The Stability of Human/Machine Learning Dynamics

Benjamin J. Chasnov, Lillian J. Ratliff, Samuel A. Burden
Electrical & and Computer Engineering
University of Washington
Seattle, WA, USA
{bchasnov,ratlifflj,sburden} @uw.edu

I. Learning as Optimization

We model learning as a continuous optimization process. When agents learn over time, they seek decision variables that incrementally improve their performance on a task. For tasks with only one agent, strong guarantees on the learning process can be provided: the stable steady-state outcome of learning is guaranteed to be locally optimal. However, with multiple agents where the objective of an agent depends on the decision variables of others, similar guarantees do not hold. The steady-state equilibria of learning processes with multiple agents do not necessarily correspond to the optimal actions for each agent individually [1].

The rich behavior of multi-agent interactions can result in counter-intuitive outcomes that do not correspond to meaningful equilibria of the original objectives. Agents can converge to non-optimal suprious attractors or get stuck in endless periodic orbits. We study this interaction theoretically, predict outcomes, and evaluate the predictions experimentally in interactive and adaptive dynamics simulators.

II. Iteractions as Games

We present a theoretical model that captures the rich dynamics of learning in multi-agent settings. The model makes predictions about real-world interactions. We present our model below, along with modifications to optimization problems that enable agents to express different learning rules while respecting their original objectives.

In an $n-$player continuous game defined by the objectives $(f_1, \ldots, f_n)$, the agent indexed by i considers two objects: an individual objective f_i, which encodes its performance as a scalar value, and a decision variable x_i, which represents its action taken from its feasible set X_i. An agent assesses its performance at time t by evaluating its objective at the current joint decision $x(t) = (x_i(t), x_{-i}(t))$ where $x_{-i}(t)$ is all other agents' decision variables in X_{-i}. At time t, an agent achieves the performance $f_i(x_i(t), x_{-i}(t)) \in \mathbb{R}$. The steady-state equilibrium $x^* = (x_i^*, x_{-i}^*)$ and learning dynamics are diagrammed in Figure 1(a),(c).

If an agent identifies a feasible decision with better performance in the neighbourhood of the current joint decision, then it will adjust its own decision in that direction. Mathematically, the agents continuously update their decisions along the steepest gradient of their objectives, represented by the coupled differential equations $\dot{x}_i(t) = D_{x_i} f_i(x_i(t), x_{-i}(t))$ for all maximizing agents $i \in [1, n]$.

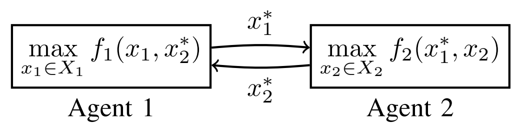

(a) A two-player Nash equilibrium.

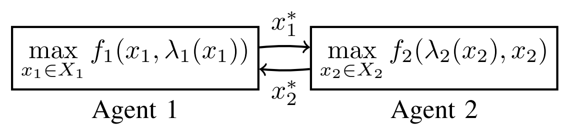

(b) A two-player conjectural equilibrium.

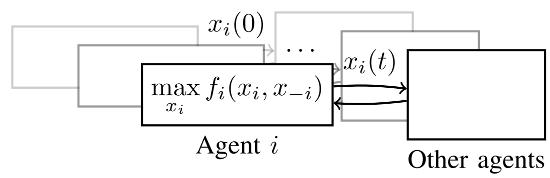

(c) The learning dynamics of agent i.

Fig. 1: The steady-state equilibria of (a) Nash play and (b) conjectural play are the joint decision variables (x_1^*, x_2^*) that satisfy each agent's optimization problem while keeping the other's decsion variable fixed. The learning dynamics (c) models an agent starting at a suboptimal decision $x_i(0)$ at time $t = 0$ and ending with decision $x_i(t)$. Do the learning dynamics lead to meaningful equilibria?

III. Experiments with Opponent Models

The behavioral model we present is verified experimentally. We also test various conjectural models, where agent i is given the ability to anticipate the responses of others using a conjecture $\lambda_i : X_i \to X_{-i}$, see Figure 1(b) and [2]. We demonstrate that conjectures consistent with the game or estimated using data can stabilize the learning dynamics, avoid spurious attractors and achieve better performance when compared to Nash. Our results provide a promising step towards a theory that can be used to ensure safe and robust interactions amongst humans and machines.

References

[1] Fudenberg, D., & Levine, D. K. (1998). "The theory of learning in games." MIT press.
[2] Figuires, Charles. (2004). "Theory of conjectural variations." Vol. 2. World Scientific.