Kantorovich-Rubinstein Duality

John Thickstun

The Wasserstein GAN [Arjovsky et al., 2017] seeks to minimize the objective

$$\arg\min_{\theta} W(p, p_\theta) = \arg\min_{\theta} \inf_{\pi \in \Pi(p,q)} \mathop{\mathbb{E}}_{(x,y)\sim\pi} \left[\|x - y\|_2\right].$$

In this form, the inner estimation of the Wasserstein distance $W(p, p_\theta)$ is intractable. But using a delicate duality argument, we are able to reformulate the Wasserstein distance as the solution to a maximization over 1-Lipschitz functions. This turns the Wasserstein GAN optimization problem into a saddle-point problem, analogous to the f-GAN. The following proof is loosely based on Basso [2015]; here we take a more concrete approach and fill in the details of the argument.

Theorem. *(Kantorovich-Rubinstein)*

$$W(p, q) = \inf_{\pi \in \Pi(p,q)} \mathop{\mathbb{E}}_{(x,y)\sim\pi} \left[\|x - y\|_2\right] = \sup_{\|h\|_L \le 1} \left[\mathop{\mathbb{E}}_{x\sim p}[h(x)] - \mathop{\mathbb{E}}_{y\sim q}[h(x)]\right].$$

Proof. Introduce Lagrange multipliers $f, g : \mathcal{X} \to \mathbb{R}$ (bounded, measurable):

$$\begin{aligned} L(\pi, f, g) = \int_{\mathcal{X}\times\mathcal{X}} \|x - y\|^2 \pi(x, y)\, dy\, dx &+ \int_{\mathcal{X}} \left(p(x) - \int_{\mathcal{X}} \pi(x, y)\, dy\right) f(x)\, dx \\ &+ \int_{\mathcal{X}} \left(q(y) - \int_{\mathcal{X}} \pi(x, y)\, dx\right) g(y)\, dy. \end{aligned}$$

Collecting terms algebraically, we can rewrite the Lagrangian as

$$L(\pi, f, g) = \mathop{\mathbb{E}}_{x\sim p}[f(x)] + \mathop{\mathbb{E}}_{y\sim q}[g(y)] + \int_{\mathcal{X}\times\mathcal{X}} \Big(\|x - y\|_2 - f(x) - g(y)\Big)\pi(x, y)\, dy\, dx.$$

And we appeal to strong duality to write

$$W(p, p_g) = \inf_{\pi} \sup_{f,g} L(\pi, f, g) = \sup_{f,g} \inf_{\pi} L(\pi, f, g).$$

Note that if $\|x - y\|_2 < f(x) + g(y)$ for some $x, y \in \mathcal{X}$ then we can concentrate mass of π at (x, y) and send $L(\pi, f, g)$ to $-\infty$. Therefore, for all x, y, we must have

$$f(x) + g(y) \le \|x - y\|_2.$$

And with this constraint, the best we can do by minimizing over π is to set $\pi = 0$:

$$\sup_{f,g} \inf_{\pi} L(\pi, f, g) = \sup_{\substack{f,g \\ f(x)+g(y)\le\|x-y\|_2}} \left[\mathop{\mathbb{E}}_{x\sim p}[f(x)] + \mathop{\mathbb{E}}_{y\sim q}[g(y)]\right] = W(p, q).$$

We now perform an amazing sleight-of-hand. First, observe that optimizing over the class of 1-Lipschitz functions give us a lower bound on the Wasserstein distance; if h is 1-Lipschitz then

$$\begin{aligned}\mathop{\mathbb{E}}_{x\sim p}[h(x)] - \mathop{\mathbb{E}}_{y\sim q}[h(y)] &= \int_{\mathcal{X}\times\mathcal{X}} (h(x) - h(y))\, \pi(x,y)\, dx\, dy \\ &\leq \int_{\mathcal{X}\times\mathcal{X}} \|x-y\|_2\, \pi(x,y)\, dx\, dy \leq W(p,q).\end{aligned}$$

It follows that

$$\sup_{\|h\|_L\leq 1} \left[\mathop{\mathbb{E}}_{x\sim p}[h(x)] - \mathop{\mathbb{E}}_{y\sim q}[h(y)]\right] \leq W(p,q).$$

Next we will show that the upper bound is attained.

Consider the function κ defined by

$$x \mapsto \inf_u \Big[\|x-u\|_2 - g(u)\Big].$$

Because g is bounded, the infimum is finite and this function is well-defined. Claim: κ is 1-Lipschitz. Given $x, y \in \mathcal{X}$, for any $u \in \mathcal{X}$ the triangle inequality give us

$$\kappa(x) \leq \|x-u\|_2 - g(y) \leq \|x-y\|_2 + \|y-u\|_2 - g(u).$$

This holds for any u and therefore

$$\kappa(x) \leq \|x-y\|_2 + \inf_u \big[\|x-u\|_2 - g(u)\big] = \|x-y\|_2 + \kappa(y).$$

I.e. $\kappa(x) - \kappa(y) \leq \|x-y\|_2$; exchanging x and y gives us $\kappa(y) - \kappa(x) \leq \|x-y\|_2$ and consequently κ is 1-Lipschitz:

$$|\kappa(x) - \kappa(y)| \leq \|x-y\|_2.$$

For any pair f, g that satisfies $f(x) + g(y) \leq \|x-y\|_2$,

$$f(x) \leq \kappa(x) \leq \|x-x\|_2 - g(x) = -g(x).$$

From these facts, we have

$$\mathop{\mathbb{E}}_{x\sim p}[f(x)] + \mathop{\mathbb{E}}_{y\sim q}[g(y)] \leq \mathop{\mathbb{E}}_{x\sim p}[\kappa(x)] - \mathop{\mathbb{E}}_{y\sim q}[\kappa(y)]$$

And we conclude that

$$\begin{aligned}W(p,q) &= \sup_{\substack{f,g \\ f(x)+g(y)\leq\|x-y\|_2}} \left[\mathop{\mathbb{E}}_{x\sim p}[f(x)] + \mathop{\mathbb{E}}_{y\sim q}[g(y)]\right] \\ &\leq \sup_{\|h\|_L\leq 1} \left[\mathop{\mathbb{E}}_{x\sim p}[h(x)] - \mathop{\mathbb{E}}_{y\sim q}[h(y)]\right] \leq W(p,q).\end{aligned}$$ □

References

Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein gan. *International Conference on Machine Learning*, 2017. (document)

Giuliano Basso. A hitchhikers guide to wasserstein distances, 2015. (document)