Equilibrium finding and learning

Notes on min-max optimization from the Simons Institute tutorial

Geelon So, agso@eng.ucsd.edu
Algorithmic game theory reading group — February 2, 2022

Source material

The material for this presentation is essentially all copied directly from the Simons Institute Learning and Games Bootcamp, from Constantinos Daskalakis' tutorial.

- Watch his talk here.
- See his slides here.

Classical learning theory framework

Model:

- nature chooses a fixed data distribution $\mathcal{D}$
- learner aims to find parameters x minimizing some loss $\ell(x; \mathcal{D})$
- learner only gets sample access $\mathcal{Z} = \{Z_1, \ldots, Z_n\}$ where $Z_i \sim \mathcal{D}$
- learner minimizes empirical estimate $\ell(x; \mathcal{Z})$

Machine learning: optimization viewpoint

Success in machine learning through casting learning as an optimization problem:

$$\min_x \ell(x)$$

- x is often very high-dimensional
- ℓ is nonconvex objective, but Lipschitz and smooth
  - access via 0th or 1st order information, e.g. $\ell(x)$ and $\nabla \ell(x)$
- optimize using gradient descent: $x \leftarrow x - \eta \nabla \ell(x)$
  - theoretical guarantee: efficiently find local minima
  - empirical finding: local minima good enough

Non-stationary (reactive) environments

Beyond the classical setting: can we learn in settings where the agent's choice x can affect its environment? The environment can be reactive due to:

- the presence of other learning agents with conflicting or aligned interests
  - markets, traffic, games, ecological/biological systems
  - distributed optimization, model of brain, evolution of language, homeostasis
- noise/adversaries/distribution shifts poisoning, corrupting, or changing the data
  - adversarial learning, gaming a decision making system
- the need to enforce constraints on the learning outcome
  - introduce constraints via competing interests (like GANs)
  - robust statistics, privacy-preserving learning, causal inference

Machine learning: equilibrium computation viewpoint

- **Minimization:** an agent is learning and decision-making in a stationary environment
- **Equilibrium computation:** an agent is learning and decision-making in an environment that is reactive

Multi-agent learning setting

Setting:

- k agents
- the ith agent gets to choose parameter x_i in:

$$\mathbf{x} \equiv (x_1, \ldots, x_k)$$

- the ith agent aims to minimize $\ell_i(x)$, where:[1]

$$\ell_i : \mathcal{X}_1 \times \ldots \times \mathcal{X}_k \to \mathbb{R}$$

Sources of tension:

- the ℓ_i's may be misaligned
- agents may be uncoordinated, have partial observability of action, payoff, and information of others

[1] In game theory, the problem is to maximize the utility $u_i = -\ell_i$.

Motivating example: GANs

How to train a Deep Generative Model?

> Set up a two-player zero-game between a player tuning the parameters x of a deep neural network (called the 'generator') and a player tuning the parameters y of a deep neural network (called the 'discriminator').

$$f(x, y) = \mathop{\mathbb{E}}_{Z \sim \mu} \left[D_y(Z) \right] - \mathop{\mathbb{E}}_{Z \sim \mathcal{N}(0,I)} \left[D_y(G_x(Z)) \right]$$

Goodfellow et al. (2020)

Motivating example: GANs

- Objective function nonconvex.
- Inputs x and y high-dimensional.
- Existence of equilibria not guaranteed.

Simultaneous gradient descent

Suppose each agent simply runs gradient descent:

$$x_i^{(t+1)} \leftarrow x_i^{(t)} - \eta_i^{(t)} \nabla_{x_i} \ell_i(\mathbf{x}^{(t)}).$$

- This dynamics might not converge, let alone to something meaningful.
- Can exhibit cycling or divergent behaviors.
- **Note:** in two-player zero-sum games, where $\ell_1 = -\ell_0$, this dynamics is called gradient descent/ascent (GDA)

Example of cycles

Example (Simultaneous gradient descent can lead to cycles)

Consider the two-player zero-sum game with $f : \mathbb{R}^2 \to \mathbb{R}$,

$$f(x, y) = x^\top y.$$

- *Objective is convex/concave with a unique Nash equilibrium* $(0, 0)$.
- *GDA defines cyclic vector field* $(-\nabla_x f, \nabla_y f) = (-y, x)$.
- *Discretization of continuous dynamics leads to divergence:*

$$x \leftarrow x - \eta \nabla_x f(x, y)$$
$$y \leftarrow y + \eta \nabla_y f(x, y)$$

since $\|(x, y)\|$ *increases at each step.*

Example of cycles

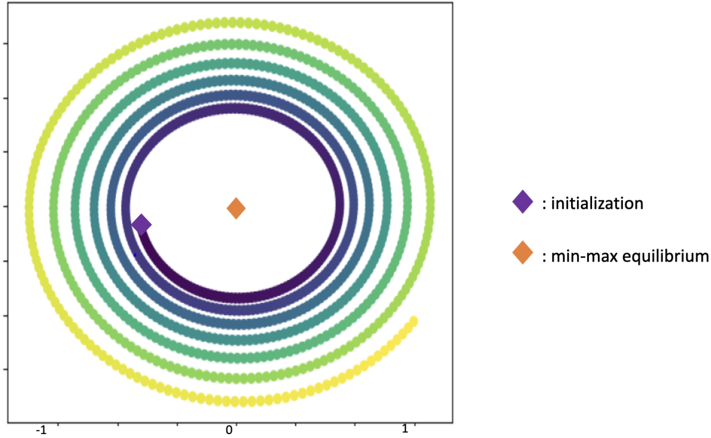

Figure 1: The vector field $(-\nabla_x f, \nabla_y f) = (-y, x)$ is tangent to the circle centered at $(0, 0)$. The continuous GDA dynamics cycles, while its discretization also expands outward.

Convergence on average to equilibrium

> "For this problem, it is known that Gradient Descent/Ascent (GDA) is a no-regret learning procedure, corresponding to follow-the-regularized leader (FTRL) with ℓ_2-regularization. As such, the average trajectory traveled by GDA converges to a min-max solution.
>
> *Daskalakis and Panageas (2018)*

Example of divergence

Example (Simultaneous gradient descent can diverge)

Define strongly-convex $f, g : \mathbb{R}^2 \to \mathbb{R}$ *by:*

$$f(x, y) = \frac{1}{2}(x + y + 1)^2 + \frac{1}{2}y^2 \qquad \textit{and} \qquad g(x, y) = \frac{1}{2}(x + y - 1)^2 + \frac{1}{2}x^2.$$

The vector field $\big(-\nabla_x f, -\nabla_y g\big)$ *has no stationary point, as* $\nabla_x f = \nabla_y g = 0$ *can never hold simultaneously:*

$$\nabla_x f(x, y) = x + y + 1 \qquad \textit{and} \qquad \nabla_y g(x, y) = x + y - 1.$$

Visualization of divergence

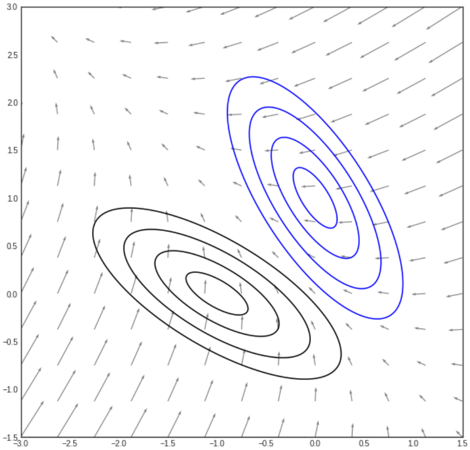

Figure 2: Here, f is black (*bottom*) and g is blue. The vector field corresponds to the dynamics $dx_t = -\eta_x \nabla_x f(x_t, y_t)$ and $dy_t = -\eta_x \nabla_y g(x_t, y_t)$.

Trouble in ‘easy’ setting

Training oscillation and garbage solutions arise even when:

- 2-player game, convex-concave, low-dimensional
- objective does not need to be learned (i.e. it is perfectly known)

Equilibrium finding in games

Question 1. Is simultaneous gradient descent (or some variant) able to arrive at the solution concepts offered by game theory?

- Nash equilibria, (coarse) correlated equilibria, regret minimization, etc.

Question 2. Are these equilibria meaningful? Do they exist? Efficiently computable?

Question 3. What is not captured by the study of equilibria?

Existence and tractability

We say the game is **convex** if each player i is minimizing a convex objective ℓ_i.

- Computing the Nash equilibrium is generally intractable (PPAD complete).
  - Computing approximate Brouwer fixed points is PPAD complete in 2+ dimensions Daskalakis et al. (2009), Chen and Deng (2009)

For nonconvex games, existence of equilibrium not guaranteed.

References

Xi Chen and Xiaotie Deng. On the complexity of 2d discrete fixed point problem. *Theoretical Computer Science*, 410(44):4448–4456, 2009.

Constantinos Daskalakis and Ioannis Panageas. Last-iterate convergence: Zero-sum games and constrained min-max optimization. *arXiv preprint arXiv:1807.04252*, 2018.

Constantinos Daskalakis, Paul W Goldberg, and Christos H Papadimitriou. The complexity of computing a nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009.

Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11): 139–144, 2020.