , denote by $z(x, t)$ the solution of the ODE

$$\partial_t z(x, t) = v(z(x, t), t), \quad z(x, 0) = x, \quad 0 \leq t \leq T.$$

Then, under suitable regularity conditions, we have that the solution of the continuity equation is given by $\rho(\cdot, t) = z(\cdot, t)\# \rho_0$. Thus the optimization problem in Eq. (8) can be written as

$$\begin{aligned} \inf_v & \int_0^T \int_{\mathbb{R}^d} \frac{1}{2} \|v(z(x, t), t)\|_2^2 \rho_0(x) dx dt \\ \text{s.t.} & \quad \partial_t z(x, t) = v(z(x, t), t), \quad z(x, 0) = x, \quad z(\cdot, T)\# \rho_0 = \rho_1. \end{aligned} \tag{9}$$

This previous problem is very similar to (4) with the following differences:

- the objective function in Eq. (4) does not have the kinetic energy of trajectories,
- the terminal constraint is imposed as a soft constraint in Eq. (4) and as a hard constraint in Eq. (9), and
- v in Eq. (4) is θ -dependent, whereas the formulation in Eq. (9) is in the non-parametric regime.

So the NF defined by Eq. (4) can be thought of as an approximation to a degenerate transportation problem that lacks transportation cost. Based on this insight one can regularize Eq. (4) by adding the transportation cost and arrive at Eq. (6) or some closely related version of it^{14,15,18,27}. It has been observed that the transportation cost (kinetic energy) regularization significantly improves the training of NFs.

JKO-flow: Wasserstein gradient flows for CNFs

While the OT-based formulation of CNFs in Eq. (6) has been found successful in some applications^{14,15,18,27}, a key difficulty arises in choosing how to balance the kinetic energy term and the KL-divergence, i.e., on selecting α . This difficulty is typical in problems where the constraints are imposed in a soft fashion. Standard training of CNFs typically involves tuning for a “large but hopefully stable enough” step size α so that the KL divergence term is sufficiently small after training. To this end, we propose an approach that avoids the need to tune α by using the fact that the solution to Eq. (6) is an approximation to a backward Euler (or proximal point) algorithm when discretizing the Wasserstein gradient flow using the Jordan–Kinderlehrer–Otto (JKO) scheme²⁰.

The seminal work in²⁰ provides a gradient flow structure of the Fokker–Planck equation using an implicit time discretization. That is, given $\alpha > 0$, density at k th iteration, $\rho^{(k)}$, and terminal density ρ_1 , one finds

$$\begin{aligned}
\rho^{(k+1)} &= \arg \min_{\rho \in \mathcal{P}_2(\mathbb{R}^d)} \frac{1}{2\alpha} W_2^2(\rho, \rho^{(k)}) + KL(\rho || \rho_1) \\
&= \arg \min_v \frac{1}{\alpha} \int_0^1 \int_{\mathbb{R}^d} \frac{1}{2} \|v(z(x, t), t)\|_2^2 \rho_0(x) dx dt + KL(z(\cdot, 1) \# \rho^{(k)} || \rho_1) \\
\text{s.t. } \partial_t z(x, t) &= v(z(x, t), t), z(x, 0) = x
\end{aligned} \tag{10}$$

for $k = 0, 1, \dots$, and $\rho^{(0)} = \rho_0$. Here, α takes the role of a step size when applying a proximal point method to the KL divergence using the Wasserstein-2 metric, and $\{\rho^{(k)}\}$ provably converges to ρ_1 ^{20,21}. Hence, repeatedly solving Eq. (9) with the KL penalty acting as a soft constraint yields an arbitrarily accurate approximation of ρ_1 . In the parametric regime each iteration takes the form

$$\arg \min_{\theta} \mathbb{E}_{x \sim \rho^{(k)}} \int_0^T \frac{1}{2} \|v_{\theta}(x, t)\|_2^2 dt + \alpha C(x, T) \quad \text{subject to } 5. \tag{11}$$

Thus we solve a sequence of problems Eq. (6), where the initial density of the current subproblem is given by the pushforward of the density generated in the previous subproblem.

Importantly, our proposed approach *does not require tuning* α . Instead, we solve a sequence of subproblems that is guaranteed to converge to ρ_1 ²⁰ prior to the neural network parameterization; see Algorithm 1. Indeed, since the traditional approach is equivalent to JKO-Flow with one iteration, JKO-Flow is generally more computationally expensive. But we crucially note that in the traditional single-shot setting, tuning α may require training the model many times as well. JKO-Flow provides a way to automate hyperparameter tuning of α . In our experiments, we observe that ten iterations of JKO-Flow leads to good results for a high-dimensional physics problem, see “Numerical experiments”. While our proposed methodology can be used in tandem with any algorithm used to solve Eq. (11), an important numerical aspect in our approach is to leverage fast computational methods that use *exact* trace estimation in Eq. (5); this approach is called OT-Flow¹⁵. Consequently, we avoid the use of stochastic approximation methods for the trace, e.g., Hutchinson’s estimator^{22,29,30}, as is typically done in CNF methods^{13,14}. A surprising result of our proposed method is that it empirically shows improved performance even with fewer number of parameters (see Fig. 3).

Algorithm 1 Proposed Algorithm

- 1: **Input:** Samples from ρ_0 , step size $\alpha > 0$, number of steps K
 - 2: Initialize θ_1 at random
 - 3: **for** $k = 1, \dots, K$ **do**
 - 4: Solve for θ_k using samples by solving (11)
 - 5: Update distribution of samples using v_{θ_k}
 - 6: **end for**
 - 7: **Output:** saved weights $\theta_1, \dots, \theta_K$
-

Related works

Density estimation. Multivariate density estimation is a fundamental problem in statistics^{31,32}, High Energy Physics (HEP)³³ and in other fields of science dealing with multivariate data. For instance, particle physicists in HEP study possible distributions from a set of high energy data. Another application of density estimation is in confidence level calculations of particles in Higgs searches at Large Electron Positron Colliders (LEP)³⁴ and discriminant methods used in the search for new particles³³. One of the main advantages of NFs over other generative models is that they provide density estimates of probability distributions using Eq. (1). That is, we do not need to apply a separate density estimation technique after generating samples from a distribution, e.g., as in GANs³⁵.

Finite flows. Finite normalizing flows^{4-6,36} use a composition of discrete transformations, where specific architectures are chosen to allow for efficient inverse and Jacobian determinant computations. NICE³⁷, RealNVP³⁸, IAF³⁹, and MAF¹⁰ use either autoregressive or coupling flows where the Jacobian is triangular, so the Jacobian determinant can be tractably computed. GLOW⁴⁰ expands upon RealNVP by introducing an additional invertible convolution step. These flows are based on either coupling layers or autoregressive transformations, whose tractable invertibility allows for density evaluation and generative sampling. Neural Spline Flows⁴¹ use splines instead of the coupling layers used in GLOW and RealNVP. Using monotonic neural networks, NAF⁴² require positivity of the weights, which UMNN⁴³ circumvent this requirement by parameterizing the Jacobian and then integrating numerically. A recent work, called Normalizing Field Flows (NFFs)⁴⁴, generalizes NFs to include learning random fields from scattered measurements; in particular, NFFs can be used to solve data-driven forward, inverse, and mixed forward/inverse stochastic partial differential equations.

Continuous and optimal transport-based flows. Modeling flows with differential equations is a natural and commonly used method^{27,45-50}. In particular, CNFs model their flow via a neural ordinary differential equation^{12,13,51}. Among the most well-known CNFs are FFJORD¹³, which estimates the determinant of the Jacobian by accumulating its trace along the trajectories, and the trace is estimated using Hutchinson’s

estimator^{22,29,30}. To promote straight trajectories, RNODE¹⁴ regularizes FFJORD with a transport cost $L(x, T)$. RNODE also includes the Frobenius norm of the Jacobian $\|\nabla v\|_F^2$ to stabilize training. The trace and the Frobenius norm are estimated using a stochastic estimator and report speedup by a factor of 2.8.

Monge-Ampère Flows¹⁸ and Potential Flow Generators¹⁷ similarly draw from OT theory but parameterize a potential function instead of the dynamics directly. OT is also used in other generative models^{52–57}. OT-Flow¹⁵ is based on a discretize-then-optimize approach⁵⁸ that also parameterizes the potential function. To evaluate the KL divergence, OT-Flow estimates the density using an *exact* trace computation following the work of¹⁹.

Wasserstein gradient flows. Our proposed method is most closely related to⁵⁹, which also employs a JKO-based scheme to perform generative modeling. But a key difference is that⁵⁹ reformulates the KL-divergence as an optimization over difference of expectations (see⁵⁹, Prop. 3.1); this makes their approach akin to GANs, where the density cannot be obtained without using a separate density estimation technique. Our proposed method is also closely related to methods that use input-convex CNNs^{60–62}. Reference⁶² focuses on the special case with KL divergence as objective function.⁶⁰ solve a sequence of subproblems different from the fluid flow formulation presented in Eq. (11). They also require an end-to-end training scheme that backpropagates to the initial distribution; this can become a computational burden when the number of time discretizations is large. Reference⁶¹ utilizes a JKO-based scheme to approximate a population dynamics given an observed trajectory and focus on applications in computational biology. Other related works include natural gradient methods⁶³ and implicit schemes based on the Wasserstein-1 distance⁶⁴.

Numerical experiments

We demonstrate the effectiveness of our proposed JKO-Flow on a series of synthetic and real-world datasets. As previously mentioned, we compute each update in Eq. (10) by solving Eq. (6) using the OT-Flow solver¹⁵, which leverages fast and exact trace computations. We also use the same architecture provided in¹⁵. Henceforth, we shall also call the traditional CNF approach the “single-shot” approach. We also clarify that α in our experiments refers to the parameter in Eq. (6).

Maximum mean discrepancy metric (MMD). Our density estimation problem requires approximating a density ρ_0 by finding a transformation f such that $f^{-1}\#\rho_1$ has density $\hat{\rho}_0$ close to ρ_0 , where ρ_1 is the standard multivariate Gaussian. However, ρ_0 is not known in real-world density estimation scenarios, such as in physics applications, all we have are samples $X = \{x_i\}_{i=1}^n$ from the unknown distribution. Consequently, we use the observed samples X and samples $\hat{X} = \{\hat{x}_j\}_{j=1}^m$, $\hat{x}_j = f^{-1}(q_j)$, generated by the CNF and samples $Q = \{q_j\}_{j=1}^m$ from ρ_1 to determine if their corresponding distributions are close in some sense. To measure the discrepancy we use a particular integral probability metric^{65–67} known as maximum mean discrepancy (MMD) defined as follows⁶⁸: let x and y be random vectors in \mathbb{R}^d with distributions μ_x and μ_y , respectively, and let \mathcal{H} be a reproducing kernel Hilbert space (RKHS) of functions on \mathbb{R}^d with Gaussian kernel (see⁶⁹ for an introduction o RKHS's)

$$k(x_i, x_j) = \exp\left(-\frac{1}{2}\|x_i - x_j\|^2\right). \quad (12)$$

Then the MMD of μ_x and μ_y is given by

$$\text{MMD}_{\mathcal{H}}(\mu_x, \mu_y) = \sup_{\|f\|_{\mathcal{H}} \leq 1} |\mathbb{E} f(x) - \mathbb{E} f(y)|.$$

It can be shown that $\text{MMD}_{\mathcal{H}}$ defines a metric on the class of probability measures on \mathbb{R}^d ^{68,70}. The squared-MMD can be written in terms of the kernel as follows:

$$\text{MMD}_{\mathcal{H}}^2(\mu_x, \mu_y) = \mathbb{E} k(x, x') + \mathbb{E} k(y, y') - 2 \mathbb{E} k(x, y),$$

where x, x' are iid μ_x independent of y, y' which are iid μ_y . An unbiased estimate of the squared-MMD based on the samples X and \hat{X} defined above is given by⁶⁸:

$$\text{MMD}_{\mathcal{H}}^2(X, \hat{X}) = \frac{1}{n(n-1)} \sum_{i \neq j} k(x_i, x_j) + \frac{1}{m(m-1)} \sum_{k \neq \ell} k(\hat{x}_k, \hat{x}_\ell) - \frac{2}{nm} \sum_{i, \ell} k(x_i, \hat{x}_\ell).$$

Note that the MMD is not used for algorithmic training of the CNF, it is only used to compare the densities ρ_0 and $\hat{\rho}_0$ based on the samples X and \hat{X} .

Synthetic 2D data set. We begin by testing our method on seven two-dimensional (2D) benchmark datasets for density estimation algorithms commonly used in machine learning^{13,43}; see Fig. 2. We generate results with JKO-Flow for different values of α and for different number of iterations. We use $\alpha = 1, 5, 10$, and 50, and for each α we use the single shot approach $k = 1$ and JKO-Flow with $k = 5$ iterations from Eq. (10). Note that in CNFs, we are interested in estimating the density (and generating samples) from ρ_0 ; consequently, once we have the optimal weights $\theta^{(1)}, \theta^{(2)}, \dots, \theta^{(5)}$, we must “flow backwards” starting with samples from the normal distribution ρ_1 . Figure 1 shows that JKO-Flow outperforms the single shot approach for different values of α . In particular, the performance for the single shot approach varies drastically for different values of α , with $\alpha = 1$ being an order of magnitude higher in MMD than $\alpha = 5$. On the other hand, JKO-Flow performs consistently regardless of the value of α for most datasets. There is one exception for the spirals dataset with $\alpha = 1$; this is

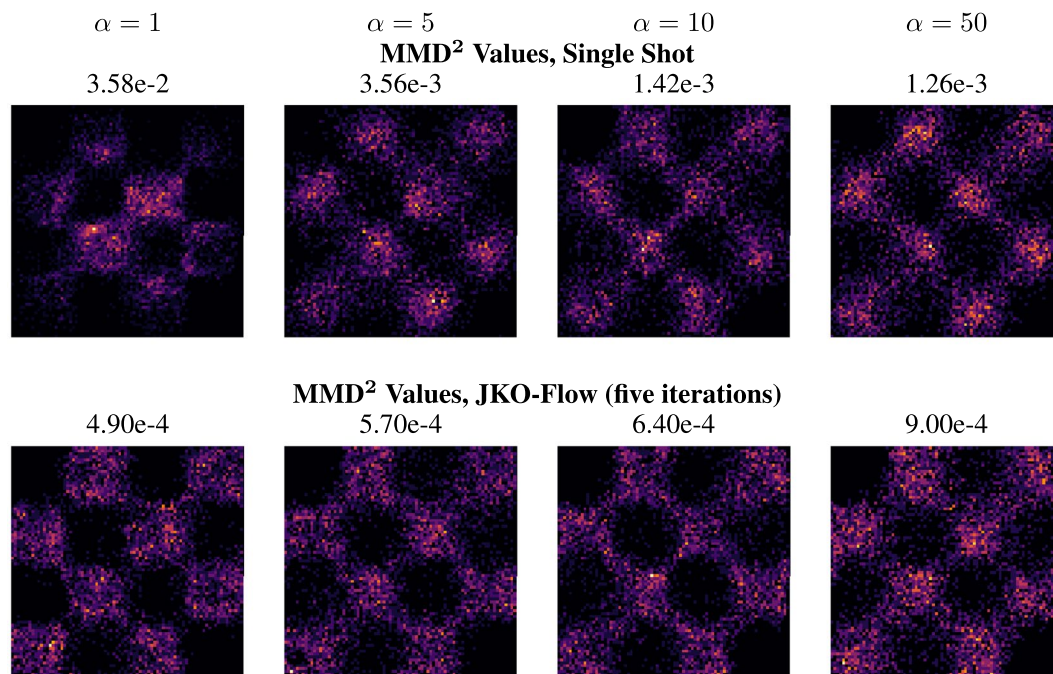


Figure 1. Checkerboard dataset: generated samples of $\hat{\rho}_0$ using the standard one-shot approach (top row). Generated using our proposed JKO-Flow using five iterations (bottom row). Here, we use $\alpha = 1, 5, 10, 50$. JKO-Flow returns consistent results regardless of the value of α .

because only five iterations are used in the JKO scheme and more are needed. When 10 iterations are used instead, we achieve a similar order of accuracy ($2.3e - 4$). As previously mentioned, this is expected as JKO-Flow is a proximal point algorithm that converges regardless of the step size α . In this case, five JKO-Flow iterations are enough to obtain this consistency. Additional plots and hyperparameter setups for different benchmark datasets with similar performance results are shown in the Supplementary Information. Table 1 summarizes the comparison between the single shot and JKO-Flow on all synthetic 2D datasets for different values of α . We also show an illustration of all the datasets, estimated densities, and generated samples with JKO-Flow in Fig. 2.

Varying network size. In addition to obtaining consistent results for different values of α , we also *empirically* observe that JKO-Flow outperforms the single shot approach for different numbers of network parameters, i.e., network size. We illustrate this in Fig. 3. This is also intuitive as we reformulate the problem of finding a

α		1	5	10	50
Dataset	Approach	MMD ²			
Checkerboard	Single shot	3.58e-2	3.56e-3	1.42e-3	1.26e-3
	JKO-flow (5 iters)	4.9e-4	5.67e-4	6.40e-4	9.00e-4
2 spirals	Single shot	7.21e-2	2.30e-2	1.84e-2	7.73e-4
	JKO-flow (5 iters)	2.10e-2	4.62e-4	1.02e-4	5.37e-5
Swiss roll	Single shot	4.74e-3	7.33e-4	2.86e-4	7.03e-4
	JKO-flow (5 iters)	5.16e-4	8.3e-5	3.27e-5	6.07e-4
8 Gaussians	Single shot	9.18e-3	2.69e-4	3.94e-4	7.10e-4
	JKO-flow (5 iters)	1.07e-4	4.13e-5	2.67e-4	7.27e-6
Circles	Single shot	9.84e-3	2.24e-4	6.51e-4	1.04e-4
	JKO-flow (5 iters)	9.49e-4	9.97e-6	2.38e-5	9.28e-5
Pinwheel	Single shot	1.18e-2	1.8e-3	1.37e-3	2.2e-5
	JKO-flow (5 iters)	4.7e-4	2.63e-4	3.84e-4	4.30e-4
Moons	Single shot	1.45e-3	2.05e-3	2.49e-4	2.42e-4
	JKO-flow (5 iters)	1.92e-4	4.65e-5	4.3e-5	1.08e-4

Table 1. Synthetic 2D data: JKO-flow performance for different values of α . JKO-flow returns consistent performance for different α .




Figure 2. Density estimation on 2D toy problems using five JKO-Flow iterations. Top: samples from the unknown distribution ρ_0 . Middle: density estimate for ρ_0 computed by inverting the flow through the five iterations of JKO-Flow from ρ_1 via Eq. (2). Bottom: samples generated by inverse JKO-Flow through five iterations where y has density ρ_1 .

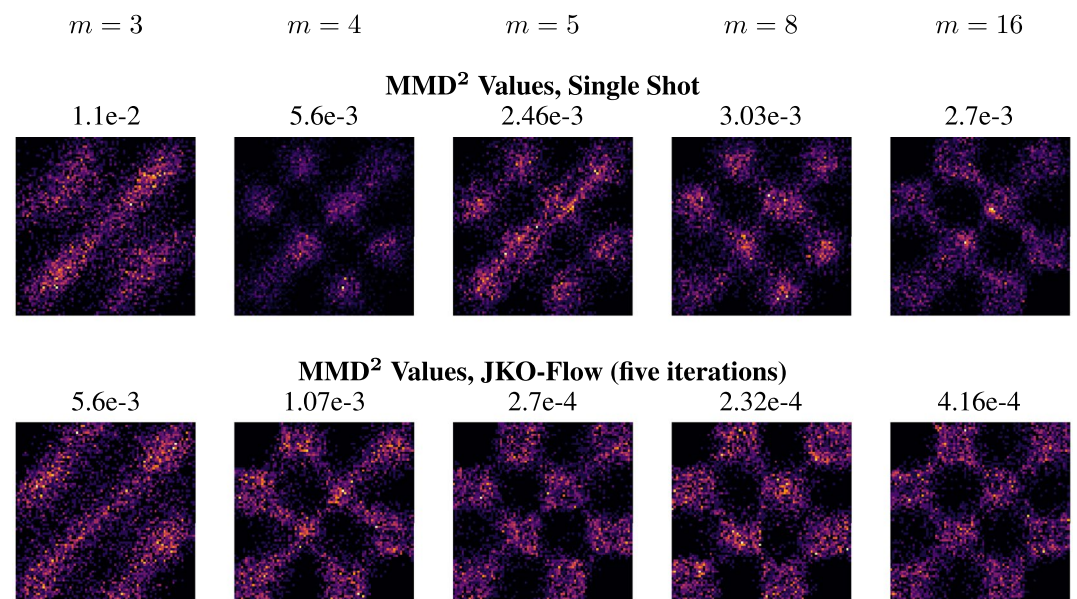


Figure 3. Checkerboard dataset: generated samples of $\hat{\rho}_0$ using the standard single shot approach (top row). Generated samples using our proposed JKO-Flow using five iterations (bottom row). Here, we fix $\alpha = 5$ and vary the network width $m = 3, 4, 5, 8$, and 16 . JKO-Flow performs competitively even with fewer parameters.

single “difficult” optimal transportation problem as a sequence of “smaller and easier” OT problems. In this setup, we vary the width of a two-layer ResNet⁷¹. In particular, we choose the widths to be $m = 3, 4, 5, 8$, and 16 . These correspond to 40, 53, 68, 125, and 365 parameters. The hyperparameter α is chosen to be the best performing value for each synthetic dataset. All datasets vary m for fixed $\alpha = 5$, except the 2 Spiral dataset, which uses $\alpha = 50$; we chose these α values as they performed the best in the fixed m experiments. Similar results are also shown for the remaining synthetic datasets in the Supplementary Information. Table 2 summarizes the comparison between the single shot and JKO-Flow on all synthetic 2D datasets.

Density estimation on a physics dataset. We train JKO-Flow on the 43-dimensional MINIBOONE dataset which is a high-dimensional, real-world physics dataset used as benchmark for high-dimensional density estimation algorithms in physics⁷². For this physics problem, our method is trained for $\alpha = 0.5, 1, 5, 10, 50$ and using 10 JKO-Flow iterations. Fig. 4 shows generated samples with JKO-Flow and the standard single-shot approach for $\alpha = 5$. Since MINIBOONE is a high-dimensional dataset, we follow¹⁵ and plot two-dimensional

m		3	4	5	8	16
Dataset	Approach	MMD ²				
Checkerboard	Single shot	1.10e-2	5.60e-3	2.46e-3	3.03e-3	2.70e-3
	JKO-flow (5 iters)	5.60e-3	1.07e-3	2.7e-4	2.32e-4	4.16e-4
2 spirals	Single shot	5.98e-3	4.54e-3	5.47e-3	1.19e-3	3.96e-3
	JKO-flow (5 iters)	1.42e-3	1.49e-5	6.11e-4	3.93e-5	2.19e-3
Swiss roll	Single shot	8.89e-3	7.71e-3	1.41e-3	1.37e-3	1.52e-3
	JKO-flow (5 iters)	1.49e-3	2.90e-4	6.13e-4	2.29e-4	8.40e-5
8 Gaussians	Single shot	2.20e-3	1.05e-3	1.04e-3	2.3e-4	5.05e-4
	JKO-flow (5 iters)	1.33e-4	9.85e-4	2.40e-5	3.96e-4	1.07e-4
Circles	Single shot	2.06e-3	1.72e-3	1.37e-3	1.69e-3	1.34e-3
	JKO-flow (5 iters)	1.94e-3	3.24e-4	7.71e-4	5.9e-5	1.01e-4
Pinwheel	Single shot	1.10e-2	4.03e-3	2.27e-3	3.80e-3	5.43e-4
	JKO-flow (5 iters)	1.20e-3	8.23e-4	1.60e-3	7.00e-5	2.69e-4
Moons	Single shot	4.98e-3	4.54e-3	5.47e-3	1.2e-3	3.96e-3
	JKO-flow (5 iters)	1.42e-3	1.5e-5	6.11e-4	3.90e-5	2.19e-3

Table 2. Synthetic 2D data: network width comparison for 1 and 5 iterations given a fixed, best performing α . JKO-Flow performs better than the single shot approach for different network sizes.

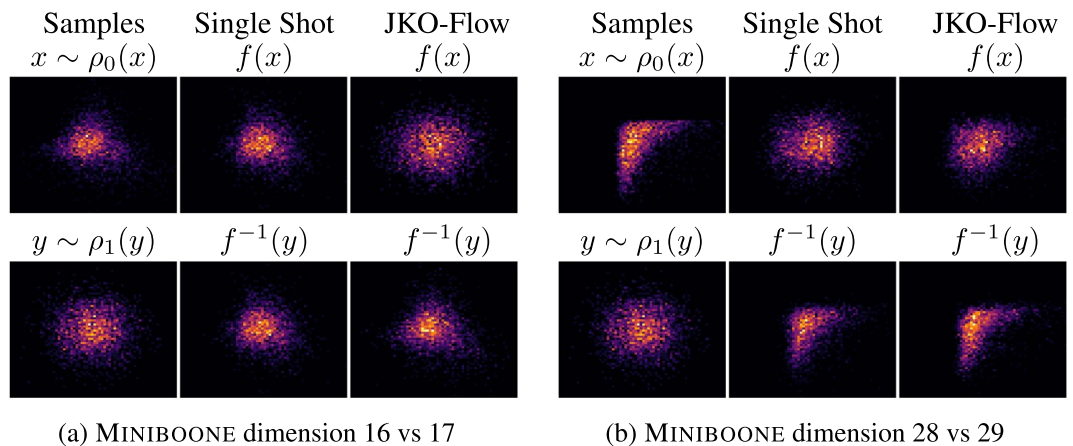


Figure 4. Generated samples for the 43-dimensional MINIBOONE dataset using the single shot approach and JKO-Flow with 10 iterations. To visualize the dataset, we show 2-dimensional slices. We show the forward flow $f(x)$ where $x \sim \rho_0$ and the generated samples $f^{-1}(y)$ where $y \sim \rho_1$.

slices. JKO-Flow generates better quality samples. Similar experiments for $\alpha = 1, 10,$ and 50 are shown in the Supplementary Information. Table 3 summarizes the results for all values of α . Note that we compute MMD values for all the dimensions as well as 2D slices; this is because we only have limited data (3000 testing samples) and the 2D slice MMD give a better indication on the improvement of the generated samples. Results show that *the MMD is consistent across all α values for JKO-Flow*. We also show the convergence (in MMD²) of the miniboone dataset across each 2D slice in Fig. 5. As expected, smaller step size α values converge slower (see $\alpha = 0.5$), but all converge to similar accuracy (unlike the single-shot).

Conclusion

We propose a new approach we call JKO-Flow to train OT-regularized CNFs without having to tune the regularization parameter α . The key idea is to embed an underlying OT-based CNF solver into a Wasserstein gradient flow framework, also known as the JKO scheme; this approach makes the regularization parameter act as a “time” variable. Thus, instead of tuning α , we repeatedly solve proximal updates for a fixed (time variable) α . In our setting, we choose OT-Flow¹⁵, which leverages exact trace estimation for fast CNF training. Our numerical experiments show that JKO-Flow leads to improved performance over the traditional approach. Moreover, *JKO-Flow achieves similar results regardless of the choice of α* . We also empirically observe improved performance when varying the size of the neural network. Future work will investigate JKO-Flow on similar problems such as deep learning-based methods for optimal control^{73–75} and mean field games^{19,76,77}.

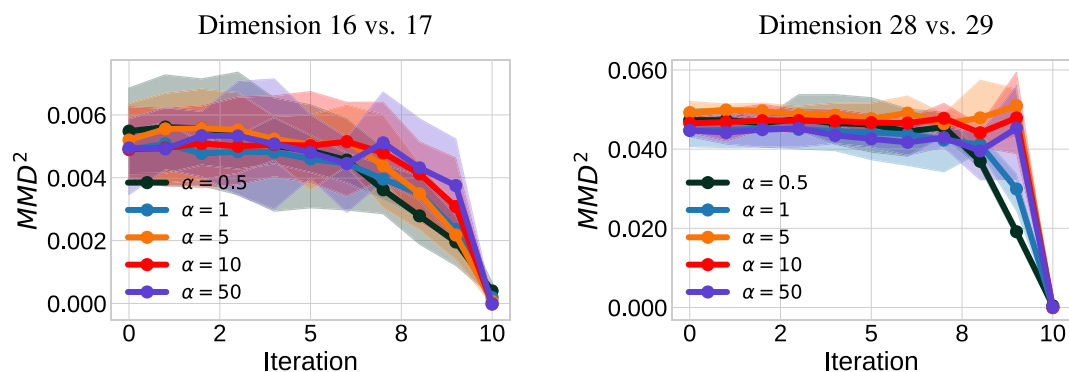


Figure 5. MMD² per iteration when using JKO-Flow after training. MMD² values vs. JKO-Flow iteration for 2-dimensional slice (dimensions 16–17 on the left and dimensions 28–29 on the right) of the MINIBOONE dataset. JKO-Flow achieves same accuracy *regardless* of the value of α .

α		0.5	1	5	10	50
Dimensions	Approach	MMD ²				
2d: 16 vs. 17	Single shot	4.62e-3	4.50e-3	5.88e-3	4.89e-3	5.2e-3
	JKO-flow (10 iters)	3.42e-4	2.94e-4	1.27e-4	1.43e-4	1.71e-4
2d: 28 vs. 29	Single shot	4.33e-2	4.60e-2	5.02e-2	4.97e-2	4.74e-2
	JKO-flow (10 iters)	8.02e-5	3.31e-5	4.43e-5	6.33e-5	9.97e-5
Full	Single shot	4.7e-2	4.51e-3	4.75e-3	4.21e-3	4.27e-3
	JKO-flow (10 iters)	4.72e-4	4.72e-4	4.71e-4	4.72e-4	4.72e-4

Table 3. MINIBOONE: comparison of single shot and JKO-flow for different values of α

Data availability

The datasets used and/or analysed during the current study available from the corresponding author on reasonable request.

Received: 13 December 2022; Accepted: 13 March 2023

Published online: 18 March 2023

References

- Brehmer, J., Kling, F., Espejo, I. & Cranmer, K. Madminer: Machine learning-based inference for particle physics. *Comput. Softw. Big Sci.* **4**(1), 1–25 (2020).
- Carleo, G. *et al.* Machine learning and the physical sciences. *Rev. Modern Phys.* **91**(4), 045002 (2019).
- F. Noé, S. Olsson, J. Köhler, & H. Wu. Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning. *Science*. **365**(6457), eaaw1147 (2019).
- I. Kobyzev, S. Prince, & M. Brubaker. Normalizing flows: An introduction and review of current methods. in *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2020).
- G. Papamakarios, E. Nalisnick, D. J. Rezende, S. Mohamed, & B. Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. [arXiv:1912.02762](https://arxiv.org/abs/1912.02762) (2019).
- D. J. Rezende & S. Mohamed. Variational inference with normalizing flows. in *International Conference on Machine Learning (ICML)*, 1530–1538 (2015).
- G. Peyré & M. Cuturi. Computational optimal transport. (2018).
- Villani, C. in *Topics in Optimal Transportation* Vol. 58 (American Mathematical Society, Providence, RI, 2003).
- R. Baptista, Y. Marzouk, R. E. Morrison, & O. Zahm. Learning non-Gaussian graphical models via Hessian scores and triangular transport. [arXiv preprint arXiv:2101.03093](https://arxiv.org/abs/2101.03093) (2021).
- G. Papamakarios, T. Pavlakou, & I. Murray. Masked autoregressive flow for density estimation. in *Advances in Neural Information Processing Systems (NeurIPS)*, 2338–2347 (2017).
- Zech, J. & Marzouk, Y. Sparse approximation of triangular transports, part II: The infinite-dimensional case. *Construct. Approximation* **55**(3), 987–1036 (2022).
- C. Chen, C. Li, L. Chen, W. Wang, Y. Pu, & L. C. Duke. Continuous-time flows for efficient inference and density estimation. in *International Conference on Machine Learning (ICML)*, 824–833 (2018).
- W. Grathwohl, R. T. Chen, J. Betterncourt, I. Sutskever, & D. Duvenaud. FFJORD: Free-form continuous dynamics for scalable reversible generative models. in *International Conference on Learning Representations (ICLR)* (2019).
- C. Finlay, J.-H. Jacobsen, L. Nurbekyan, & A. M. Oberman. How to train your neural ODE: The world of Jacobian and kinetic regularization. in *International Conference on Machine Learning (ICML)*, 3154–3164 (2020).
- D. Onken, S. Wu Fung, X. Li, & L. Ruthotto. Ot-flow: Fast and accurate continuous normalizing flows via optimal transport. in *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35 (2021).
- Villani, C. *Optimal Transport: Old and New* Vol. 338 (Springer Science & Business Media, New York, 2008).
- L. Yang and G. E. Karniadakis. Potential flow generator with L_2 optimal transport regularity for generative models. in *IEEE Transactions on Neural Networks and Learning Systems* (2020).

18. L. Zhang, W. E. & L. Wang. Monge–Ampère flow for generative modeling. [arXiv:1809.10188](https://arxiv.org/abs/1809.10188) (2018).
19. Ruthotto, L., Osher, S. J., Li, W., Nurbekyan, L. & Fung, S. W. A machine learning framework for solving high-dimensional mean field game and mean field control problems. *Proc. Natl. Acad. Sci.* **117**(17), 9183–9193 (2020).
20. Jordan, R., Kinderlehrer, D. & Otto, F. The variational formulation of the Fokker–Planck equation. *SIAM J. Math. Anal.* **29**(1), 1–17 (1998).
21. A. Salim, A. Korba, & G. Luise. The Wasserstein proximal gradient algorithm. in *Advances in Neural Information Processing Systems* (H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, & H. Lin, eds.), Vol. 33, 12356–12366. (Curran Associates, Inc., 2020).
22. L. Tenorio. *An Introduction to Data Analysis and Uncertainty Quantification for Inverse Problems*. (SIAM, 2017).
23. M. Burger, S. Osher, J. Xu, & G. Gilboa. Nonlinear inverse scale space methods for image restoration. in *Variational, Geometric, and Level Set Methods in Computer Vision* (N. Paragios, O. Faugeras, T. Chan, & C. Schnörr, eds.), 25–36 (Springer, 2005).
24. Osher, S., Burger, M., Goldfarb, D., Xu, J. & Yin, W. An iterative regularization method for total variation-based image restoration. *Multisc. Model. Simulat.* **4**(2), 460–489 (2005).
25. Peyré, G. & Cuturi, M. Computational optimal transport. *Foundations Trends Mach. Learn.* **11**(5–6), 355–607 (2019).
26. F. Santambrogio. *Optimal Transport for Applied Mathematicians*, Vol. 87 of *Progress in Nonlinear Differential Equations and their Applications*. Birkhäuser/Springer, Cham, 2015. Calculus of variations, PDEs, and modeling.
27. M. Welling & Y. W. Teh. Bayesian learning via stochastic gradient Langevin dynamics. in *International Conference on Machine Learning (ICML)*, 681–688 (2011).
28. Benamou, J.-D. & Brenier, Y. A computational fluid mechanics solution to the Monge–Kantorovich mass transfer problem. *Numerische Mathematik* **84**(3), 375–393 (2000).
29. Avron, H. & Toledo, S. Randomized algorithms for estimating the trace of an implicit symmetric positive semi-definite matrix. *J. ACM (JACM)* **58**(2), 1–34 (2011).
30. Hutchinson, M. F. A stochastic estimator of the trace of the influence matrix for Laplacian smoothing splines. *Commun. Stat.-Simulat. Comput.* **19**(2), 433–450 (1990).
31. Scott, D. W. *Multivariate Density Estimation: Theory, Practice, and Visualization* (Wiley, New York, 2015).
32. Silverman, B. W. *Density Estimation for Statistics and Data Analysis* (CRC Press, New York, 1986).
33. Cranmer, K. Kernel estimation in high-energy physics. *Comput. Phys. Commun.* **136**(3), 198–207 (2001).
34. Collaboration, O. & Abbiendi, G. Search for neutral Higgs bosons in collisions at 189 gev. *Eur. Phys. J. C-Particles Fields.* **12**(4), 567–586 (2000).
35. Goodfellow, I. et al. Generative adversarial networks. *Commun. ACM* **63**(11), 139–144 (2020).
36. Tabak, E. G. & Turner, C. V. A family of nonparametric density estimation algorithms. *Commun. Pure Appl. Math.* **66**(2), 145–164 (2013).
37. L. Dinh, D. Krueger, & Y. Bengio. NICE: Non-linear independent components estimation. in *International Conference on Learning Representations (ICLR)* (Y. Bengio & Y. LeCun, eds.) (2015).
38. L. Dinh, J. Sohl-Dickstein, & S. Bengio. Density estimation using real NVP. in *International Conference on Learning Representations (ICLR)* (2017).
39. D. P. Kingma, T. Salimans, R. Jozefowicz, X. Chen, I. Sutskever, & M. Welling. Improved variational inference with inverse autoregressive flow. in *Advances in Neural Information Processing Systems (NeurIPS)*, 4743–4751 (2016).
40. D. P. Kingma & P. Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. in *Advances in Neural Information Processing Systems (NeurIPS)*, 10215–10224 (2018).
41. C. Durkan, A. Bekasov, I. Murray, & G. Papamakarios. Neural spline flows. in *Advances in Neural Information Processing Systems (NeurIPS)*, 7509–7520 (2019).
42. C.-W. Huang, D. Krueger, A. Lacoste, & A. Courville. Neural autoregressive flows. in *International Conference on Machine Learning (ICML)* 2078–2087 (2018).
43. A. Wehenkel & G. Louppe. Unconstrained monotonic neural networks. in *Advances in Neural Information Processing Systems (NeurIPS)*, 1543–1553 (2019).
44. Guo, L., Wu, H. & Zhou, T. Normalizing field flows: Solving forward and inverse stochastic differential equations using physics-informed flow models. *J. Comput. Phys.* **461**, 111202 (2022).
45. C.-W. Huang, R. T. Chen, C. Tsirigotis, & A. Courville. Convex potential flows: Universal probability distributions with optimal transport and convex optimization. [arXiv preprint arXiv:2012.05942](https://arxiv.org/abs/2012.05942) (2020).
46. Neal, R. M. MCMC using Hamiltonian dynamics. *Handbook Markov Chain Monte Carlo* **2**(11), 2 (2011).
47. Y. Park, D. Maddix, F.-X. Aubet, K. Kan, J. Gasthaus, & Y. Wang. Learning quantile functions without quantile crossing for distribution-free time series forecasting. in *International Conference on Artificial Intelligence and Statistics*, 8127–8150. (PMLR, 2022).
48. Ruthotto, L. & Haber, E. An introduction to deep generative modeling. *GAMM-Mitteilungen* **44**(2), e202100008 (2021).
49. T. Salimans, D. Kingma, & M. Welling. Markov chain Monte Carlo and variational inference: Bridging the gap. in *International Conference on Machine Learning (ICML)*, 1218–1226 (2015).
50. Suykens, J., Verrelst, H. & Vandewalle, J. On-line learning Fokker–Planck machine. *Neural Process. Lett.* **7**, 81–89 (1998).
51. T. Q. Chen, Y. Rubanova, J. Bettencourt, & D. K. Duvenaud. Neural ordinary differential equations. in *Advances in Neural Information Processing Systems (NeurIPS)*, 6571–6583, (2018).
52. G. Avraham, Y. Zuo, & T. Drummond. Parallel optimal transport GAN. in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 4406–4415 (2019).
53. Lei, N., Su, K., Cui, L., Yau, S.-T. & Gu, X. D. A geometric view of optimal transportation and generative model. *Comput. Aided Geometr. Design* **68**, 1–21 (2019).
54. J. Lin, K. Lensink, & E. Haber. Fluid flow mass transport for generative networks. [arXiv:1910.01694](https://arxiv.org/abs/1910.01694) (2019).
55. T. Salimans, H. Zhang, A. Radford, & D. N. Metaxas. Improving GANs using optimal transport. in *International Conference on Learning Representations (ICLR)* (2018).
56. M. Sanjabi, J. Ba, M. Razaviyayn, & J. D. Lee. On the convergence and robustness of training gans with regularized optimal transport. in *Advances in Neural Information Processing Systems (NeurIPS)*, 7091–7101 (2018).
57. A. Tanaka. Discriminator optimal transport. in *Advances in Neural Information Processing Systems (NeurIPS)*, 6816–6826 (2019).
58. D. Onken & L. Ruthotto. Discretize-optimize vs. optimize-discretize for time-series regression and continuous normalizing flows. [arXiv:2005.13420](https://arxiv.org/abs/2005.13420) (2020).
59. J. Fan, Q. Zhang, A. Taghvaei, & Y. Chen. Variational Wasserstein gradient flow. in *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research* (K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvari, G. Niu, & S. Sabato, eds.), 6185–6215. PMLR, 17–23 Jul 2022.
60. D. Alvarez-Melis, Y. Schiff, & Y. Mroueh. Optimizing functionals on the space of probabilities with input convex neural networks. [arXiv preprint arXiv:2106.00774](https://arxiv.org/abs/2106.00774) (2021).
61. C. Bunne, L. Papaxanthos, A. Krause, & M. Cuturi. Proximal optimal transport modeling of population dynamics. in *International Conference on Artificial Intelligence and Statistics*, 6511–6528. (PMLR, 2022).
62. Mokrov, P. et al. Large-scale Wasserstein gradient flows. *Adv. Neural Inform. Process. Syst.* **34**, 15243–15256 (2021).
63. L. Nurbekyan, W. Lei, & Y. Yang. Efficient natural gradient descent methods for large-scale optimization problems. [arXiv preprint arXiv:2202.06236](https://arxiv.org/abs/2202.06236) (2022).

64. H. Heaton, S. W. Fung, A. T. Lin, S. Osher, & W. Yin. Wasserstein-based projections with applications to inverse problems. arXiv preprint [arXiv:2008.02200](https://arxiv.org/abs/2008.02200) (2020).
65. Müller, A. Integral probability metrics and their generating classes of functions. *Adv. Appl. Probability* **29**(2), 429–443 (1997).
66. Rachev, S. T., Klebanov, L. B., Stoyanov, S. V. & Fabozzi, F. *The Methods of Distances in the Theory of Probability and Statistics* Vol. 10 (Springer, New York, 2013).
67. Zolotarev, V. M. Metric distances in spaces of random variables and their distributions. *Math. USSR-Sbornik* **30**(3), 373 (1976).
68. Gretton, A., Borgwardt, K. M., Rasch, M. J., Schölkopf, B. & Smola, A. A kernel two-sample test. *J. Mach. Learn. Res. (JMLR)* **13**(25), 723–773 (2012).
69. Paulsen, V. I. & Raghupathi, M. *An Introduction to the Theory of Reproducing Kernel Hilbert Spaces* Vol. 152 (Cambridge University Press, Cambridge, 2016).
70. K. Fukumizu, A. Gretton, X. Sun, & B. Schölkopf. Kernel measures of conditional dependence. *Adv. Neural Inform. Process. Syst.* **20** (2007).
71. K. He, X. Zhang, S. Ren, & J. Sun. Deep residual learning for image recognition. in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 770–778 (2016).
72. B. Roe. MiniBooNE particle identification. UCI Machine Learning Repository (2010).
73. W. H. Fleming & H. M. Soner. *Controlled Markov Processes and Viscosity Solutions*, Vol. 25 of *Stochastic Modelling and Applied Probability*, 2nd Edn. (Springer, 2006).
74. D. Onken, L. Nurbekyan, X. Li, S. W. Fung, S. Osher, & L. Ruthotto. A neural network approach applied to multi-agent optimal control. in *2021 European Control Conference (ECC)*, 1036–1041. (IEEE, 2021).
75. D. Onken, L. Nurbekyan, X. Li, S. W. Fung, S. Osher, & L. Ruthotto. A neural network approach for high-dimensional optimal control applied to multiagent path finding. in *IEEE Transactions on Control Systems Technology* (2022).
76. Agrawal, S., Lee, W., Fung, S. W. & Nurbekyan, L. Random features for high-dimensional nonlocal mean-field games. *J. Comput. Phys.* **459**, 111136 (2022).
77. Lin, A. T., Fung, S. W., Li, W., Nurbekyan, L. & Osher, S. J. Alternating the population and control neural networks to solve high-dimensional stochastic mean-field games. *Proc. Natl. Acad. Sci.* **118**(31), e2024713118 (2021).

Acknowledgements

LN and SO were partially funded by AFOSR MURI FA9550-18-502, ONR N00014-18-1-2527, N00014-18-20-1-2093 and N00014-20-1-2787.

Author contributions

A.V., L.N., and S.W.F. performed the research. All Authors wrote the manuscript.

Competing Interests

The authors declare no competing interests.

Additional information

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1038/s41598-023-31521-y>.

Correspondence and requests for materials should be addressed to A.V.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2023