Wasserstein of Wasserstein Loss for Learning Generative Models

ICML 2019

Yonatan Dukler*, Wuchen Li*, Alex Tong Lin*, Guido Montúfar*
{ydukler, wcl, atlin, montufar}@math.ucla.edu

June 11, 2019

Wasserstein GAN

- In the Wasserstein GAN framework [2, 1], given $\mathcal{P}_{\text{data}}$, $G(\Theta)$, and $F(\Phi, x)$ trained as

$$\min_{\Theta} \max_{\Phi} \ \mathbb{E}_{x \sim \mathcal{P}_{G(\Theta)}} F(\Phi, x) - \mathbb{E}_{x \sim \mathcal{P}_{\text{data}}} F(\Phi, x) + \underbrace{\lambda \mathbb{E}_{\hat{x}}(\|\nabla_x F(\Phi, x)\|_{L^2} - 1)^2}_{\text{Lipschitz enforcing term}}.$$

- Based on the Kanatrovich-Rubenstein duality

$$\mathcal{W}_{\mathcal{D}}(\mu, \nu) = \sup_{f \in \mathcal{C}} \left\{ \int f d\nu dx - \int f d\mu dx \right\}$$

where $\mathcal{C} := \{f : Lip_d(f) \leq 1\}$ and $d(x, y)$ is a ground metric on samples (high dim. images) [3].

Adding Prior Knowledge When Choosing $d(x, y)$

$$\boxed{\text{C} := \{\text{f} : \mathbf{Lip}_d(f) \leq 1\} \text{ and } \mathbf{d}(\mathbf{x}, \mathbf{y})}$$

(A) Geometry of natural images: Lie in lower dimensional manifold than entire image space.

(B) Natural images satisfy well known symmetries and transformations cf. data-augmentation techniques.

A good metric should satisfy these priors.

Wasserstein metric on images

- Image as a distribution over pixels
- Pixel space is a graph
- Cost of transport based on graph's weights

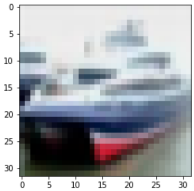

Raster image

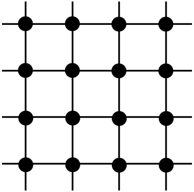

Pixel space graph

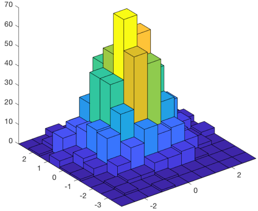

Distribution

Wasserstein of Wasserstein GAN (WWGAN)

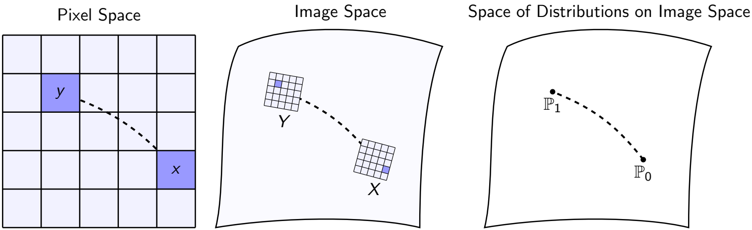

WWGAN

- Wasserstein distance on images is computationally expansive
- Instead utilize Wasserstein-2 metric's Riemannian structure.
- **WWGAN:** Replacing L^2 gradient with Wasserstein-2 gradient ($\nabla_x^{\mathcal{W}_2} F$) in WGAN
- Penalty is computed efficiently via convolutions. (add'l. comp. cost $\sim 8\%$)

WWGAN is more robust to natural perturbations

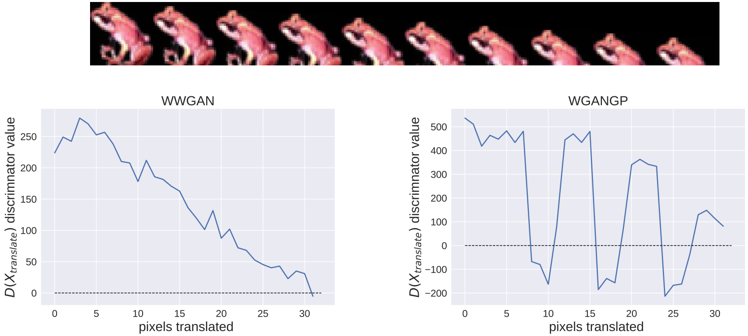

References

Martin Arjovsky, Soumith Chintala, and Léon Bottou. "Wasserstein GAN". In: *arXiv preprint arXiv:1701.07875* (2017).

Ishaan Gulrajani et al. "Improved Training of Wasserstein GANs". In: *Advances in Neural Information Processing Systems*. 2017, pp. 5767–5777.

Jonas Adler and Sebastian Lunz. "Banach Wasserstein GAN". In: *Advances in Neural Information Processing Systems*. 2018, pp. 6755–6764.