



MOTIVATION

Continuous Normalizing Flows (CNFs)

A normalizing flow [1] is an invertible mapping $f: \mathbb{R}^d \rightarrow \mathbb{R}^d$ between an arbitrary probability distribution and a standard normal distribution whose densities we denote by ρ_0 and ρ_1 , respectively.

By change of variables,

$$\log \rho_0(\mathbf{x}) = \log \rho_1(f(\mathbf{x})) + \log |\det \nabla f(\mathbf{x})| \quad \text{for all } \mathbf{x} \in \mathbb{R}^d. \quad (1)$$

In CNFs, f solves the neural ordinary differential equation (ODE) [2, 3]

$$\partial_t \begin{bmatrix} \mathbf{z}(\mathbf{x}, t) \\ \ell(\mathbf{x}, t) \end{bmatrix} = \begin{bmatrix} \mathbf{v}(\mathbf{z}(\mathbf{x}, t), t; \boldsymbol{\theta}) \\ \text{tr}(\nabla \mathbf{v}(\mathbf{z}(\mathbf{x}, t), t; \boldsymbol{\theta})) \end{bmatrix}, \quad \begin{bmatrix} \mathbf{z}(\mathbf{x}, 0) \\ \ell(\mathbf{x}, 0) \end{bmatrix} = \begin{bmatrix} \mathbf{x} \\ 0 \end{bmatrix}, \quad (2)$$

where, for artificial time $t \in [0, T]$,

- \mathbf{x} maps to $f(\mathbf{x}) = \mathbf{z}(\mathbf{x}, T)$ following trajectory $\mathbf{z}: \mathbb{R}^d \times [0, T] \rightarrow \mathbb{R}^d$
- $\mathbf{v}: \mathbb{R}^d \times [0, T] \rightarrow \mathbb{R}^d$ is a neural network layer parameterized by $\boldsymbol{\theta}$
- $\ell(\mathbf{x}, T) = \log \det \nabla f(\mathbf{x})$, derived from Jacobi's Formula [2]

Microscale: an arbitrary sample \mathbf{x} maps to a normally distributed $f(\mathbf{x})$

Macroscale: ρ_0 maps to ρ_1

CNFs are trained by solving the optimization problem [1, 3]

$$\min_{\boldsymbol{\theta}} \mathbb{E}_{\rho_0(\mathbf{x})} \left\{ C(\mathbf{x}, T) := \frac{1}{2} \|\mathbf{z}(\mathbf{x}, T)\|^2 - \ell(\mathbf{x}, T) + \frac{d}{2} \log(2\pi) \right\} \quad \text{s.t.} \quad (2). \quad (3)$$

High Training Costs

- Many functions evaluations are needed to solve (2)
- Computing the trace with automatic differentiation (AD) requires vector-Jacobian products with all d standard basis vectors, costing $\mathcal{O}(d^2)$ FLOPs total

Our Contributions

Optimal Transport (OT) Incorporating OT, we regularize the CNF so it has a unique solution (Figure 1).

Analytic Exact Trace We derive formulae for an exact trace computation with complexity $\mathcal{O}(d)$ FLOPs.




Figure 1: While a CNF can have curved trajectories, OT-Flow's are straight (modification of Fig. 1 in [3, 4]).

OPTIMAL TRANSPORT

L_2 Transport Costs Add transport costs

$$L(\mathbf{x}, T) = \int_0^T \frac{1}{2} \|\mathbf{v}(\mathbf{z}(\mathbf{x}, t), t)\|^2 dt.$$

to (3) to penalize the arc-length of the trajectories.

Potential Function By the Pontryagin maximum principle [5], there exists a scalar potential function $\Phi: \mathbb{R}^d \times [0, T] \rightarrow \mathbb{R}$ such that

$$\mathbf{v}(\mathbf{x}, t; \boldsymbol{\theta}) = -\nabla \Phi(\mathbf{x}, t; \boldsymbol{\theta}).$$

Idea: Analogous to classical physics, samples move to minimize their potential

\Rightarrow We parameterize potential Φ instead of \mathbf{v} .

HJB Regularizer At optimality, Φ satisfies the Hamilton-Jacobi-Bellman (HJB) equation [6]

$$-\partial_t \Phi(\mathbf{x}, t) = -\frac{1}{2} \|\nabla \Phi(\mathbf{z}(\mathbf{x}, t), t)\|^2,$$

$$\Phi(\mathbf{x}, T) = 1 + \log(\rho_0(\mathbf{x})) - \log(\rho_1(\mathbf{z}(\mathbf{x}, T))) - \ell(\mathbf{x}, T)$$

To penalize sub-optimality, use HJB regularizer

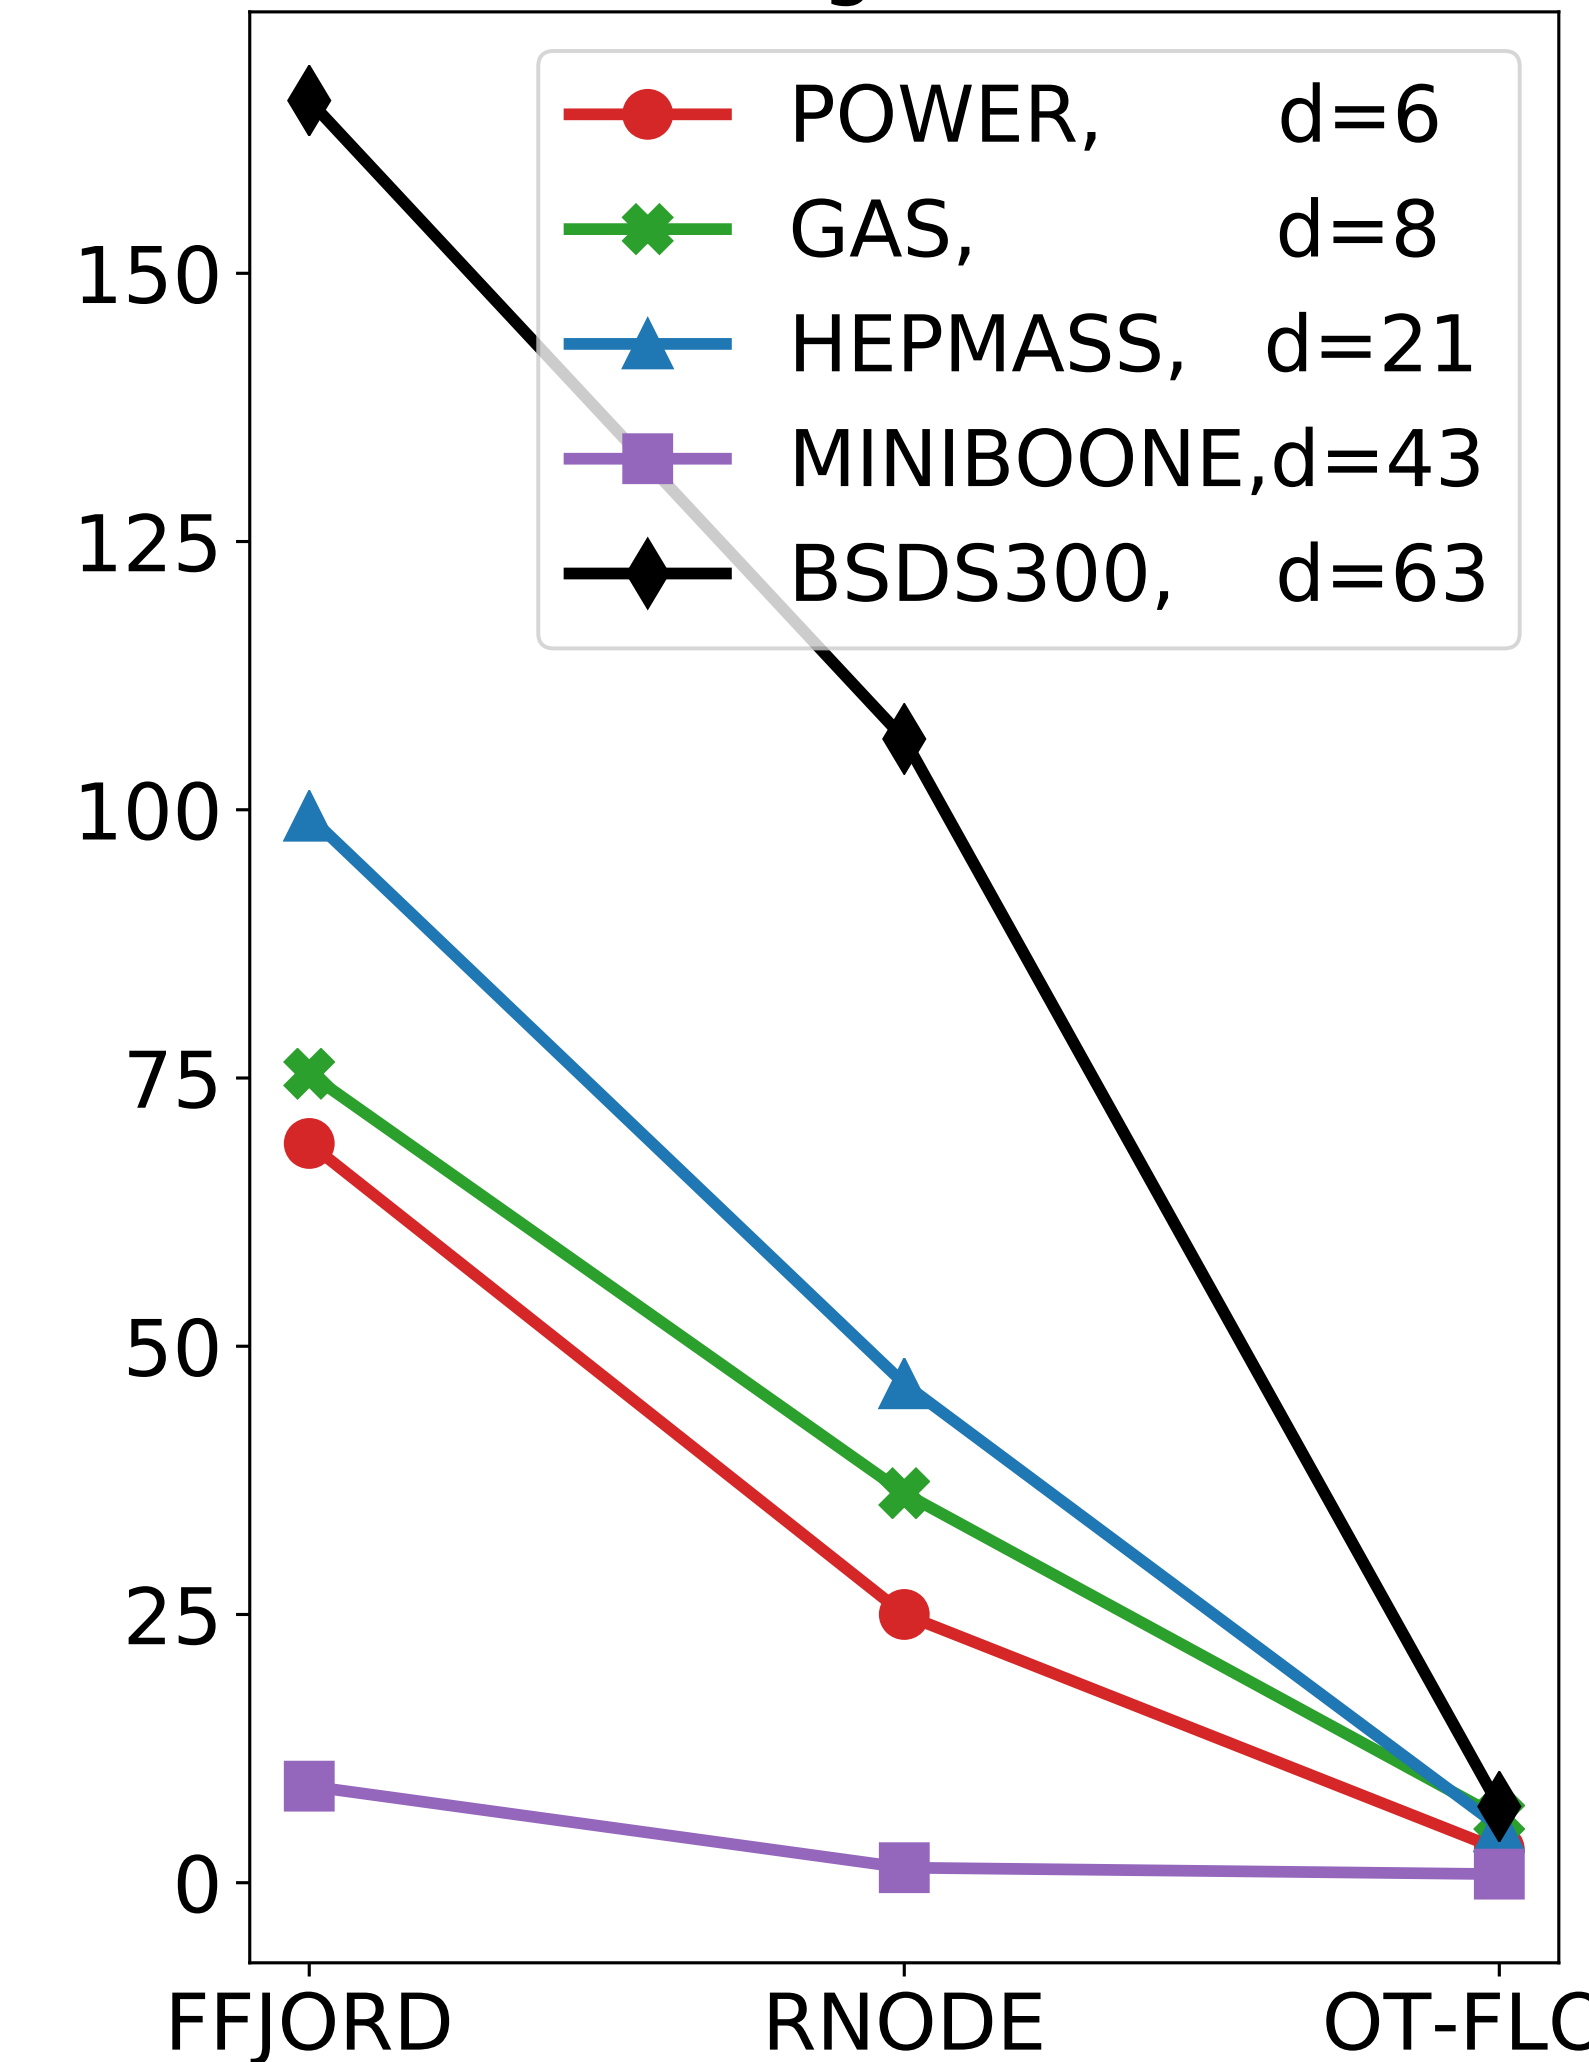
$$R(\mathbf{x}, T) = \int_0^T \left| \partial_t \Phi(\mathbf{z}(\mathbf{x}, t), t) - \frac{1}{2} \|\nabla \Phi(\mathbf{z}(\mathbf{x}, t), t)\|^2 \right| dt$$

OT-Flow Optimization Problem

$$\min_{\boldsymbol{\theta}} \mathbb{E}_{\rho_0(\mathbf{x})} \left\{ C(\mathbf{x}, T) + L(\mathbf{x}, T) + R(\mathbf{x}, T) \right\} \quad \text{s.t.} \quad (2)$$

RESULTS

Training Time (hr)



8x speedup in training and 24x speedup in inference relative to the state-of-the-art RNODE [4] on five real datasets of dimensionality $d = 6, 8, 21, 43, 63$

REFERENCES

- [1] D Rezende and S Mohamed. *Variational Inference with Normalizing Flows*. ICML, 2015.
- [2] Chen et al. *Neural Ordinary Differential Equations*. NeurIPS, 2018.
- [3] W Grathwohl et al. *FFJORD: Free-form Continuous Dynamics for Scalable Reversible Generative Models*. ICLR, 2019.
- [4] C Finlay et al. *How to Train your Neural ODE: the World of Jacobian and Kinetic Regularization*. ICML, 2020.
- [5] L Evans. *An Introduction to Mathematical Optimal Control Theory Version 0.2*. 2013.
- [6] L Evans. *Partial differential equations and Monge-Kantorovich mass transfer*. Current Developments in Mathematics, 1997.

ACKNOWLEDGMENTS

National Science Foundation award DMS 1751636, Binational Science Foundation Grant 2018209, AFOSR Grants 20RT0237 and FA9550-18-1-0167, AFOSR MURI FA9550-18-1-0502, ONR Grant No. N00014-18-1-2527, a gift from UnitedHealth Group R&D, and a GPU donation by NVIDIA Corporation.

Special Thanks: IPAM Long Program MLP 2019 organizers & staff

CODE

github.com/EmoryMLIP/OT-Flow

EXACT TRACE

We devise an efficient analytic method for computing the trace by exploiting

$$\text{tr}(\nabla^2 \Phi(\mathbf{s}; \boldsymbol{\theta})) = \text{tr}(\mathbf{E}^\top \nabla_{\mathbf{s}}^2 \Phi(\mathbf{s}; \boldsymbol{\theta}) \mathbf{E}) \quad \text{for } \mathbf{E} = \text{eye}(d+1, d).$$

In runtime, the exact trace is competitive with estimators used in other CNFs.

In convergence, the exact trace during training achieves 1) quicker validation convergence than using an estimator and 2) less variance in training loss than using an estimator

