Solving Neural Min-Max Games: The Role of Architecture, Initialization & Dynamics

Deep Patel and Emmanouil-Vasileios Vlatakis-Gkaragkounis

Department of Computer Science, University of Wisconsin-Madison, WI, USA

Rise of Multi-Agent Learning Applications

- Many emerging applications — such as adversarial training, AI alignment, and robust optimization — can be framed as zero-sum games between neural nets with Nash equilibria (NE) capturing the desirable system behavior.
- Much of the remarkable progress stems from the capacity of deep networks to operate effectively in environments with large, often continuous, state and action space (e.g. Go, autonomous driving, Texas Hold'em Poker, and StarCraft II).

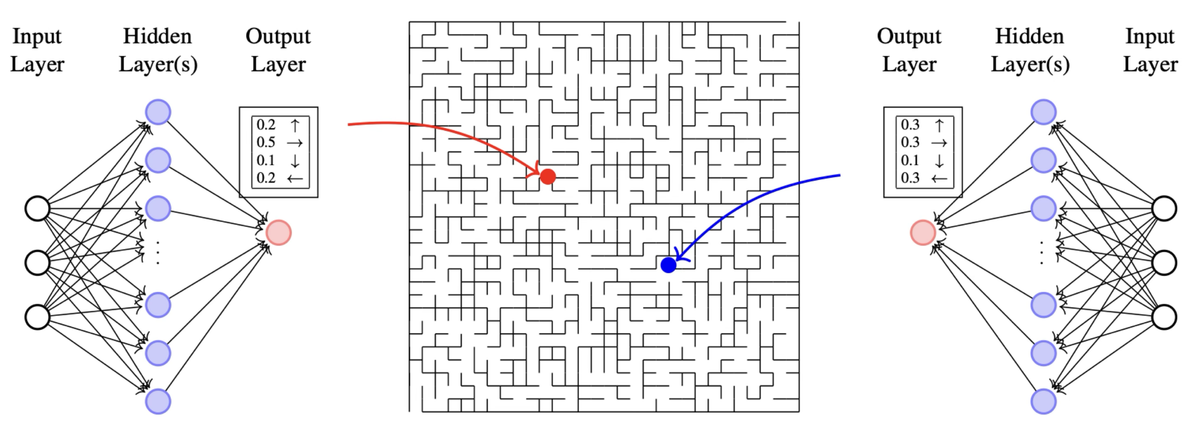

Figure 1: *Illustration of a maze environment where each agent must reason over a vast space of action sequences. Instead of explicitly constructing and searching the full decision tree, a neural network implicitly encodes both the value of paths and the policy for navigation, learning an effective strategy dynamically without ever uncovering the complete structure of the maze.*

Question At The Heart Of Our Work

How can two neural networks be designed and trained to compute a solution to a zero-sum game?

More Specifically...

How many parameters should the two neural networks have so that vanilla methods like AltGDA can converge to a saddle point?

From `MIN` to `MIN-MAX`

- Prior works primarily focus on degree of overparameterization needed for gradient descent to reach global minima (classical `MIN` setting) [1].
- Moreover, making Gradient Descent (GD) analogs 'work' for `MIN-MAX` problems is hard due to 'cycling' behaviour.
- Hence, focus on a class of min-max games that capture as many of the current deep learning applications as possible and yet avoid cycles.
- ***Hidden Games!***

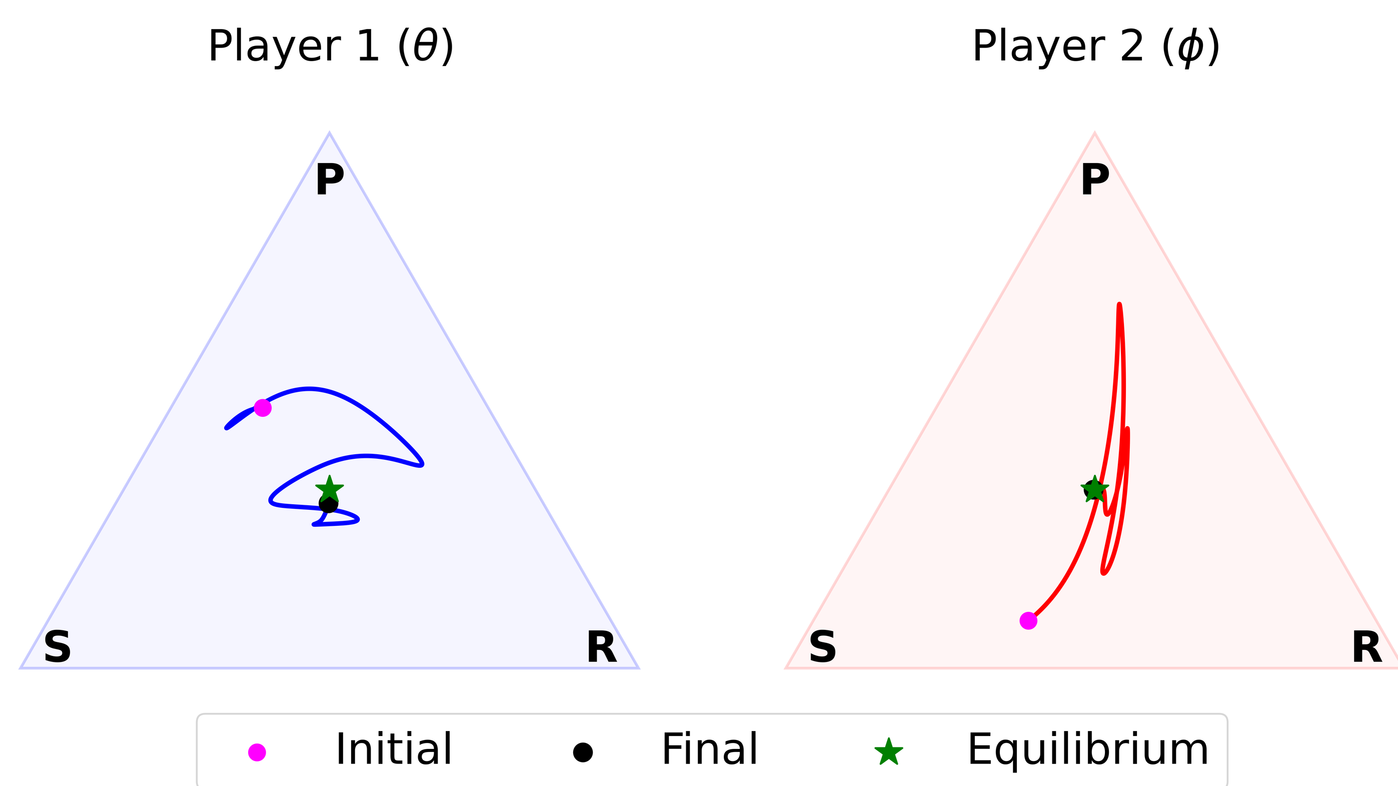

Figure 2: *An AltGDA trajectory of players' latent space strategies in an ℓ_2-regularized hidden game of Rock-Paper-Scissors.*

Hidden Games

$$(\theta^\star, \phi^\star) \in \arg\min_{\theta\in\mathbb{R}^m} \arg\max_{\phi\in\mathbb{R}^n} \mathbb{E}_{(x,x')\sim P_{xx'}}\left[L(F(x;\theta), G(x';\phi))\right]$$

- The `MIN-MAX` problem above is called a hidden-convex-hidden-concave game if the loss L is convex (concave) in $F(\cdot;\theta)$ ($G(\cdot;\phi)$).
- Note that, in general, the loss L is non-convex (non-concave) in θ (ϕ).
- Captures various applications such as GANs, Parametric Distributionally Robust Optimization, Robust Reinforcement Learning, Domain Invariant Representation Learning (DIRL).
- Some known convergence results when the `MIN-MAX` objective is regularized [2] or satisfies two-sided Polyak-Łojasiewicz (PŁ) condition [3].

Hidden Convex Optimization

$$\min_{\theta\in\Theta} F(\theta) := H(c(\theta)) \qquad (\bigstar)$$

- This problem is (μ_c, μ_H)-hidden convex if the function H is μ_H -strongly-convex ($\mu_H \geq 0$) and the map c is invertible with the inverse mapping c^{-1} as $1/\mu_c$-Lipschitz.
- The minimum singular value of the Jacobian for mapping $c(\theta)$ ($\sigma^2_{\min}(\mathbf{J}(\theta))$) in Equation ($\bigstar$) above corresponds to the constant μ_c.

Fact. [4] (μ_c, μ_H)-hidden-strongly-convex function satisfies PŁ-condition with modulus $\mu_H\mu_c^2$. (Recall: A function f is said to satisfy μ-PŁ-condition if it satisfies the following: $f(x) - f(x^\star) \leq \frac{1}{2\mu}\|\nabla f(x)\|_2^2$ where $x^\star$ is the global minimizer of f.)

Ensuring Hidden (Strong) Convexity and Saddle-Point Convergence

- Ensure $\sigma^2_{\min}(\mathbf{J}(\theta_t)) > 0 \ \ \forall t \in \{0, \ldots, T\} \implies$ objective satisfies PŁ-condition $\forall t \leq T$.
- Adapt analysis of AltGDA [3] to prove saddle-point convergence.

Our Results: Input Games

$$\min_{x_{Alice}\in\mathcal{D}_F} \max_{x_{Bob}\in\mathcal{D}_G} L(F(x_{Alice};\theta), G(x_{Bob};\phi)) \qquad (\bullet)$$

- Parameters θ, ϕ are fixed.
- Optimizing over inputs.
- Example: adversarial example generation.

§**Theorem 1.** For bilinear objectives (payoff matrix A) with $\epsilon-$ ℓ_2-regularization, w.h.p. AltGDA converges to ϵ-saddle point if the Gaussian-randomly-initialized mappings F and G (1-hidden-layer neural networks) satisfy

$$\sigma^2_{F/G} = \tilde{\Theta}\left(\frac{poly(1/width_{F/G})}{\sigma_{\max}(A)}\right)$$

Our Results: Neural Games

$$\min_{\theta\in\mathbb{R}^m} \max_{\phi\in\mathbb{R}^n} \mathbb{E}_{(x,x')\sim P_{xx'}}\left[L(F(x;\theta), G(x';\phi))\right] \qquad (\blacksquare)$$

- Optimize over parameters θ, ϕ for given data.
- GANs, DIRL, etc.

§**Theorem 2.** For separable hidden-strongly-convex-strongly-concave min-max objectives with bilinear coupling, w.h.p. AltGDA converges to a saddle point if the Gaussian-random initializations and hidden-layer width of the networks F and G satisfy

$$\sigma_{1,F/G} \cdot \sigma_{2,F/G} \lesssim \frac{1}{\sqrt{d_{in,F/G} \cdot \text{width}_{F/G}}}$$

$$\text{width}_{F/G} = \tilde{\Omega}\left(\mu^2_{\theta/\phi}\frac{n^3}{d_{in,F/G}}\right)$$

Proof Outline

- Choose Gaussian random initializations (θ_0, ϕ_0) such that the Jacobian for networks F and G is 'well-conditioned' w.h.p.
- Define radius R of a Euclidean ball $\mathcal{B}((\theta_0, \phi_0), R)$ such that the Jacobian remains well-conditioned within it.
- Compute path length bound of AltGDA iterates (θ_t, ϕ_t) in terms of P_0, a special Lyapunov potential P_0 at time $t = 0$.
- Find sufficient conditions on hidden layer width of networks F, G to ensure this path length is smaller than the ball radius R.

Conclusion

- To our knowledge, first overparameterization condition (sufficient condition) for saddle-point convergence in a special class of games – Separable Hidden-Strongly-Convex-Strongly-Concave with Bilinear Coupling.
- Results hold for shallow neural networks with differentiable activation functions (e.g. GeLU).

Future Work

- The width (and hence the overparameterization) condition on the shallow neural networks is a *sufficient* condition. Is it also *necessary*?
- Analysis assumes differentiable activation functions. Extend to non-differentiable activation functions (e.g. ReLU).
- Connect results with those for extensive-form games.
- Extend to Hidden MVIs for polyhedral settings.

References

[1] Yongtao Wu, Fanghui Liu, Grigorios Chrysos, and Volkan Cevher.
On the convergence of encoder-only shallow transformers.
Advances in Neural Information Processing Systems, 36:52197–52237, 2023.

[2] Fivos Kalogiannis, Emmanouil-Vasileios Vlatakis-Gkaragkounis, Ian Gemp, and Georgios Piliouras.
Solving zero-sum convex markov games.
In *Forty-second International Conference on Machine Learning*, 2025.

[3] Junchi Yang, Negar Kiyavash, and Niao He.
Global convergence and variance reduction for a class of nonconvex-nonconcave minimax problems.
Advances in neural information processing systems, 33:1153–1165, 2020.

[4] Ilyas Fatkhullin, Niao He, and Yifan Hu.
Stochastic optimization under hidden convexity.
arXiv preprint arXiv:2401.00108, 2023.