Learning Simulation-Based Games from Data

Enrique Areyan Viqueira[1,2], Cyrus Cousins[1], Amy Greenwald[1,2] & Eli Upfal[1]

1. Brown University
Providence, Rhode Island, USA

2. NEC-AIST AI Cooperative Research Laboratory
Tokyo, Japan

Overview

- A **game** among **strategic** agents consists of **actions** and **payoffs** to players depending on players joint actions.
- A fundamental concern is to **predict** the outcome in a game, i.e., which actions rational players will choose. We consider **Nash equilibria** as our prediction.
- Our research focus is on **learning** games from **noisy** observations of the game's payoffs.
- Our main results (1) **bounds on all** Nash equilibria of a learned game, (2) algorithms to **efficiently** learn games.

Approximating Games

- A game G is **compatible** to a game G' if they have the same players and actions but possibly different payoffs.
- Compatible games G and G' **are very close** if their payoffs don't differ by much, say ε > 0. Example:

Game G	Strategy C1	Strategy C2
Strategy R1	1, 3	2+ε, -3
Strategy R2	3, 4-ε	0, 5

Game G'	Strategy C1	Strategy C2
Strategy R1	1+ε, 3	2, -3
Strategy R2	3, 4	0, 5-ε

Main Result on Approximation

$Nash_\varepsilon(G)$ is the set of ε-Nash of Game G. We show:

$$Nash(G) \subseteq Nash_{2\varepsilon}(G') \subseteq Nash_{4\varepsilon}(G)$$

- We think of G as our **ground-truth game** and G' as an **empirical game** built from observational data.
- First containment shows **perfect recall**, all ground-truth Nash are in the approximation.
- Second containment shows **approximately perfect precision**, all Nash of the approximation are close to Nash in the ground truth

Learning Games and Empirical Results

- We assume access to a **simulator** capable of producing any number of samples for any possible payoff (any cell in the game's matrix)
- We propose and evaluate two **PAC-Learner** algorithms to learn empirical games. Global and Progressive Sampling.
- **Progressive Sampling** is a novel algorithm that samples dynamically, saving on samples where fewer data are necessary to confidently learn.
- We empirically demonstrate pruning **substantially saves** on sampling.
(results to the right. Feel **free** to ask **me** any question! :-).

A new **methodology** to **learn all** equilibria of games from **data**.

Experimental Results

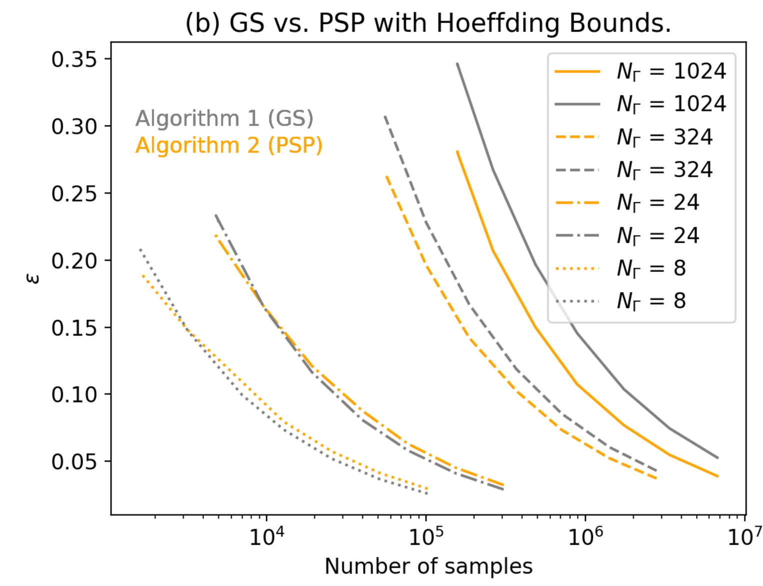

- **Pruning** significantly **reduces** the number of samples required to achieve a desired **accuracy** as compared to global sampling.

More on Statistical Bounds

- **Hoeffding's inequality**. Given the desired error accuracy ε>0, and the desired failure probability δ > 0, Hoeffding's inequality provides the number of samples needed to obtain an empirical game G that is ε close to G' with probability 1 – δ. The number of samples is a function of the size of the game (#players and #strategy profiles).
- **Rademacher Complexity**. An alternative to Hoeffding's inequality that is independent of the size of the game but depends on the sampled data. More research on Rademacher complexity for learning games is current ongoing work.

Applications

We can learn all equilibria. Great! …
Wait, why do we care?

Empirical Mechanism Design

- **Mechanism Design:** the science of designing the rules of a game (system) such that the strategic interaction among participants leads to desirable outcomes.
- **Parametric Mechanism Design**: the mechanism designer can optimize parameters of the system, e.g., reserve prices in auctions.
- **Assuming**: for every parameter of the system, participants play a set of known actions and their interaction leads to an equilibrium (or close).
- **The punch line**: our methodology allows control for any possible equilibria that might be played, allowing the designer to optimize with confidence.
- **Example application**: electronic advertisement exchange systems such as Google AdWords©, Amazon Sponsored Brands and Sponsored Products©, etc.

BROWN

AIST
NATIONAL INSTITUTE OF ADVANCED INDUSTRIAL SCIENCE AND TECHNOLOGY (AIST)