

Deterministic Calibration with Simple Rules

Dean Foster

University of Pennsylvania

and

Sham Kakade

University of Pennsylvania

January 31, 2005

The problem: Learning Nash equilibria

Current methods are slow and involve exhaustive search

Can a fast method be found?

How about for special form games?

Measuring complexity

Two definitions of speed of convergence:

- total CPU used
- number of rounds of play

History

	Forecast probability	Forecast
Blackwell	CE Calibration (F. and Vohra, '97)	CE No regret (F. and Vohra '97) (Hart and Mascherian '96)
Exhaustive search	NE Hypothesis testing (F. and Young '03)	NE Regret (F. and Young '03) (Germano & Vohra '04)
Public methods	NE Weak calibration yesterday's talk (Kakade and F. '04)	NE Weak utility today's talk (Kakade and F. '04)

	Forecast probability	Forecast
Blackwell (\rightarrow CE)	$(1/\epsilon)^{a^n}$	$(a/\epsilon)^n$
Exhaustive search (\rightarrow Nash)	$\gg (1/\epsilon)^{a^n}$	$\gg (1/\epsilon)^{a^n}$
Public methods (\rightarrow Nash)	$(1/\epsilon)^{a^n}$ $2^{ \mathcal{I} }$	$(1/\epsilon)^{a^n}$ $ \mathcal{I} ^{\log \log \mathcal{I} }$ (with)

n = number of players

a = number of actions per player

ϵ = desired accuracy

$|\mathcal{I}| = a^n$ = input size (a is fixed)

(CE: Blackwell gives fast approx algo. NE: slow, few results known.)

- X_t sequence to be forecast by p_t
- Weak calibration, means

$$\sum_{t=1}^T (X_t - p_t) w(p_t) \rightarrow 0$$

- $w()$ is any smooth function.
- What Sham talked about yesterday.

- Today's twist: Use other testing functions. Eg

$$\sum_{t=1}^T (X_t - p_t) w(p_t, X_{t-1}) \rightarrow 0$$

Would test for Markov patterns.

- Game setting for calibration
 - $X_{i,t}$ is the observable that player i cares about
 - $p_{i,t}$ is a forecast of $X_{i,t}$

- Individual calibration:

$$(\forall i) \quad \sum_{t=1}^T (X_{i,t} - p_{i,t}) w(p_{i,t}) \rightarrow 0$$

- Public calibration:

$$(\forall i) \quad \sum_{t=1}^T (X_{i,t} - p_{i,t}) w(\vec{p}_t) \rightarrow 0$$

The game model

- Player i uses $p_{i,t}$ to predict the round t
- Player i then use smooth decision rule $s_i(p_{i,t})$ to probability of their play in round t .
- Player i then randomly action S_i from this distrib

- Game setup:
 - Take $X_i = S_{-i}$ (i.e. all actions but player i)
 - $p_{i,t}$ is forecast of $X_{i,t}$

- Individual calibration:

$$(\forall i) \quad \sum_{t=1}^T (X_{i,t} - p_{i,t}) w(p_{i,t}) \rightarrow 0$$

- Public calibration:

$$(\forall i) \quad \sum_{t=1}^T (X_{i,t} - p_{i,t}) w(\vec{p}_t) \rightarrow 0$$

- Suppose players play a smooth best reply to forecast
 - Traditional calibration \rightarrow correlated equilibria
 - Public calibration \rightarrow Nash equilibria
- Speed of convergence is related to dimension of the “space” of the testing functions
 - For individual: dimension $(1/\epsilon)^{a^n}$
 - For public: dimension is $(1/\epsilon)^{na^n}$
 - Hence convergence is slow in both cases.
- Need lower dimensional space, but what can be c

- Truth \approx prediction
 - via calibration
- Truth is independent
 - Given \vec{p} each player is in fact playing independent
- ϵ -rationality
 - ϵ -BR to prediction
 - p_i includes information about what all other pl
- Independence + ϵ -rationality = ϵ -NE.

What can be changed?

- Take $X_{i,t}$ to be the vector of potential payoffs
 - \vec{S}_{-i} is the vector of everyone else's play
 - $u_{i,t}(k) = u_i(k, \vec{S}_{-i,t})$
 - $X_{i,t} = (u_{i,t}(1), \dots, u_{i,t}(a))$

- Utility model

- $p_{i,t}$ is an estimate of $X_{i,t}$ made at time $t - 1$
- For CE we need

$$(\forall i) \quad \sum_{t=1}^T (X_{i,t} - p_{i,t}) w(p_{i,t}) \rightarrow 0$$

- For NE we need

$$(\forall i) \quad \sum_{t=1}^T (X_{i,t} - p_{i,t}) w(\vec{p}_t) \rightarrow 0$$

- For CE: number of rounds is $O((n/\epsilon)^a)$
- For NE: number of rounds is $O((n/\epsilon)^{an})$
- Looks almost polynomial in length of input
 - $|I| = a^n = \text{input size}$ (a is fixed)
 - number of rounds is $O(|\mathcal{I}|^{\log \log |\mathcal{I}|})$
 - “pseudo Poly”.
- Although exp in a , little known computationally.

Graphical Models for Game Theory

- Undirected graph capturing local (strategic) interactions (Kearns, Littman, & Singh)
 - Each “player” represented by a vertex
 - Payoff to i , is only a function of neighbors actions
 - Compact (yet general) representation of game
 - Assume max degree is d , then representation is of $O(a^n)$.
- Can graphical games be learned faster than general

- $X_{i,t}$ need only capture plays of neighbors
 - $N(i)$ is the set of neighbors of i (assume $|N(i)| < \infty$)
 - $S_{N(i)-i}$ is actions of all neighbors excluding self
 - $u_{i,t} = u_i(S_{i,t}, S_{N(i)-i})$
 - $p_{i,t}$ is forecast of $X_{i,t}$

- Same proof as before shows that for a NE we need

$$(\forall i) \quad \sum_{t=1}^T (X_{i,t} - p_{i,t}) w(\vec{p}_t) \rightarrow 0$$

- But we desire to do better for structured games.

(This is $(1/\epsilon)^{na^d}$, while the representation of a game is na^d .)

- We don't need to check $w(\vec{p}_t)$
- Instead we can check only

$$(\forall i) \sum_{t=1}^T (X_{i,t} - p_{i,t}) w(\vec{p}_{N(i),t}) \rightarrow$$

where $\vec{p}_{N(i),t}$ is a vector of all the p 's of all the ne

- Since this is all that matters in $u_i()$, rationality against the entire \vec{p} .
- Complexity: $n(1/\epsilon)^{a^{2d}}$
- The complexity is $|\mathcal{I}|$.
- NOTE TO SELF: No matter how excited you are, complexity, never, write it as $|\mathcal{I}|$!

A even smaller observable set

- X_i = personal utility
- p_i = forecast of personal utility
- $w()$ is local:

$$(\forall i) \quad \sum_{t=1}^T (X_{i,t} - p_{i,t}) w(\vec{p}_{N(i),t}) \rightarrow 0$$

- Converges to NE.
- Complexity: $n(1/\epsilon)^{a^d}$

- X_i = action taken
- p_i = forecast of own action
- decisions are made based on other peoples forecast
- $w()$ is local:

$$(\forall i) \quad \sum_{t=1}^T (X_{i,t} - p_{i,t}) w(\vec{p}_{N(i),t}) \rightarrow 0$$

- Converges to NE.
- Complexity: $n(1/\epsilon)^{a^d}$
- Violations can cause the system to crumble

Speed of convergence:

- Complexity: $n(1/\epsilon)^{da^d}$
- Recall, game representation is na^d
- Hence, the max degree is the bottleneck!
- Can get better results with utility forecasts: $n(1/\epsilon)^{da^d}$

CPU time:

- For tree games, fast per round computation
- Total CPU time comparable to NashProp
- For general graphs, could be hard to make forecasts

Future directions

- Analyze the CPU complexity
 - Have we just pushed the difficulty back to the step?
- Look at other games with simple structure
- Look at linear weightings rather than local weightings

See reverse side of handout for related re