ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Finding the Mixed Nash Equilibria of Generative Adversarial Networks

Ya-Ping Hsieh — EPFL
Chen Liu — EPFL
Volkan Cevher — EPFL

① Motivation

Problem formulation for training Wasserstein GANs (WGAN):

$$\min_{\theta\in\Theta}\max_{\mathbf{w}\in\mathcal{W}} \mathbb{E}_{X\sim\mathbb{P}_{real}}[f_{\mathbf{w}}(X)] - \mathbb{E}_{X\in\mathbb{P}_\theta}[f_{\mathbf{w}}(X)]$$

But, training GANs is not easy.

- Highly unstable training.
- Lack of theory about the existence of saddle points.

② Proposal: Mixed Nash Equilibria (NE)

WGAN with mixed strategy [a]

$$\min_{\nu\in\mathcal{M}(\Theta)}\max_{\mu\in\mathcal{M}(\mathcal{W})} \mathbb{E}_{\mathbf{w}\sim\nu}\mathbb{E}_{X\sim\mathbb{P}_{real}}[f_{\mathbf{w}}(X)] - \mathbb{E}_{\mathbf{w}\sim\mu}\mathbb{E}_{\theta\sim\nu}\mathbb{E}_{X\sim\mathbb{P}_\theta}[f_{\mathbf{w}}(X)]$$

Define $g(\mathbf{w}) = \mathbb{E}_{X\sim\mathbb{P}_{real}}[f_{\mathbf{w}}(X)]$, $G\nu(\mathbf{w}) = \mathbb{E}_{\theta\sim\nu,X\sim\mathbb{P}_\theta}[f_{\mathbf{w}}(X)]$ and $\langle\nu,h\rangle = \mathbb{E}_\nu h$, WGAN with mixed strategy can be reformulated

$$\min_{\nu\in\mathcal{M}(\Theta)}\max_{\mu\in\mathcal{M}(\mathcal{W})} \langle\mu,g\rangle - \langle\mu,G\nu\rangle$$

This is exactly an *infinite* dimensional two-player game.

How to find Mixed NE? ▷ **Entropic Mirror Descent (MD)**

Start with a two-player game with *finite* actions

$$\min_{\mathbf{p}\in\Delta_m}\max_{\mathbf{q}\in\Delta_n} \langle\mathbf{q},\mathbf{a}\rangle - \langle\mathbf{q},\mathbf{A}\mathbf{p}\rangle$$

Entropic MD learns an $O(T^{-1/2})$-NE in T iterations.

$$\begin{cases}\mathbf{p}_{t+1} = MD_\eta(\mathbf{p}_t, -\mathbf{A}^T\mathbf{q}_t)\\ \mathbf{q}_{t+1} = MD_\eta(\mathbf{q}_t, -\mathbf{a}+\mathbf{A}\mathbf{p}_t)\end{cases}$$

Here, MD_η is the MD iterate defined by entropy function $\phi(\mathbf{z}) = \sum_{i=1}^d z_i\log z_i$ and its Fenchel dual $\phi^*(\mathbf{y}) = \log\sum_{i=1}^d e^{y_i}$.

$$\mathbf{z}' \equiv MD_\eta(\mathbf{z},\mathbf{b}) = \nabla\phi^*(\nabla\phi(\mathbf{z}) - \eta\mathbf{b})$$

$$\mathbf{z}'_i = \frac{\mathbf{z}_i e^{-\eta\mathbf{b}_i}}{\sum_{i=1}^d \mathbf{z}_i e^{-\eta\mathbf{b}_i}}$$

Can **Entropic MD** generalize to *infinite* dimensional two-player game formulation of WGAN and enjoy the same convergence rate?

Answer is Yes!

---

[a] Although we use WGAN here as example, our methods can be applied to other GANs as well.

③ Main Theory

```
ENTROPIC MD FOR MIX-STRATEGY WGAN
  Initial distribution $\mu_1$, $\nu_1$, learning rate $\eta$ given
  for $t = 1, 2, \dots, T$ do
      $\nu_{t+1} = MD_\eta(\nu_t, -G^T\mu_t)$            ▷ Update $\nu$
      $\mu_{t+1} = MD_\eta(\mu_t, -g + G\nu_t)$          ▷ Update $\mu$
  return $\bar\mu = \frac{1}{T}\sum_{t=1}^T \mu_t$ and $\bar\nu = \frac{1}{T}\sum_{t=1}^T \nu_t$.
```

Main Theorem (informal)

Assmue $-G^T\mu_t$ and $-g+G\nu_t$ are bounded and smooth.

1. If we have access to deterministic values of $-G^T\mu_t$ and $-g+G\nu_t$, then Entropic MD achieves $O(T^{-1/2})$ - NE in T iterations with proper learning rate.
2. If we have access to unbiased stochastic values with bounded variance of $-G^T\mu_t$ and $-g+G\nu_t$, then Entropic MD achieves $O(T^{-1/2})$ - NE in expectation in T iterations with proper learning rate.

④ From Theory to Practice

Four steps to make the algorithm implementable.

1. Using the property of Shannon entropy, the property measure of μ_t and ν_t can be expressed in terms of the history.

$$d_{\mu_t} = \frac{exp\{(t-1)g - G\sum_{s=1}^{t-1}\nu_s\}d_{\mathbf{w}}}{\int exp\{(t-1)g - G\sum_{s=1}^{t-1}\nu_s\}d_{\mathbf{w}}}, d_{\nu_t} = \frac{exp\{G^T\sum_{s=1}^{t-1}\mu_s\}d_\theta}{\int exp\{G^T\sum_{s=1}^{t-1}\mu_s\}d_\theta}$$

2. We use unbiased empirical average to estimate expectation over $\{\mu_t\}_{t=1}^T$ and $\{\nu_t\}_{t=1}^T$ if we can sample from them.
3. For distributions of density function $e^{-h}d_{\mathbf{z}}$, we can use Stochastic Gradient Langevin Dynamics (SGLD) to obtain samples.

$$\mathbf{z}_{t+1} = \mathbf{z}_t - \gamma\hat\nabla h(\mathbf{z}_t) + \sqrt{2\gamma}\epsilon\xi_t$$

γ, ϵ and ξ_t are step size, thermal noise and Gaussian noise.

4. To avoid memory overflow because of storing model samples in each iteration, we summary each distribution only by its mean.

⑤ Experiments

We name our method Mirror-GAN and compare it with other method to train GANs.