Gradient play in smooth games

Mazumdar et al. (2020)

Geelon So, agso@eng.ucsd.edu
Algorithmic game theory reading group — April 1, 2022

Game setting

- $\mathcal{I} = \{1, \ldots, n\}$ is a set of n **players**
- Each player controls their own **decision variable** $x_i \in X_i \subset \mathbb{R}^{m_i}$
- Each player aims to minimize their own **cost function** $f_i : \prod X_i \to \mathbb{R}$,

$$f_i(x) \equiv f_i(x_1, \ldots, x_n) \equiv f_i(x_i, x_{-i}).$$

  - Assume that each f_i is smooth.
- Let $\omega(x) = \big(\nabla_1 f_1(x), \ldots, \nabla_n f_n(x)\big)$ be the **player derivatives**.
- Let $D\omega(x)$, the Jacobian of ω, be the **game Hessian**,

$$D\omega(x) = \begin{bmatrix} D_1^2 f_1(x) & \cdots & D_{n1} f_1(x) \\ \vdots & \ddots & \vdots \\ D_{1n} f_n(x) & \cdots & D_n^2 f_n(x) \end{bmatrix}.$$

Game-theoretic solution concept: local Nash equilibria

Definition

A strategy $x = (x_1, \ldots, x_n)$ *is a* ***local Nash equilibrium*** *for the game* $(f_1, \ldots, f_n)$ *if for each player* $i \in \mathcal{I}$*, there is an open set* $W_i \subset X_i$ *containing* x_i *such that:*

$$f_i(x_i, x_{-i}) \leq f_i(w_i, x_{-i}) \qquad \forall w_i \in W_i \setminus \{x_i\}.$$

If all the inequalities as strict, then x *is a* ***strict local Nash equilibrium****.*

- In a neighborhood of x, no individual player is incentivized to unilaterally deviate.

Game-theoretic solution concept: differential Nash equilibrium

Definition

A strategy x is a ***differentiable Nash equilibrium*** *for the game $(f_1, \ldots, f_n)$ if for each i,*

- *the first derivative is zero, $\nabla_i f_i(x) = 0$*
- *the second derivative is positive definite $D_i^2 f_i(x) \succ 0$*

If we further have $\det\big(D\omega(x)\big) \neq 0$, then we say that it is ***non-degenerate****.*

Equivalence of local and non-degenerate differential NE

Theorem (Ratliff et al. (2014))

Local Nash equilibria are generically non-degenerate differential Nash equilibria.

- Games where LNE $\neq$ NDDNE form a measure zero set in the space of C^2-games.

Sufficient condition for isolated NE

Theorem (Ratliff et al. (2013))

Non-degenerate differential Nash equilibria are isolated strict local Nash equilibria.

Comparison with optimization

Optimization	**Game theory**
local minimum	local Nash equilibrium
$\nabla f(x) = 0$ and $\nabla^2 f \succ 0$	$\nabla_i f_i(x) = 0$ and $\nabla_i^2 f_i(x) \succ 0$
$\det\big(\nabla^2 f\big) \neq 0 \implies$ isolated	$\det\big(D\omega(x)\big) \neq 0 \implies$ isolated

Equilibrium computation

Question. How can differential Nash equilibria be computed?

- Often, gradient descent used in optimization.
- We consider simultaneous gradient descent for games.

Gradient play, or simultaneous gradient descent

Consider the following learning dynamic:

$$x_i^{(t+1)} \leftarrow x_i^{(t)} - \eta_i \cdot \nabla_i f\left(x_i^{(t)}, x_{-i}^{(t)}\right).$$

- Each player simply takes a step in the direction that decreases their cost the fastest.
- The learning dynamic is a *discretization* of the continuous dynamics:

$$\dot{x} = -\omega(x).$$

Dynamical system solution concept

Definition

A point $x \in X$ *is a* ***locally asymptotically stable equilibrium*** *of the continuous-time dynamics* $\dot{x} = -\omega(x)$ *if* $\omega(x) = 0$ *and* $\mathrm{Re}(\lambda) > 0$ *for all eigenvalues* λ *of* $D\omega(x)$.

Properties of locally asymptotically stable equilibria

Let x be a locally asymptotically stable equilibrium. Then:

1. x is isolated (i.e. there is a neighborhood in which no other equilibrium exists)
2. the dynamics $\dot{x} = -\omega(x)$ is locally exponentially attracting
3. an appropriately discretized dynamics converges at rate $O(1/t)$

Dynamical system solution concept

In the following, let λ_i be eigenvalues of $D\omega(x)$ ordered such that:

$$\mathrm{Re}(\lambda_1) \leq \mathrm{Re}(\lambda_2) \leq \cdots$$

Definition

*A point $x \in X$ is a **saddle point** of $\dot{x} = -\omega(x)$ if $\omega(x) = 0$ and $\mathrm{Re}(\lambda_1) \leq 0$. A saddle point is **strict** if the real part of all eigenvalues are nonzero, and there is also some $\mathrm{Re}(\lambda_i) > 0$.*

Relationship between dynamical and game-theoretic solutions?

Question. Are the dynamical equilibria related to game-theoretic equilibria?

- Analogous questions in optimization:
  - Does gradient descent only converge to local minima? (Yes, Lee et al. (2016))
  - How fast is convergence?
  - What happens if there is noise (i.e. stochastic gradient descent)?

Result 1: convergence to non-game-theoretic solution

Gradient play can converge to irrelevant (game theoretic) solution

Proposition

In the class of general-sum continuous games, there exists a continuum of games $\mathcal{G}$ *such that* $\mathrm{LASE}(\omega) \not\subset \mathrm{NDDNE}(\mathcal{G}.$ *Moreover,* $\mathrm{LASE}(\omega) \not\subset \mathrm{local\ NE}(\mathcal{G}.$

Proof: construct a counterexample

Consider the two-player game:

$$f_1(x_1, x_2) = \frac{1}{2}ax_1^2 + bx_1x_2 \qquad \text{and} \qquad f_2(x_1, x_2) = \frac{1}{2}dx_2^2 + cx_1x_2.$$

The game Hessian is given by:

$$D\omega(x) = \begin{bmatrix} a & b \\ c & d \end{bmatrix}.$$

- If $a > 0$ and $d < 0$, there is a unique stationary point $x = (0, 0)$.
  - It is not a differential or local Nash equilibrium, since $d < 0$.
  - If $a > -d$ and $ad > cb$, then eigenvalues of $D\omega$ have positive real parts; x is a LASE.

□

Result 2: non-convergence to some local Nash equilibria

NDDNEs are LASEs or SSPs

Proposition

A non-degenerate differential Nash equilibrium is either a locally asymptotic stable equilibrium or a strict saddle point of $\dot{x} = -\omega(x)$,

$$\text{NDDNE}(\mathcal{G}) \subset \text{LASE}(\mathcal{G}) \cup \text{SSP}(\mathcal{G}).$$

- LASE: all eigenvalues λ satisfy $\text{Re}(\lambda) > 0$.
- SSP: $\text{Re}(\lambda) \neq 0$, and there are $\text{Re}(\lambda_j) < 0$ and $\text{Re}(\lambda_j) > 0$.

Note: this implies that a NDDNE is not strictly unstable or strictly marginally stable (eigenvalues are imaginary).

Proof that $\mathrm{NDDNE}(\mathcal{G}) \subset \mathrm{LASE}(\mathcal{G}) \cup \mathrm{SSP}(\mathcal{G})$

Let $x \in \mathrm{NDDNE}(\mathcal{G})$ be a non-degenerate differential Nash equilibrium. Need to show that (i) some eigenvalue has $\mathrm{Re}(\lambda) > 0$ and that (ii) no eigenvalue has $\mathrm{Re}(\lambda) = 0$.

(i) Claim: $\sum \lambda_i = \mathrm{tr}\big(D\omega(x)\big) > 0$.

- note that $\mathrm{tr}\big(D\omega(x)\big) = \sum \mathrm{tr}\big(\nabla_i^2 f_i(x)\big)$
- differential Nash equilibrium condition $\nabla_i^2 f_i(x) \succ 0$ implies $\mathrm{tr}\big(\nabla_i^2 f_i(x)\big) > 0$

(ii) Claim: $\mathrm{Re}(\lambda_i) \neq 0$ for all i.

- note that $\det\big(D\omega(x)\big) = \prod \lambda_i$
- non-degenerateness condition $\det\big(D\omega(x)\big) \neq 0$ implies $\mathrm{Re}(\lambda_i) \neq 0$ [why?]

□

Characterization of real and imaginary eigenvalues

5 The main difference between imaginary and real eigenvalues is that imaginary eigenvalues are imaginary, whereas real eigenvalues are real. – Gerry Myerson Jul 5, 2016 at 13:18

An example of a SSP

Consider the two-player game from before:

$$f_1(x_1, x_2) = \frac{1}{2}ax_1^2 + bx_1x_2 \qquad \text{and} \qquad f_2(x_1, x_2) = \frac{1}{2}dx_2^2 + cx_1x_2.$$

The game Hessian is given by:

$$D\omega(x) = \begin{bmatrix} a & b \\ c & d \end{bmatrix}.$$

- We obtain a strict saddle point if $a, d > 0$ but $\det\big(D\omega(x)\big) < 0$.

Non-convergence result overview

Let player $i \in \mathcal{I}$ have learning rate γ_i. Put $\gamma = (\gamma_1, \ldots, \gamma_n)$. Consider the discretization:

$$x_{t+1} = g(x_t),$$

where $g(x_t) = x_t - \gamma \odot \omega(x)$ where $\odot$ is element-wise multiplication.

- We saw that some local Nash equilibria are strict saddle points.
- Simultaneous gradient descent avoids strict saddle points almost surely.
  - Makes use of the **stable manifold theorem**.

Summary: gradient play will almost surely avoid certain local NE, while it may converge to dynamical equilibria that are game-theoretically irrelevant.

Stable manifold theorem: linear example

Consider the following discrete-time linear dynamical system on $\mathbb{R}^4$ where:

$$x_{t+1} = \phi(x_t) \qquad \text{and} \qquad \phi = \begin{bmatrix} 0.5 & & & \\ & 0 & 1 & \\ & 1 & 0 & \\ & & & 2 \end{bmatrix}.$$

- Decompose $\mathbb{R}^4 = E_1 \oplus E_{2,3} \oplus E_4$ into its generalized eigenspaces.
  - Points $x = (x_1, 0, 0, 0) \in E_1$ contract to 0 exponentially quickly
  - Points $x = (0, x_2, x_3, 0) \in E_{2,3}$ have limit cycles
  - Points $x = (0, 0, 0, x_4) \in E_4$ diverge to ∞ exponentially quickly
- If $\phi^t(x)$ remains in a small ball around 0 for all $t > 0$, then x must be in $E_1 \oplus E_{2,3}$.

Stable manifold theorem

Let $\phi : \mathbb{R}^d \to \mathbb{R}^d$ be smooth and let x_0 be a fixed point. Assume that ϕ is a local diffeomorphism at x_0, so that ϕ is smoothly invertible around x_0.

Theorem (Stable manifold theorem)

Let ϕ and x_0 as above. Then the tangent space at x_0 has decomposition

$$\mathbb{R}^d = E^{\text{stable}} \oplus E^{\text{center}} \oplus E^{\text{unstable}}$$

into $D\phi(x_0)$-invariant generalized eigenspaces corresponding to eigenvalues $|\lambda|$ less than 1, equal to 1, and greater than 1. There is an invariant disc $W \subset E^{\text{stable}} \oplus E^{\text{center}}$ called the ***local stable center manifold*** *and a ball B around x_0 such that:*

- $\dim_{\text{manifold}}(W) = \dim(E^{\text{stable}} \oplus E^{\text{center}})$ *and* $\phi(W) \cap B \subset W$,
- *if* $\phi^t(x) \in B$ *for all* $t \geq 0$, *then* $x \in W$.

Non-convergence to SSPs

Assume the following:

- $f_i : X \to \mathbb{R}$ is smooth and $\|D\omega(x)\|_2 \leq L$ for all x.
- $\gamma_i \in (0, 1/L)$ and $\gamma = (\gamma_1, \ldots, \gamma_n)$.
- $X = \prod X_i$ is open and convex.
- Let $g(x) = x - \gamma \odot \omega(x)$ and $Dg(x) = I - \Gamma D\omega(x)$ where $\Gamma = \mathrm{diag}(\gamma_1, \ldots, \gamma_n)$.

Theorem

If $g(X) \subset X$, then the set of initial conditions from which simultaneous gradient descent converges to a strict saddle point has measure zero.

Proof of non-convergence to SSPs

Let $x^* \in X$ be a SSP and let W be its local stable center manifold. Note that W has measure zero since it has positive codimension is $\dim(E^{\text{unstable}})$: since x^* is a strict saddle point, $I - \gamma D\omega(x^*)$ as an eigenvalue greater than 1.

- Assume that g is a diffeomorphism for now.
- If $g^t(x_0)$ converges to x^*, the stable manifold theorem implies:

$$x_0 \in \bigcup_{t=0}^{\infty} g^{-t}(W).$$

- Thus, the set of initial conditions converging to x^* has measure zero since it is a countable union of measure zero sets.
  - Showing that the set of initial conditions converging to any SSP has measure zero is slightly complicated if there are uncountably many SSPs [is this possible in $\mathbb{R}^d$?]. However, the fix is to find countable a cover.

Proof of non-convergence to SSPs: g is a diffeomorphism

(i) Claim: g is invertible.
- If $g(x) = g(y)$, then $y - x = \Gamma\big(\omega(y) - \omega(x)\big)$.
- However, ω is L-Lipschitz while $\|\gamma\|_{\text{op}} \leq \max|\gamma_i| < 1/L$.
- Together, $\|y - x\| < 1/L \cdot L\|x - y\|$, a contradiction.

(ii) Claim: g is a local diffeomorphism.
- Show that Dg is invertible then appeal to implicit function theorem.
- Since $Dg = I - \Gamma D\omega$, suffices to show that 1 is not an eigenvalue of $\Gamma D\omega$:

$$\|\Gamma D\omega\|_{\text{op}} \leq \underbrace{\max|\gamma_i|}_{<1/L} \cdot \underbrace{\|D\omega\|_{\text{op}}}_{\leq L} < 1.$$

Together: (i) shows invertibility while (ii) shows inverse is locally smooth. Thus g is smooth and has smooth inverse.

□

Other limiting behavior

Because $\dot{x} = -\omega(x)$ is not a gradient flow, it can exhibit more complicated behaviors including **limit cycles** and **chaos**.

- Mazumdar et al. (2020) also show avoidance of linearly unstable limit cycles.

Stochastic approximation

Stochastic gradient play

Assume that each player updates with the rule:

$$x_i^{(t+1)} \leftarrow x_i^{(t)} - \gamma_i^{(t)} \cdot \left(\nabla_i f_i(x^{(t)}) + Z_i^{(t+1)}\right),$$

where $Z_i^{(t+1)}$ is mean-zero noise with finite variance.

Stochastic gradient play does not converge to SSPs

Theorem

Consider a game $(f_1, \ldots, f_n)$ on $X = \mathbb{R}^m$. Suppose each player applies the stochastic gradient play update each step, with uniform learning rates $\gamma_i^{(t)} = \gamma^{(t)}$ satisfying $\sum \left(\gamma^{(t)}\right)^2$. Further suppose that there exists some $b > 0$ such that for all unit vectors v,

$$\mathbb{E}[(Z_i^{(t)} \cdot v)_+ \mid \mathcal{F}_t] \geq b.$$

Then, the iterates do not converge to strict saddle points almost surely.

- The condition $\mathbb{E}[(Z_i^{(t)} \cdot v)_+ \mid \mathcal{F}_t] \geq b$ prevents the case where "noise forces the stochastic dynamics onto the stable manifold".

Summary

- We can analyze the behavior of gradient-based learning in games through a dynamical systems perspective.
- The stability depends on the game Hessian $D^2 f = D\omega$.
- Gradient play can converge to irrelevant game-theoretic solutions.
- Gradient play can almost surely avoid local Nash equilibria.
- Implications for gradient-based learning in multi-agent reinforcement learning, multi-armed bandits, generative adversarial networks, online optimization?

References

Jason D Lee, Max Simchowitz, Michael I Jordan, and Benjamin Recht. Gradient descent converges to minimizers. *arXiv preprint arXiv:1602.04915*, 2016.

Eric Mazumdar, Lillian J Ratliff, and S Shankar Sastry. On gradient-based learning in continuous games. *SIAM Journal on Mathematics of Data Science*, 2(1):103–131, 2020.

Lillian J Ratliff, Samuel A Burden, and S Shankar Sastry. Characterization and computation of local nash equilibria in continuous games. In *2013 51st Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 917–924. IEEE, 2013.

Lillian J Ratliff, Samuel A Burden, and S Shankar Sastry. Genericity and structural stability of non-degenerate differential nash equilibria. In *2014 American Control Conference*, pages 3990–3995. IEEE, 2014.