Computational Game Theory (CIS 620/OPIM 952)
A Sampling Argument for Approximate NE
M. Kearns Jan 21, 2002

Setting: 2-player, m-action games

Approach:

- First show that there always $\exists$ approx. NE (P,Q) s.t. P & Q have small support
- Then search small-support P & Q "exhaustively"
- Result will be a "pseudopolynomial" algorithm for computing approx. NE

Defn. $\forall$ distribution P, support$(P) = \{i : P_i \neq 0\}$

Note that if (P,Q) is a NE, then

- $i \in$ support$(P) \Rightarrow i$ a best response to Q
- $j \in$ support$(Q) \Rightarrow j$ is a b.r. to P

So let (P,Q) be a NE in which the exp. payoffs are (U,V) $(U,V \in \mathbb{R})$.

Suppose we sample P & Q $\underline{\underline{t}}$ times each, and create empirical (sampled) distributions (P',Q').

Some notation:

- for any i, let
$$M_1(i,Q) \triangleq E_{j\sim Q}[M_1(i,j)]$$
- for any j, let
$$M_2(P,j) \triangleq E_{i\sim P}[M_2(i,j)]$$
- $$M(P,Q) \triangleq E_{i\sim P,\, j\sim Q}[M(i,j)]$$

Now $\forall$ fixed i, if sample size t is large enough, $M_1(i,Q') \simeq M_1(i,Q)$.

Why? The return that i gets in a single round against Q is a random variable with mean $M_1(i,Q)$.

$M_1(i,Q')$ is the average of this r.v. over t trials. So certainly

$$\lim_{t\to\infty} M_1(i,Q') = M_1(i,Q)$$

But we need something stronger.

Claim $\forall \varepsilon > 0$:

$$\Pr[\,|M_1(i,Q') - M_1(i,Q)| \geq \varepsilon\,] \leq e^{-\varepsilon^2 t/3}$$

$\hookrightarrow$ This prob. over the t draws of Q'.

(assumes $|M_1(i,j)| \leq 1$, but generalizes)

Large Deviation Bound

This was for a <u>single</u>, <u>fixed</u> i. But by $Pr[A \text{ or } B] \le Pr[A] + Pr[B]$ (<u>Union Bound</u>) we have:

$$\underbrace{Pr}\left[\,|M_1(i,Q') - M_1(i,Q)| \ge \varepsilon \text{ for } \underline{\text{some}}\ i\,\right] \le m e^{-\varepsilon^2 t/3}$$

→ again, over t samples of Q'

$+$

and also

$$\underbrace{Pr}\left[\,|M_2(P',j) - M_2(P,j)| \ge \varepsilon \text{ for } \underline{\text{some}}\ j\,\right] \le m e^{-\varepsilon^2 t/3}$$

→ draw of P'

$$\overline{}\ \ \text{''}$$

we'll come back to this... $\Rightarrow$ $2me^{-\varepsilon^2 t/3}$.

<u>Note:</u>

$\forall i\ |M_1(i,Q') - M_1(i,Q)| \le \varepsilon \Rightarrow \forall P''\ |M_1(P'',Q') - M_1(P'',Q)| \le \varepsilon$

$\forall j\ |M_2(P',j) - M_2(P,j)| \le \varepsilon \Rightarrow \forall Q''\ |M_2(P',Q'') - M_2(P,Q'')| \le \varepsilon$

So against P', can't get better than $V+\varepsilon$
against Q', can't get better than $U+\varepsilon$

Furthermore, since any $i \in \text{support}(P)$ gets at least $U-\varepsilon$ against Q', P' gets at least $U-\varepsilon$ against Q' (and Q' gets at least $V-\varepsilon$ against P').

<u>Thus</u>: (P',Q') is a 2ε-NE
(and has supports of size "only" t).

<u>Now</u>: to make sure such (P', Q') really <u>exist</u>, just need $2me^{-\varepsilon^2 t/3} < 1$, e.g.:

$$2me^{-\varepsilon^2 t/3} < 1/2$$

$$e^{-\varepsilon^2 t/3} < 1/4m$$

$$-\varepsilon^2 t/3 < \log(1/4m)$$

$$\varepsilon^2 t/3 > \log(4m)$$

$$\boxed{t > \frac{3}{\varepsilon^2} \underbrace{\log(4m)}_{\ll m}}$$

So if t is at least this large, we have a 50% chance of sampling (P', Q') that are 2ε-NE $\Longrightarrow$ must <u>exist</u> such a (P', Q')!

The <u>Probabilistic</u> <u>Method</u>

OK... but how do we *find* (P', Q')?
Remember, we don't know (P, Q) to begin!

Ugly but effective: discretized "exhaustive search":
We will search over *all* possible support sets of size t. This gives

$$\binom{m}{t} \times \binom{m}{t} \sim m^{2t} \text{ choices.}$$

For each such support set, we will *discretize* the possible distributions - only considering probabilities that are multiples of some small value Δ. For any *fixed* support sets, the # of such distributions is

$$\underbrace{\left(\frac{1}{\Delta}\right) \times \left(\frac{1}{\Delta}\right) \times \cdots \times \left(\frac{1}{\Delta}\right)}_{t \text{ times}} = \left(\frac{1}{\Delta}\right)^t$$

or $\left(\frac{1}{\Delta}\right)^{2t}$ for the pair of distributions.

Now claim that if a t-support set has a 2ε-NE, then we will find a

$(2\varepsilon + 2t\Delta)$-NE among the Δ-discretized distributions

So if we want to end up with an ε'-NE,

Solve $\quad 2\varepsilon + 2t\Delta \leq \varepsilon'$

e.g. $\varepsilon = \varepsilon'/4 \quad \Delta = \frac{\varepsilon'}{4t}$ works.

Total computation time goes like:

$$m^{2t} \times \left(\frac{1}{\Delta}\right)^{2t} = \left(\frac{m}{\Delta}\right)^{2t}$$

$$= \left(\frac{4tm}{\varepsilon'}\right)^{2t}$$

$$= O\left(\left(\frac{1}{(\varepsilon')^3} m \log(m)\right)^{\frac{48}{(\varepsilon')^2}\log(4m)}\right)$$

Upshot: for fixed ε',

can compute ε'-NE in time

$$O\left(\underbrace{(m\log(m))^{c \cdot \log(m)}}_{\text{subexponential in } m\ldots}\right)$$

(But there must be a better idea...)