On Last-Iterate Convergence Beyond Zero-Sum Games

Ioannis Anagnostides, Ioannis Panageas, Gabriele Farina, Tuomas Sandholm

Carnegie Mellon University, University of California Irvine

Main Focus: No-Regret Learning Dynamics in Games

- What if learning agents play against each other in games?
- Traditional no-regret learning fails to converge to a **Nash equilibrium**
- Recent progress using simple variants **(optimistic mirror descent)**

Our Questions

- How fast is the convergence?
- What if players are using different updates rules?
- Convergence beyond zero-sum games?
- Guarantees in terms of the social welfare?

New Regret-Based Framework for Last-Iterate Convergence

- So far, the analysis of last-iterate convergence differs from the analysis of the players' **regrets**
- We show last-iterate guarantees using regret bounds
- We inherit the robustness stemming from a regret-based analysis

Implications

- Tight $O(1/\sqrt{T})$ rates for games with **nonnegative sum of regrets**
- Captures MVI property, and the weak MVI property
- Last-iterate convergence implies optimal $O(1)$ regret for each player
- Our guarantees apply even if players use different learning rules
- **Beyond zero-sum**: Games that satisfy an approximate minimax theorem

No-Regret Learning Can Outperform the Price of Anarchy

- Variants of optimistic mirror descent **outperform** the (robust) price of anarchy when they do not converge to Nash equilibria
- The price of anarchy corresponds to the social welfare of the worst Nash equilibrium in the game
- Cycling behavior can improve efficiency in games

Questions?

ioannis.anagnostides@gmail.com