

Structured Model Learning

Generative Adversarial Networks

Boris Flach
Czech Technical University in Prague

- ◆ Generative adversarial networks
- ◆ Wasserstein distance for distributions
- ◆ Wasserstein GANs

Generative adversarial networks

(Goodfellow et al., 2014)

As before with VAEs

- ◆ Let \mathcal{X} denote the image space and \mathcal{Z} denote a **noise** space. We may assume e.g. $\mathcal{Z} = \mathbb{R}^n$. We fix a simple distribution $Z \sim \mathcal{N}(0, \mathbb{I})$ on it.
- ◆ Given a sample $\mathcal{S}^m = \{x^j \in \mathcal{X} \mid j = 1, \dots, m\}$, we want to learn a parametrised mapping $g_w: \mathcal{Z} \rightarrow \mathcal{X}$ such that the joint model $p_w(x, z)$ is likely to generate images similar to those in \mathcal{S}^m .
- ◆ We consider a classifier network with a single output neuron. This network should discriminate natural images (coming from \mathcal{S}^m) from synthetic images generated by $p_w(x, z)$. Its output $d_v(x)$ is considered as probability of the class “natural”.

Its loss is defined by

$$L(w, v) = \mathbb{E}_{x \sim \mathcal{S}^m} \left[\log d_v(x) \right] + \mathbb{E}_{z \sim p(z)} \left[\log(1 - d_v \circ g_w(z)) \right]$$

and the optimisation task is

$$\min_w \max_v L(v, w).$$

Simplified analysis of GANs

Let us analyse this saddle point task under the assumptions

- ◆ infinite training set \mathcal{S}^m , $m \rightarrow \infty$, denote the distribution by $p_*(x)$
- ◆ the generator model has infinite(!) capacity

The maximum w.r.t. the discriminator is achieved at $d^*(x) = \frac{p_d(x)}{p_d(x)+p_g(x)}$ because

$$L(g, d) = \int p_*(x) \log d(x) dx + \int p(z) \log(1 - d \circ g(z)) dz =$$
$$\int p_*(x) \log(d(x)) + p_g(x) \log(1 - d(x)) dx$$

Then the minimum w.r.t. the generator is achieved at $p_g = p_*$ because

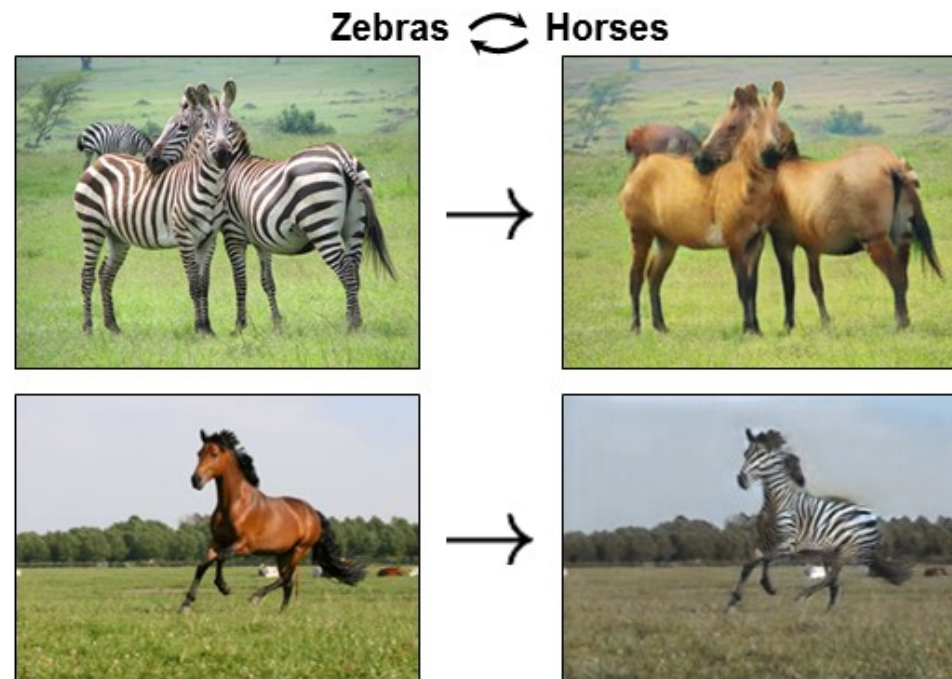
$$C(g) = \max_d L(g, d) = \mathbb{E}_{x \sim p_*} \left[\log \frac{p_*(x)}{p_*(x) + p_g(x)} \right] + \mathbb{E}_{x \sim p_g} \left[\log \frac{p_g(x)}{p_*(x) + p_g(x)} \right]$$

Practical issues of GANs

Quite impressive results achieved by deep convolutional GANs, but there are technical problems

- ◆ what happens with $\log(1 - d(g(z)))$ early in learning, when g is poor?
- ◆ mode collapse: what happens if g maps all $z \in \mathcal{Z}$ on one x ?
- ◆ searching saddle points of functions is complicated

Learning GANs is difficult and sensitive to hyper-parameter and architecture setup.



(Zhu et al., 2017) Cycle GANs

Wasserstein distance for distributions

Many distances for distributions are problematic when

- ◆ one distribution is discrete and the other is absolutely continuous (has a density)
- ◆ the distributions have disjoint support

Consider the following situation: let \mathbb{P}_* be a fixed distribution on \mathcal{X} , let Z be a random variable on \mathcal{Z} and let $g_w: \mathcal{Z} \rightarrow \mathcal{X}$ be a parametrised family of mappings. Denoting the distribution of $X = g_w(Z)$ by \mathbb{P}_w , we want a distribution distance $W(\mathbb{P}_*, \mathbb{P}_w)$ that behaves nicely as a function of w .

Defintion 1. The Wasserstein distance of distributions \mathbb{P}_1 and \mathbb{P}_2 is defined by

$$W(\mathbb{P}_1, \mathbb{P}_2) = \inf_{\gamma \in \Pi(\mathbb{P}_1, \mathbb{P}_2)} \mathbb{E}_{(x,y) \sim \gamma} \|x - y\|,$$

where $\Pi(\mathbb{P}_1, \mathbb{P}_2)$ is the set of all joint distributions $\gamma(x, y)$ with marginals \mathbb{P}_1 and \mathbb{P}_2

Remarks:

- ◆ this distance is sometimes called earth-mover distance,
- ◆ $\gamma(x, y)$ is interpreted as a transport plan and W is the cost of the optimal plan

Wasserstein GANs

(Arjovsky et al., 2017)

Theorem 1. (w/o proof) Let g_w be a neural network composed of affine transformations and smooth Lipschitz pointwise non-linearities. Let $p(z)$ be a prior over z with bounded expectation of its norm. Then $W(\mathbb{P}_*, \mathbb{P}_w)$ is continuous and almost everywhere differentiable in w .

How to compute the distance? The task is convex and its dual is

$$W(\mathbb{P}_*, \mathbb{P}_w) = \sup_{\|f\|_L \leq 1} \left[\mathbb{E}_{x \sim \mathbb{P}_*} [f(x)] - \mathbb{E}_{x \sim \mathbb{P}_w} [f(x)] \right],$$

where the supremum is over all 1-Lipschitz functions $f: \mathcal{X} \rightarrow \mathbb{R}$. Under some mild conditions on g_w

- ◆ the problem $\max_{\|f\|_L \leq 1} \mathbb{E}_{x \sim \mathbb{P}_*} [f(x)] - \mathbb{E}_{x \sim \mathbb{P}_w} [f(x)]$ has a solution and
- ◆ the gradient is $\nabla_w W(\mathbb{P}_*, \mathbb{P}_w) = -\mathbb{E}_{z \sim p(z)} [\nabla_w f(g_w(z))]$

Wasserstein GANs

- ◆ both, the generator mapping $g_w: \mathcal{Z} \rightarrow \mathcal{X}$ and the critic $f_v: \mathcal{X} \rightarrow \mathbb{R}$ are implemented as neural networks.
- ◆ f_v must be a 1-Lipschitz function (or at least a K -Lipschitz function).

In the inner loop of the algorithm (given the current w) the critic is learned on x -batches sampled from the data and z -batches sampled from the prior noise distribution. An additional gradient norm penalty is used to enforce the Lipschitz requirement. The critic is learned till optimality.

In the outer loop w is updated based on a z -batch sampled from $p(z)$ by using the gradient formula given above.

