Nested Optimization in Games

Rozhina Ghanavi

University of Toronto

November 2, 2019

Different types of games

- Simultaneous games
- Sequential games, Stackelberg games: consists of a leader and a follower, the follower observes the leader's quantity choice and choose action based on that.

$$\min_{x_1 \in X_1} \left\{ f_1\left(x_1, x_2\right) \middle| x_2 \in \arg\min_{y \in X_2} f_2\left(x_1, y\right) \right\} \tag{1}$$

Motivations

- Why we are interested in games?
  Use cases in ML: GANs, adverserial training, and primal-dual RL.
- What is the problem?
  Simple gradient based methods are not working and we are looking for other optimization methods.

GANs

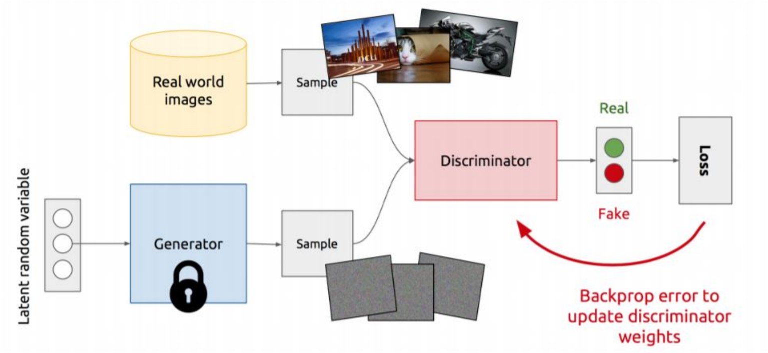

from Binglin, Shashan, and Bhargav.

GANs

$$\min_G \max_D V(G, D) \tag{2}$$

$$V(G, D) = \int_{\boldsymbol{x}} p_{\text{data}}(\boldsymbol{x}) \log(D(\boldsymbol{x}))dx + \int_z p_{\boldsymbol{z}}(\boldsymbol{z}) \log(1 - D(g(\boldsymbol{z})))dz \tag{3}$$

- Equilibrium no longer consist of a single loss, hence nested optimization.

GAN optimization algorithm

- GAN optimization is based on gradient descent ascent (GDA).
- Update the discriminator by ascending gradient:

$$\nabla_{\theta_d} \frac{1}{m} \sum_{i=1}^{m} \left[\log D\left(\boldsymbol{x}^{(i)}\right) + \log\left(1 - D\left(G\left(\boldsymbol{z}^{(i)}\right)\right)\right) \right] \tag{4}$$

- Update the generator by descending gradient:

$$\nabla_{\theta_g} \frac{1}{m} \sum_{i=1}^{m} \log\left(1 - D\left(G\left(z^{(i)}\right)\right)\right) \tag{5}$$

Convergence of Learning Dynamics in Stackelberg Games

T. Fiez, B. Chasnov, and L. J. Ratliff

Games setting

- They considered a sequential Stackelberg game (pure strategy: Stackelberg equilibrium).
- This game consists of a leader and a follower.

Finite-Time High-Probability Guarantees

The follower converges to:

$$\mathrm{P}\left(\|x_{2,n} - z_n\| \leq \varepsilon, \forall n \geq \bar{n} | x_{2,n_0}, z_{n_0} \in B_{q_0}\right) \to 1 \tag{6}$$

where, $z_k = r\left(x_{1,k}\right)$ and $r(x)$ is the implicit function.

Finite-Time High-Probability Guarantees

The leader converges to:

$$\mathrm{P}\left(\left\|x_{1,n} - x_1\left(\hat{t}_n\right)\right\| \leq \varepsilon, \forall n \geq \bar{n} | x_{n_0}, x_{n_0} \in B_{q_0}\right) \to 1 \tag{7}$$

Take away point, we converge to a neighborhood of a Stackelberg equilibrium in finite-time, with a good probability!!

Conclusions

- Shows that there exist stable attractors of simultaneous gradient play that are Stackelberg equilibria and not Nash equilibria.

Conclusions

- Shows that there exist stable attractors of simultaneous gradient play that are Stackelberg equilibria and not Nash equilibria.
- A finite-time high probability bound for local convergence to a neighborhood of a stable Stackelberg equilibrium in general-sum games.

On Solving Minimax Optimization Locally: A Follow-the-Ridge Approach

Under blind review at ICLR 2020

Games Setting

- Differentiable sequential games,
- Two players,
- zero-sum, minimax,

$$\min_{\mathbf{x}\in\mathbb{R}^n} \max_{\mathbf{y}\in\mathbb{R}^m} f(\mathbf{x},\mathbf{y}) \tag{8}$$

How to solve minimax optimization?

- Gradient descent-ascent (GDA)
  - Problem 1. The goal is to converge to local minimax points, but GDA fails.
    Problem 2. Strong rotation around fixed points. Requires small learning rate.
- Follow-the-Ridge (FR), proposed by this paper.
  - Solves both issues.

Follow the ridge (FR)

- GDA tends to drift away from the ridge.
- How to solve it?
  By definition, a local minimax has to lie on a ridge.
  So, follow the ridge!

FR algorithm

Algorithm Follow-the-Ridge (FR). Differences from gradient descent-ascent are shown in blue.

Require: Learning rate $\eta_{\mathbf{x}}$ and $\eta_{\mathbf{y}}$; number of iterations T.

1: **for** $t = 1, ..., T$ **do**

2: $\quad \mathbf{x}_{t+1} \leftarrow \mathbf{x}_t - \eta_{\mathbf{x}} \nabla_{\mathbf{x}} f(\mathbf{x}_t, \mathbf{y}_t)$ ▷ gradient descent

3: $\quad \mathbf{y}_{t+1} \leftarrow \mathbf{y}_t + \eta_{\mathbf{y}} \nabla_{\mathbf{y}} f(\mathbf{x}_t, \mathbf{y}_t) + \eta_{\mathbf{x}} \mathbf{H}_{\mathbf{yy}}^{-1} \mathbf{H}_{\mathbf{yx}} \nabla_{\mathbf{x}} f(\mathbf{x}_t, \mathbf{y}_t)$ ▷ modified gradient ascent

---

FR algorithm

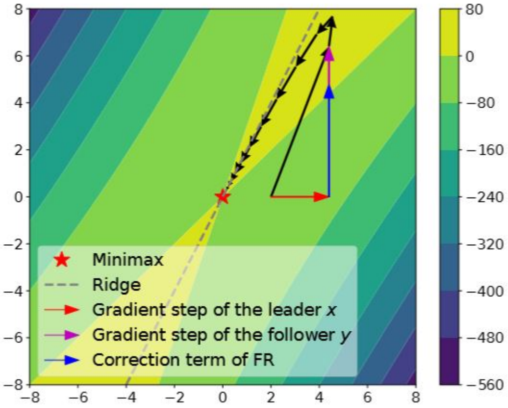

FR results

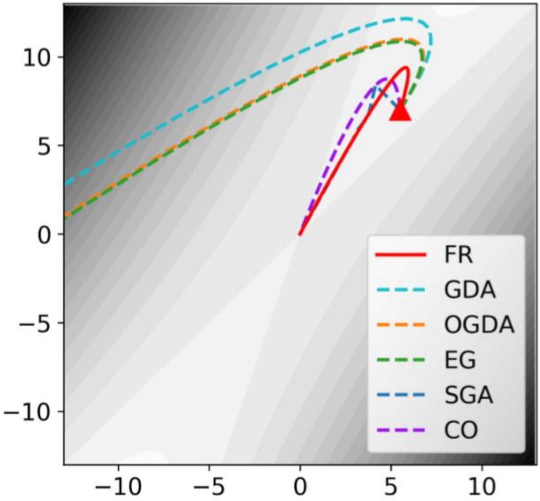

from the paper.

Conclusion

- It addresses the rotational behaviour of gradient dynamics and allows larger learning rate than GDA.

Conclusion

- It addresses the rotational behaviour of gradient dynamics and allows larger learning rate than GDA.
- Standard acceleration techniques can be added.

Conclusion

- It addresses the rotational behaviour of gradient dynamics and allows larger learning rate than GDA.
- Standard acceleration techniques can be added.
- In general we were so hyped about using GD in neural networks because we knew they are converging, this method can be viewed as a similar way to think about GANs