

Generative Modeling with Optimal Transport Maps

Litu Rout¹, Alexander Korotin²³, Evgeny Burnaev²³

¹Space Applications Centre (SAC), ISRO

²Skolkovo Institute of Science and Technology

³Artificial Intelligence Research Institute (AIRI)

International Conference on Learning Representations (ICLR 2022)

Overview

Introduction

Optimal Transport in Generative Models

Proposed Optimal Transport Modeling

Experimental Results

Conclusion

Introduction


Optimal Transport in Generative Models

Proposed Optimal Transport Modeling

Experimental Results

Conclusion

Introduction to Optimal Transport



(a) Monge formulation (1785). (b) Kantorovich formulation (1942).

Monge Problem (MP):

- $\text{Cost}(\mu, \nu) \stackrel{\text{def}}{=} \inf_{T \# \mu = \nu} \int_{\mathcal{X}} c(x, T(x)) d\mu(x)$.
- The feasible set of solutions can be empty.
- MP does not allow mass splitting.

Kantorovich Problem (KP):

- $\text{Cost}(\mu, \nu) \stackrel{\text{def}}{=} \inf_{\pi \in \Pi(\mu, \nu)} \int_{\mathcal{X} \times \mathcal{Y}} c(x, y) d\pi(x, y)$.

Kantorovich Dual Problem (Kantorovich, 1942, KDP) :

$$\text{Cost}(\mu, \nu) = \sup_{(u,v)} \left\{ \int_{\mathcal{X}} u(x) d\mu(x) + \int_{\mathcal{Y}} v(y) d\nu(y) : u(x) + v(y) \leq c(x, y) \right\},$$

Using c -transform, i.e., $v^c(x) = \inf_{y \in \mathcal{Y}} \{c(x, y) - v(y)\}$, KDP becomes (Villani, 2008, §5):

$$\text{Cost}(\mu, \nu) = \sup_v \left\{ \int_{\mathcal{X}} v^c(x) d\mu(x) + \int_{\mathcal{Y}} v(y) d\nu(y) \right\}$$

- KDP allows optimization over one functional u or v .

Introduction


Optimal Transport in Generative Models

Proposed Optimal Transport Modeling


Experimental Results

Conclusion

Optimal Transport in Generative Models



(a) OT cost as the loss for the generator.



(b) OT map as the generative model.

Optimal Transport as the Generative Map




Figure 3: The pipeline of most prevalent approaches (Taghvaei & Jalali, 2019; Makkuva et al., 2020; Korotin et al., 2021).

- Compute OT maps in latent spaces of autoencoders.
- OT maps are not considered in ambient spaces.
- Lack scalability due to poor expressivity of ICNNs.

Overview

Introduction

Optimal Transport in Generative Models

Proposed Optimal Transport Modeling

Experimental Results

Conclusion

Equal Dimensions of Input and Output

With $\psi(y) \stackrel{\text{def}}{=} \frac{1}{2}\|y\|^2 - v(y)$, the c -transform of Kantorovich potential $v(y)$ becomes:

$$\begin{aligned} v^c(x) &= \inf_{y \in \mathbb{R}^D} \left\{ \frac{1}{2}\|x - y\|^2 - v(y) \right\} \\ &= \frac{1}{2}\|x\|^2 - \sup_{y \in \mathbb{R}^D} \{ \langle x, y \rangle - \psi(y) \} = \frac{1}{2}\|x\|^2 - \bar{\psi}(x). \end{aligned} \tag{1}$$

Substituting (1), KDP simplifies to

$$\text{Constant}(\mu, \nu) - \inf_{\psi} \left\{ \sup_T \int_{\mathcal{X}} \{ \langle x, T(x) \rangle - \psi(T(x)) \} d\mu(x) + \int_{\mathcal{Y}} \psi(y) d\nu(y) \right\} \tag{2}$$

- We prove that the OT map T^* is the solution of the inner maximization problem (2), see **Lemma 4.2** in our paper.

Unequal Dimensions of Input and Output




Figure 4: The scheme of our approach.

$$\inf_{\psi} \sup_G \left\{ \int_{\mathcal{X}} \{ \langle Q(x), G(x) \rangle - \psi(G(x)) \} d\mu(x) + \int_{\mathcal{Y}} \psi(y) d\nu(y) \right\} \quad (3)$$

- We prove that the OT map G^* and potential ψ^* are the solutions of the saddle point problem (3).

Introduction


Optimal Transport in Generative Models

Proposed Optimal Transport Modeling


Experimental Results

Conclusion

Generative Modeling: Qualitative Results



(a) CelebA, 64×64 , RGB, FID: 6.5



(b) CelebA, 128×128 , RGB, FID: 24.58

Generative Modeling: Quantitative Results

Table 1: Results on CelebA, 64x64 dataset.

Model	Related Work	FID ↓
DCGAN	Radford et al. (2016)	52.0
DRAGAN	Kodali et al. (2017)	42.3
BEGAN	Berthelot et al. (2017)	38.9
NVAE	Vahdat & Kautz (2020)	13.4
NCP-VAE	Aneja et al. (2021)	5.2
WGAN	Arjovsky et al. (2017)	41.3
WGAN-GP	Gulrajani et al. (2017)	30.0
WGAN-QC	Liu et al. (2019)	12.9
AE-OT	An et al. (2020a)	28.6
W2GN+AE	Korotin et al. (2021)	17.2
AE-OT-GAN	An et al. (2020b)	7.8
OTM	Ours	6.5

Unpaired Restoration



(a) Noisy

(b) Pushforward

(c) Original

Model	Denosing	Colorization	Inpainting
Input	166.59	32.12	47.65
WGAN-GP	25.49	7.75	16.51
OTM-GP (ours)	10.95	5.66	9.96
OTM (ours)	5.92	5.65	8.13

Overview

Introduction

Optimal Transport in Generative Models

Proposed Optimal Transport Modeling

Experimental Results

Conclusion

Summary


- We proposed to fit the OT map in quadratic transport cost $\mathcal{W}_2^2(\mu, \nu)$ which acts as a generative map.
- We developed an end-to-end solution for equal and unequal dimensions of input and output distributions.
- We demonstrated OTM in unpaired restoration tasks: denoising, colorization, and inpainting.

Conclusion

Generative Modeling with Optimal Transport Maps

Stop by our poster to learn more about our research.

arxiv.org/pdf/2110.02999.pdf



github.com/LituRout/OptimalTransportModeling

- Dongsheng An, Yang Guo, Na Lei, Zhongxuan Luo, Shing-Tung Yau, and Xianfeng Gu. Ae-ot: A new generative model based on extended semi-discrete optimal transport. In International Conference on Learning Representations, 2020a. URL <https://openreview.net/forum?id=HkldyTNYwH>.
- Dongsheng An, Yang Guo, Min Zhang, Xin Qi, Na Lei, and Xianfang Gu. Ae-ot-gan: Training gans from data specific latent distribution. In European Conference on Computer Vision, pp. 548–564. Springer, 2020b.

Jyoti Aneja, Alexander Schwing, Jan Kautz, and Arash Vahdat.

Ncp-vae: Variational autoencoders with noise contrastive priors. In Advances in Neural Information Processing Systems Conference, 2021.

Martin Arjovsky, Soumith Chintala, and Léon Bottou.

Wasserstein generative adversarial networks. In International conference on machine learning, pp. 214–223. PMLR, 2017.

David Berthelot, Thomas Schumm, and Luke Metz. Began:

Boundary equilibrium generative adversarial networks. arXiv preprint arXiv:1703.10717, 2017.

- Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron Courville. Improved training of Wasserstein GANs. In Proceedings of the 31st International Conference on Neural Information Processing Systems, pp. 5769–5779, 2017.
- Leonid Vitalevich Kantorovich. On the translocation of masses. In Dokl. Akad. Nauk SSSR, volume 37, pp. 199–201, 1942.
- Naveen Kodali, Jacob Abernethy, James Hays, and Zoltan Kira. On convergence and stability of GANs. arXiv preprint arXiv:1705.07215, 2017.

Alexander Korotin, Vage Egiazarian, Arip Asadulaev, Alexander Safin, and Evgeny Burnaev. Wasserstein-2 generative networks. In International Conference on Learning Representations, 2021. URL

https://openreview.net/forum?id=bEoxzW_EXsa.

Huidong Liu, Xianfeng Gu, and Dimitris Samaras. Wasserstein gan with quadratic transport cost. In Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), October 2019.

Ashok Makkuva, Amirhossein Taghvaei, Sewoong Oh, and Jason Lee. Optimal transport mapping via input convex neural networks. In International Conference on Machine Learning, pp. 6672–6681. PMLR, 2020.

Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. In 4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings, 2016.
URL <http://arxiv.org/abs/1511.06434>.

Amirhossein Taghvaei and Amin Jalali. 2-wasserstein approximation via restricted convex potentials with application to improved training for gans. arXiv preprint arXiv:1902.07197, 2019.

Arash Vahdat and Jan Kautz. Nvae: A deep hierarchical variational autoencoder. In Advances in Neural Information Processing Systems Conference, 2020.

Cédric Villani. Optimal transport: old and new, volume 338. Springer Science & Business Media, 2008.