Algorithmic Game Theory and Applications

Lecture 3: Nash's Theorem

Kousha Etessami

The Brouwer Fixed Point Theorem

We will use the following to prove Nash's Theorem.

Theorem(Brouwer, 1909) Every continuous function $f : D \mapsto D$ mapping a compact and convex, nonempty subset $D \subseteq \mathbb{R}^m$ to itself has a "fixed point", i.e., there is $x^* \in D$ such that $f(x^*) = x^*$.

Explanation:

- A "continuous" function is intuitively one whose graph has no "jumps". I.e., any "sufficiently little (non-zero) change" in x can change $f(x)$ by at most "as little (non-zero) change as desired".
- For our current purposes, we don't need to know exactly what "compact and convex" means.

(See the appendix of this lecture for definitions.)

We only state the following fact:

Fact The set of profiles $X = X_1 \times \ldots \times X_n$ is a compact and convex subset of R^m.
(Where $m = \Sigma_{i=1}^n m_i$, recalling that $m_i = |S_i|$.)

Simple cases of Brouwer's Theorem

To see a simple example of what Brouwer's theorem says, consider the interval $[0,1] = \{x \mid 0 \leq x \leq 1\}$.

$[0,1]$ is compact and convex.
(More generally, $[0,1]^n$ is compact and convex.)

For a continuous $f : [0,1] \mapsto [0,1]$, you can "visualize" why the theorem is true:

The "visual proof" in the 1-dimensional case:

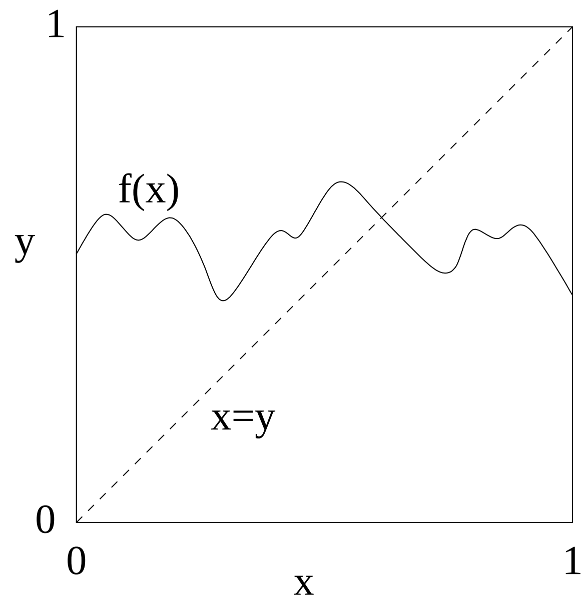

For $f : [0,1]^2 \mapsto [0,1]^2$, the theorem is already far less obvious: "the crumpled sheet experiment".

brief remarks

- Brouwer's Theorem is a deep and important result in topology.
- It is not very easy to prove, and we won't prove it.
- If you are desperate to see a proof, there are many. See, e.g., any of these:
  - [Milnor'66] (Differential Topology). (uses, e.g., Sard's Theorem).
  - [Scarf'67 & '73, Kuhn'68, Border'89], uses **Sperner's Lemma**.
  - [Rotman'88] (Algebraic Topology). (uses homology, etc.)
  - Possibly my favorite proof:
    [D. Gale'79] , uses the fact that HEX (a finite, extensive form game of perfect information, re-invented by Nash) is a "win-lose" game, i.e., any n-dimensional Hex game has a winner (i.e., can not end in a draw).

proof of Nash's theorem

Proof: (Nash's 1951 proof)

We will define a continuous function $f : X \mapsto X$, where $X = X_1 \times \ldots \times X_n$, and we will show that if $f(x^*) = x^*$ then $x^* = (x_1^*, \ldots, x_n^*)$ must be a Nash Equilibrium.

By Brouwer's Theorem, we will be done.

(In fact, it will turn out that x^* is a Nash Equilibrium if and only if $f(x^*) = x^*$.)

We start with a claim.

Claim: A profile $x^* = (x_1^*, \ldots, x_n^*) \in X$ is a Nash Equilibrium if and only if, for every player i, and every pure strategy $\pi_{i,j}$, $j \in S_i$,

$$U_i(x^*) \geq U_i(x^*_{-i}; \pi_{i,j})$$

proof of claim: If x^* is a NE then, it is obvious by definition that $U_i(x^*) \geq U_i(x^*_{-i}, \pi_{i,j})$.

For the other direction: by calculation it is easy to see that for any mixed strategy $x_i \in X_i$,

$$U_i(x^*_{-i}; x_i) = \sum_{j=1}^{m_i} x_i(j) * U_i(x^*_{-i}; \pi_{i,j})$$

I.e., the payoff of Player i is the "weighted average" of the payoffs of each of its pure strategies, j, weighted by the probability $x_i(j)$ of that strategy.

By assumption, $U_i(x^*) \geq U_i(x^*_{-i}; \pi_{i,j})$, for all j.

So, clearly $U_i(x^*) \geq U_i(x^*_{-i}; x_i)$, for any $x_i \in X_i$, because a "weighted average" of things no bigger than $U_i(x^*)$ can't be bigger than $U_i(x^*)$.

Hence, each x_i^* is a best response strategy to x^*_{-i}. In other words, x^* is a Nash Equilibrium. ∎

So, rephrasing our goal, we want to find $x^* = (x_1^*, \ldots, x_n^*)$ such that

$$U_i(x^*_{-i}; \pi_{i,j}) \leq U_i(x^*)$$

i.e., such that

$$U_i(x^*_{-i}; \pi_{i,j}) - U_i(x^*) \leq 0$$

for all players $i \in N$, and all $j = 1, \ldots, m_i$.

For a mixed profile $x = (x_1, x_2, \ldots, x_n) \in X$: let

$$\varphi_{i,j}(x) = \max\{0, U_i(x_{-i}; \pi_{i,j}) - U_i(x)\}$$

Intuitively, $\varphi_{i,j}(x)$ measures "how much better off" player i would be if he/she picked $\pi_{i,j}$ instead of x_i (and everyone else remained unchanged).

Define $f : X \mapsto X$ as follows: For $x = (x_1, x_2, \ldots, x_n) \in X$, let

$$f(x) = (x'_1, x'_2, \ldots, x'_n)$$

where for all i, and $j = 1, \ldots, m_i$,

$$x'_i(j) = \frac{x_i(j) + \varphi_{i,j}(x)}{1 + \sum_{k=1}^{m_i} \varphi_{i,k}(x)}$$

Facts:

1. If $x \in X$, then $f(x) = (x'_1, \ldots, x'_n) \in X$.

2. $f : X \mapsto X$ is continuous.

(These facts are not hard to check.)

Thus, by Brouwer, there exists $x^* = (x^*_1, x^*_2, \ldots, x^*_n) \in X$ such that $f(x^*) = x^*$.

Now we have to show x^* is a NE.

For each i, and for $j = 1, \ldots, m_i$,

$$x_i^*(j) = \frac{x_i^*(j) + \varphi_{i,j}(x^*)}{1 + \sum_{k=1}^{m_i} \varphi_{i,k}(x^*)}$$

thus,

$$x_i^*(j)(1 + \sum_{k=1}^{m_i} \varphi_{i,k}(x^*)) = x_i^*(j) + \varphi_{i,j}(x^*)$$

hence,

$$x_i^*(j) \sum_{k=1}^{m_1} \varphi_{i,k}(x^*) = \varphi_{i,j}(x^*)$$

We will show that in fact this implies $\varphi_{i,j}(x^*)$ must be equal to 0 for all j.

Claim: For any mixed profile x, for each player i, there is some j such that $x_i(j) > 0$ and $\varphi_{i,j}(x) = 0$.

<u>*Proof of claim:*</u> For any $x \in X$,

$$\varphi_{i,j}(x) = \max\{0, U_i(x_{-i}; \pi_{i,j}) - U_i(x)\}$$

Since $U_i(x)$ is the "weighted average" of $U_i(x_{-i}; \pi_{i,j})$'s, based on the "weights" in x_i, there must be some j used in x_i, i.e., with $x_i(j) > 0$, such that $U_i(x_{-i}; \pi_{i,j})$ is no more than the weighted average. I.e.,

$$U_i(x_{-i}; \pi_{i,j}) \leq U_i(x)$$

I.e.,

$$U_i(x_{-i}; \pi_{i,j}) - U_i(x) \leq 0$$

Therefore,

$$\varphi_{i,j}(x) = \max\{0, U_i(x_{-i}; \pi_{i,j}) - U_i(x)\} = 0$$

■

Thus, for such a j, $x_i^*(j) > 0$ and

$$x_i^*(j) \sum_{k=1}^{m_1} \varphi_{i,k}(x^*) = 0 = \varphi_{i,j}(x^*)$$

But, since $\varphi_{i,k}(x^*)$'s are all ≥ 0, this means $\varphi_{i,k}(x^*) = 0$ for all $k = 1, \ldots, m_i$. Thus,

For all players i, and for $j = 1, \ldots, m_i$,

$$U_i(x^*) \geq U_i(x^*_{-i}; \pi_{i,j})$$

Q.E.D. (Nash's Theorem)

In fact, since $U_i(x^*)$ is the "weighted average" of $U_i(x^*_{-i}, \pi_{i,j})$'s, we see that

Useful Corollary for Nash Equilibria:
$U_i(x^*) = U_i(x^*_{-i}, \pi_{i,j})$, whenever $x_i^*(j) > 0$.

Rephrased: In a Nash Equilibrium x^*, if $x_i^*(j) > 0$ then $U_i(x^*_{-i}; \pi_{i,j}) = U_i(x^*)$; i.e., each such $\pi_{i,j}$ is itself a "best response" to x^*_{-i}.

This is a subtle but very important point.
It will be useful later when we try to compute NE's.

Remarks

- The proof using Brouwer gives ostensibly no clue how to compute a Nash Equilibrium. It just says it exists!

- We will come back to the question of computing Nash Equilibria in general games later in the course.

- We start next time with a special case: 2-player zero-sum games (e.g., of the Rock-Paper-Scissor's variety). These have an elegant theory (von Neumann 1928), predating Nash.

- To compute solutions for 2p-zero-sum games, Linear Programming will come into play.
  Linear Programming is a very important tool in algorithms and optimization. Its uses go FAR beyond solving zero-sum games. So it will be a good opportunity to learn about LP.

supplementary reading: evolutionary biology as a game

- One way to view how we might "arrive" at a Nash equilibrium is through a process of evolution.

- John Maynard Smith (1972-3,'82) introduced game theoretic ideas into evolutionary biology with the concept of an Evolutionarily Stable Strategy.

- Your extra reading is from Straffin(1993) which gives a very amusing introduction to this.

- Intuitively, a mixed strategy can be viewed as percentages in a population that exhibit different behaviors (strategies).

- Their behaviors effect each other's survival, and thus each strategy has a certain survival value dependent on the strategy of others.

- The population is in "evolutionary equilibrium" if no "mutant" strategy could invade it and "take over".

a glossary for your reading

- **Definition** A 2-player game is **symmetric** if $S_1 = S_2$, and for all $s_1, s_2 \in S_1$, $u_1(s_1, s_2) = u_2(s_2, s_1)$.
- **Definition** In a 2p-sym-game, mixed strategy x_1^* is an **Evolutionarily Stable Strategy** (ESS), if:
  1. x_1^* is a best response to itself, i.e., $x^* = (x_1^*, x_1^*)$ is a <u>symmetric</u> Nash Equilibrium, &
  2. If $x_1' \neq x_1^*$ is any other best response to x_1^*, then $U_1(x_1', x_1') < U_1(x_1^*, x_1')$.

  Nash (1951, p. 289) also proves that every symmetric game has a <u>symmetric NE</u>, (x_1^*, x_1^*). (However, not every symmetric game has a ESS.)
- Given a profile $x \in X$ in an n-player game, the **"(purely utilitarian) social welfare"** is:
$$U_1(x) + U_2(x) + \ldots + U_n(x).$$
- A profile $x \in X$ is **pareto efficient** (a.k.a., **pareto optimal**) if there is no other profile x' such that $U_i(x) \leq U_i(x')$ for all players i, <u>and</u> such that for some player k, $U_k(x) < U_k(x')$.
- (Prisoner's Dilemma shows that NE's need not optimize social welfare nor be Pareto optimal.)

How hard is it to detect an ESS?

- It turns out that even deciding whether a 2-player symmetric game has an ESS is hard. It is both NP-hard and coNP-hard, and contained in Σ_2^P:

  K. Etessami & A. Lochbihler, "The computational complexity of Evolutionarily Stable Strategies", *International Journal of Game Theory*, vol. 31(1), pp. 93–113, 2008.

  (And, more recently, it has been shown Σ_2^P-complete, see:

  V. Conitzer, *"The exact computational complexity of Evolutionary Stable Strategies"*, in *Proceeding of Web and Internet Economics* (WINE), pages 96-108, 2013.)

- For simple 2×2 2-player symmetric games, there is a simple way to detect whether there is an ESS, and if so to compute one (described in the supplementary reading from Straffin).

- There is a huge literature on ESS and on *"Evolutionary Game Theory"*. See, e.g., the following book: J. Weibull, *Evolutionary Game Theory*, 1997.

Appendix: continuity, compactness, convexity

Definition For $x, y \in \mathbb{R}^n$, $\texttt{dist}(x, y) = \sqrt{\sum_{i=1}^{n}(x(i) - y(i))^2}$ denotes the Euclidean distance between points x and y.
A function $f : D \subseteq \mathbb{R}^n \mapsto \mathbb{R}^n$ is **continuous at a point** $x \in D$ if for all $\epsilon > 0$, there exists $\delta > 0$, such that for all $y \in D$: if $\texttt{dist}(x, y) < \delta$ then $\texttt{dist}(f(x), f(y)) < \epsilon$.
f is called **continuous** if it is continuous at every point $x \in D$.

Definition A set $K \subseteq \mathbb{R}^n$ is **convex** if for all $x, y \in K$ and all $\lambda \in [0, 1]$, $\lambda x + (1 - \lambda)y \in K$.

Rather than stating a general definition of compactness for arbitrary topological spaces, we use the following fact as a definition, restricted to Euclidean space:

Fact A set $K \subseteq \mathbb{R}^n$ is **compact** if and only if it is **closed** and **bounded**. (So, we need to define "closed" and "bounded".)

Definition A set $K \subseteq \mathbb{R}^n$ is **bounded** iff there is some non-negative integer M, such that $K \subseteq [-M, M]^n$.
(i.e., K "fits inside" a finite n-dimensional box.)

Definition A set $K \subseteq \mathbb{R}^n$ is **closed** iff for all sequences $x_0, x_1, x_2, \ldots$, where for all $i \geq 0$, $x_i \in K$, if there exists $x \in \mathbb{R}^n$ such that $x = \lim_{i \to \infty} x_i$ (i.e., for all $\epsilon > 0$, there exists integer $k > 0$ such that $\texttt{dist}(x, x_m) < \epsilon$ for all $m > k$), then $x \in K$.
(In other words, if a sequence of points is in K then its limit (if it exists) must also be in K.)