Game Theory in Graph-Based Computer Vision and Pattern Recognition

Marcello Pelillo

Ca' Foscari University, Venice

pelillo@dsi.unive.it

Summer School on Graphs in Computer Graphics, Image and Signal Analysis
Bornholm, Denmark, August 2011

Outline

Lecture 1: Monday (10:15 – 11:00)
Introduction to the basic concepts of game theory

Lecture 2: Tuesday (11:15 – 12:00)
Evolutionary games and graph-based data clustering

Exercises: Tuesday (13:00 – 14:30)
Reading groups and mini-presentations

Lecture 3: Thursday (09:15 – 10:00)
Graph labeling problems and graph transduction

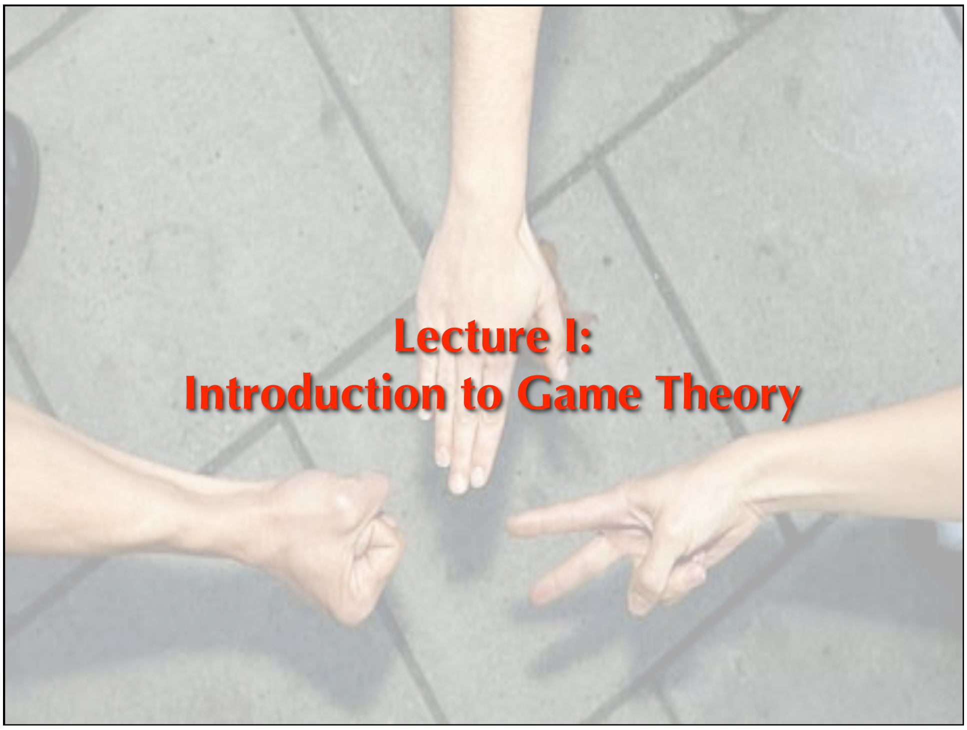

What is Game Theory?

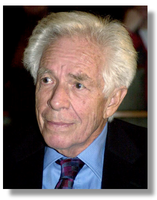

"The central problem of game theory was posed by von Neumann as early as 1926 in Göttingen. It is the following: If n players, $P_1, \ldots, P_n$, play a given game Γ, how must the i^{th} player, P_i, play to achieve the most favorable result for himself?"

Harold W. Kuhn
Lectures on the Theory of Games (1953)

A few cornerstones in game theory

1921–1928: Emile Borel and John von Neumann give the first modern formulation of a mixed strategy along with the idea of finding minimax solutions of normal-form games.

1944, 1947: John von Neumann and Oskar Morgenstern publish *Theory of Games and Economic Behavior*.

1950–1953: In four papers John Nash made seminal contributions to both non-cooperative game theory and to bargaining theory.

1972–1982: John Maynard Smith applies game theory to biological problems thereby founding "evolutionary game theory."

late 1990's –: Development of algorithmic game theory…

Normal-form Games

We shall focus on finite, non-cooperative, simultaneous-move games in **normal form**, which are characterized by:

- ✓ A set of **players**: $I = \{1, 2, \dots, n\}$ $(n \geq 2)$
- ✓ A set of **pure strategy profiles**: $S = S_1 \times S_2 \times \dots \times S_n$ where each $S_i = \{1, 2, \dots, m_i\}$ is the (finite) set of "pure" strategies (actions) available to the player i
- ✓ A **payoff function**: $\pi : S \rightarrow \Re^n$, $\pi(s) = (\pi_1(s),\dots,\pi_n(s))$, where $\pi_i(s)$ $(i=1\dots n)$ represents the "payoff" (or utility) that player i receives when strategy profile s is played

Each player is to choose one element from his strategy space in the absence of knowledge of the choices of the other players, and "payments" will be made to them according to the function $\pi_i(s)$.

Players' goal is to maximize their own returns.

Two Players

In the case of two players, payoffs can be represented as two $m_1 \times m_2$ matrices (say, A for player 1 and B *for player 2*):

$$A = (a_{hk}) \qquad a_{hk} = \pi_1(h,k)$$

$$B = (b_{hk}) \qquad b_{hk} = \pi_2(h,k)$$

Special cases:

- ✓ Zero-sum games: $A + B = 0$ ($a_{hk} = -b_{hk}$ for all h and k)
- ✓ Symmetric games: $B^T = A$
- ✓ Doubly-symmetric games: $A = A^T = B^T$

Example 1:
Prisoner's Dilemma

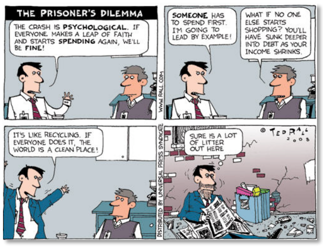

		Prisoner 2	
		Confess (defect)	Deny (cooperate)
Prisoner 1	Confess (defect)	-10 , -10	-1 , -25
	Deny (cooperate)	-25 , -1	-3 , -3

How to "Solve" the Game?

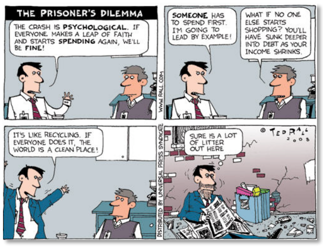

Dominated strategy !

		Prisoner 2	
		Confess (defect)	Deny (cooperate)
Prisoner 1	Confess (defect)	-10 , -10	-1 , -25
	Deny (cooperate)	-25 , -1	-3 , -3

Dominated strategy !

Example 2:
Battle of the Sexes

		Wife	
		Soccer	Ballet
Husband	Soccer	2 , 1	0 , 0
	Ballet	0 , 0	1 , 2

Example 3:
Rock-Scissors-Paper

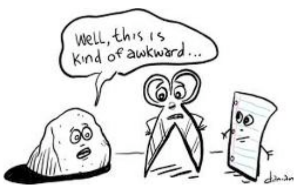

		You		
		Rock	Scissors	Paper
Me	Rock	0 , 0	1 , -1	-1 , 1
	Scissors	-1 , 1	0 , 0	1 , -1
	Paper	1 , -1	-1 , 1	0 , 0

Mixed Strategies

A **mixed strategy** for player i is a probability distribution over his set S_i of pure strategies, which is a point in the (m_i-1)-dimensional **standard simplex**:

$$\Delta_i = \left\{ x_i \in R^{m_i} \ : \ \forall h = 1 \ldots m_i : x_{ih} \geq 0, \text{ and } \sum_{h=1}^{m_i} x_{ih} = 1 \right\}$$

The set of pure strategies that is assigned positive probability by mixed strategy $x_i \in \Delta_i$ is called the **support** of x_i:

$$\sigma(x_i) = \left\{ h \in S_i : x_{ih} > 0 \right\}$$

A **mixed strategy profile** is a vector $x = (x_1,\ldots,x_n)$ where each component $x_i \in \Delta_i$ is a mixed strategy for player $i \in I$.

The **mixed strategy space** is the multi-simplex $\Theta = \Delta_1 \times \Delta_2 \times \ldots \times \Delta_n$

Standard Simplices

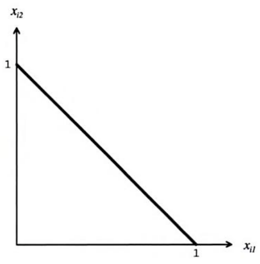

$m_i = 2$

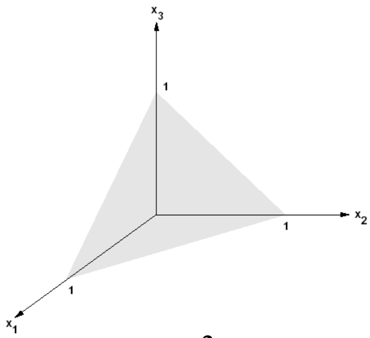

$m_i = 3$

Note: Corners of standard simplex correspond to pure strategies.

Mixed-Strategy Payoff Functions

In the standard approach, all players' randomizations are assumed to be independent.

Hence, the probability that a pure strategy profile $s = (s_1,\ldots,s_n)$ will be used when a mixed-strategy profile x is played is:

$$x(s) = \prod_{i=1}^{n} x_{is_i}$$

and the expected value of the payoff to player i is:

$$u_i(x) = \sum_{s\in S} x_i(s)\pi_i(s)$$

In the special case of two-players games, one gets:

$$u_1(x) = \sum_{h=1}^{m_1}\sum_{k=1}^{m_2} x_{1h}a_{hk}x_{2k} = x_1^T A x_2 \qquad u_2(x) = \sum_{h=1}^{m_1}\sum_{k=1}^{m_2} x_{1h}b_{hk}x_{2k} = x_1^T B x_2$$

where A and B are the payoff matrices of players 1 and 2, respectively.

Best Replies

Notational shortcut. Here, and in the sequel, if $z\in\Theta$ and $x_i\in\Delta_i$, the notation (x_i,z_{-i}) stands for the strategy profile in which player $i\in I$ plays strategy x_i, while all other players play according to z.

Player i's **best reply** to the strategy profile x_{-i} is a mixed strategy $x_i^*\in\Delta_i$ such that

$$u_i(x_i^*,x_{-i}) \geq u_i(x_i,x_{-i})$$

for all strategies $x_i \in\Delta_i$.

The best reply is not necessarily unique. Indeed, except in the extreme case in which there is a unique best reply that is a pure strategy, the number of best replies is always infinite.

Indeed:

- ✓ When the support of a best reply includes two or more pure strategies, any mixture of these strategies must also be a best reply
- ✓ Similarly, if there are two pure strategies that are individually best replies, any mixture of the two is necessarily also a best reply

Nash Equilibria

The Nash equilibrium concept is motivated by the idea that a theory of rational decision-making should not be a self-destroying prophecy that creates an incentive to deviate for those who believe it.

A strategy profile $x \in \Theta$ is a **Nash equilibrium** if it is a best reply to itself, namely, if:

$$u_i(x_i, x_{-i}) \geq u_i(z_i, x_{-i})$$

for all $i = 1 \ldots n$ and all strategies $z_i \in \Delta_i$.

If strict inequalities hold for all $z_i \neq x_i$ then x is said to be a **strict Nash equilibrium**.

Theorem. A strategy profile $x \in \Theta$ is a Nash equilibrium if and only if for every player $i \in I$, every pure strategy in the support of x_i is a best reply to x_{-i}.

It follows that every pure strategy in the support of any player's equilibrium mixed strategy yields that player the same payoff.

Finding Pure-strategy Nash Equilibria

		Player 2		
		Left	Middle	Right
Player 1	Top	3 , 1	2 , 3	10 , 2
	High	4 , 5	3 , 0	[illegible] , 4
	Low	2 , 2	5 , 4	12 , 3
	Bottom	5 , 6	4 , 5	9 , 7

Nash equilibrium!

Multiple Equilibria in Pure Strategies

Nash equilibrium!

Nash equilibrium!

		Wife	
		Soccer	Ballet
Husband	Soccer	2 , 1	0 , 0
	Ballet	0 , 0	1 , 2

No Equilibrium in Pure Strategies

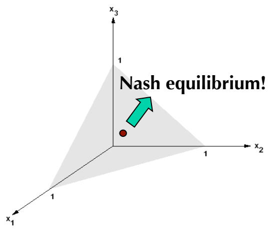

ium!

No Nas

		You		
		Rock	Scissors	Paper
Me	Rock	0 , 0	1 , -1	-1 , 1
	Scissors	-1 , 1	0 , 0	1 , -1
	Paper	1 , -1	-1 , 1	0 , 0

Existence of Nash Equilibria

Theorem (Nash, 1951). Every finite normal-form game admits a mixed-strategy Nash equilibrium.

Idea of proof.

1. Define a continuous map T on Δ such that the fixed points of T are in one-to-one correspondence with Nash equilibria.
2. Use Brouwer's theorem to prove existence of a fixed point.

"Together with factoring, *the complexity of finding a Nash equilibrium is in my opinion the most important concrete open question on the boundary of P today."*

Christos Papadimitriou
Algorithms, games, and the internet (2001)

Evolution and the Theory of Games

"We repeat most emphatically that our theory is thoroughly static. A dynamic theory would unquestionably be more complete and therefore preferable.
But there is ample evidence from other branches of science that it is futile to try to build one as long as the static side is not thoroughly understood."

John von Neumann and Oskar Morgenstern
Theory of Games and Economic Behavior (1944)

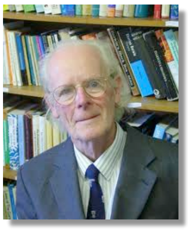

"Paradoxically, it has turned out that game theory is more readily applied to biology than to the field of economic behaviour for which it was originally designed."

John Maynard Smith
Evolution and the Theory of Games (1982)

Evolutionary Games

Introduced by John Maynard Smith (1973, 1974, 1982) to model the evolution of behavior in animal conflicts.

Assumptions:

- ✓ A large population of individuals belonging to the same species which compete for a particular limited resource
- ✓ This kind of conflict is modeled as a symmetric two-player game, the players being pairs of randomly selected population members
- ✓ Players do not behave "rationally" but act according to a pre-programmed behavioral pattern
- ✓ Reproduction is assumed to be asexual
- ✓ Utility is measured in terms of Darwinian fitness, or reproductive success

Interpreting Mixed Strategies

There are two ways to interpret the notion of a mixed strategy into the evolutionary framework:

1. Each individual is hard-wired to play a **pure strategy**, but some portion of the population plays one strategy while the rest of the population plays another.

2. Each individual is hard-wired to play a particular **mixed strategy** – that is, they are genetically configured to choose randomly from among certain options with certain probabilities.

It turns out that the two interpretations are mathematically equivalent.

In defining the "static" notions of evolutionary game theory, it is customary to focus on the second idea; the "dynamical" aspects are instead more conveniently dealt with using the first.

Evolutionary Stability

A strategy is **evolutionary stable** if it is resistant to invasion by new strategies.

Formally, assume:

- ✓ A small group of "invaders" appears in a large populations of individuals, all of whom are pre-programmed to play strategy $x \in \Delta$
- ✓ Let $y \in \Delta$ be the strategy played by the invaders
- ✓ Let ε be the share of invaders in the (post-entry) population ($0 < \varepsilon < 1$)

The payoff in a match in this bimorphic population is the same as in a match with an individual playing mixed strategy:

$$w = \varepsilon y + (1 - \varepsilon)x \in \Delta$$

hence, the (post-entry) payoffs got by the incumbent and the mutant strategies are $u(x,w)$ and $u(y,w)$, respectively.

Evolutionary Stable Strategies

Definition. A strategy $x \in \Delta$ is said to be an **evolutionary stable strategy** (ESS) if for all $y \in \Delta - \{x\}$ there exists $\delta \in (0,1)$, such that for all $\varepsilon \in (0, \delta)$ we have:

$$\underbrace{u[x, \varepsilon y + (1 - \varepsilon)x]}_{\textit{incumbent}} > \underbrace{u[y, \varepsilon y + (1 - \varepsilon)x]}_{\textit{mutant}}$$

Theorem. A strategy $x \in \Delta$ is an ESS if and only if it meets the following first- and second-order best-reply conditions:

1. $u(y,x) \leq u(x,x)$ for all $y \in \Delta$
2. $u(y,x) = u(x,x) \Rightarrow u(y,y) < u(x,y)$ for all $y \in \Delta - \{x\}$

Note. From the conditions above, we have:

- ✓ $\Delta^{ESS} \subseteq \Delta^{NE}$
- ✓ If $x \in \Delta$ is a strict Nash equilibrium, then x is an ESS

Existence of ESS's

Unlike Nash equilibria existence of ESS's is not guaranteed.

- ✓ Unique Na … /3,1/3)^T
- ✓ Hence, all
- ✓ Let the "m
- ✓ But $u(y,y) = u(x,y)$, hence $\Delta^{ESS} = \emptyset$

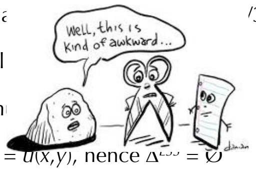

		You		
		Rock	Scissors	Paper
Me	Rock	0 , 0	1 , -1	-1 , 1
	Scissors	-1 , 1	0 , 0	1 , -1
	Paper	1 , -1	-1 , 1	0 , 0

Complexity Issues

Two questions of computational complexity naturally present themselves:

- ✓ What is the complexity of determining whether a given game has an ESS (and of finding one)?
- ✓ What is the complexity of recognizing whether a given x is an ESS for a given game?

Theorem (Etessami and Lochbihler, 2004). Determining whether a given two-player symmetric game has an ESS is both NP-hard and coNP-hard.

Theorem (Nisan, 2006). Determining whether a (mixed) strategy x is an ESS of a given two-player symmetric game is coNP-hard.

Replicator Dynamics

Let $x_i(t)$ the population share playing pure strategy *i* at *time t*. The **state** of the population at time *t* is: $x(t) = (x_1(t),\dots,x_n(t)) \in \Delta$.

Replicator dynamics (Taylor and Jonker, 1978) are motivated by Darwin's principle of natural selection:

$$\frac{\dot{x}_i}{x_i} \propto \text{payoff of pure strategy i} - \text{average population payoff}$$

which yields:

$$\begin{aligned} \dot{x}_i &= x_i\left[u(e^i,x) - u(x,x)\right] \\ &= x_i\left[(Ax)_i - x^T A x\right] \end{aligned}$$

Notes.

- ✓ Invariant under positive affine transformations of payoffs (i.e., $u \leftarrow \alpha u + \beta$, with $\alpha > 0$)
- ✓ Standard simplex Δ is invariant under replicator dynamics, namely, $x(0) \in \Delta \Rightarrow x(t) \in \Delta$, for all $t > 0$ (so is its interior and boundary)

Replicator Dynamics and ESS's

Theorem (Nachbar, 1990; Taylor and Jonker, 1978). A point $x \in \Delta$ is a Nash equilibrium if and only if *x* is the limit point of a replicator dynamics trajectory starting from the interior of Δ.
Furthermore, if $x \in \Delta$ is an ESS, then it is an asymptotically stable equilibrium point for the replicator dynamics.

The opposite need not be true.

$$A = \begin{bmatrix} 0 & 6 & -4 \\ -3 & 0 & 5 \\ -1 & 3 & 0 \end{bmatrix}$$

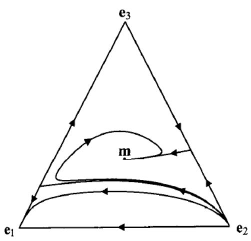

- ✓ The point $m = (1/3,1/3,1/3)^T$ is asymptotically stable (its eigenvalues have negative parts).
- ✓ But $e^1 = (1,0,0)^T$ is an ESS.
- ✓ Hence *m* cannot be an ESS (being in the interior, it would have to be the unique ESS).

Doubly Symmetric Games

In a doubly symmetric (or partnership) game, the payoff matrix A is symmetric ($A = A^T$).

Fundamental Theorem of Natural Selection (Losert and Akin, 1983).
For any doubly symmetric game, the average population payoff $f(x) = x^T A x$ is strictly increasing along any non-constant trajectory of replicator dynamics, namely, $d/dt f(x(t)) \geq 0$ for all $t \geq 0$, with equality if and only if $x(t)$ is a stationary point.

Characterization of ESS's (Hofbauer and Sigmund, 1988)

For any doubly simmetric game with payoff matrix A, the following statements are equivalent:

a) $x \in \Delta^{ESS}$

b) $x \in \Delta$ is a strict local maximizer of $f(x) = x^T A x$ over the standard simplex Δ

c) $x \in \Delta$ is asymptotically stable in the replicator dynamics

Discrete-time Replicator Dynamics

A well-known discretization of replicator dynamics, which assumes non-overlapping generations, is the following (assuming a non-negative A):

$$x_i(t+1) = x_i(t) \frac{A(x(t))_i}{x(t)^T A x(t)}$$

which inherits most of the dynamical properties of its continuous-time counterpart (*e.g.*, the fundamental theorem of natural selection).

MATLAB implementation

```
distance=inf;
while distance>epsilon
    old_x=x;
    x = x.*(A*x);
    x = x./sum(x);
    distance=pdist([x,old_x]');
end
```

References

Texts on (classical) game theory

J. von Neumann and O. Morgerstern. *Theory of Games and Economic Behavior*. Princeton University Press (1944, 1953).

D. Fudenberg and J. Tirole. *Game Theory*. MIT Press (1991).

M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press (1994).

Texts on evolutionary game theory

J. Weibull. *Evolutionary Game Theory*. MIT Press (1995).

J. Hofbauer and K. Sigmund. *Evolutionary Games and Population Dynamics*. Cambridge University Press (1998).

Computationally-oriented texts

N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani (Eds.) *Algorithmic Game Theory*. Cambridge University Press (2007).

Y. Shoham and K. Leyton-Brown. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press (2009).

On-line resources

http://gambit.sourceforge.net/	a library of game-theoretic algorithms
http://gamut.stanford.edu/	a suite of game generators for testing game algorithms

Reading Groups on Tuesday

Matching

M. Pelillo, K. Siddiqi, and S. W. Zucker. Matching hierarchical structures using association graphs. *ECCV 1998* (longer version in *PAMI 1999*).

M. Pelillo. Replicator equations, maximal cliques, and graph isomorphism. *NIPS 1998* (longer version in *Neural Computation 1999*).

M. Pelillo. Matching free trees with replicator equations. *NIPS 2001* (longer version in *PAMI 2002*).

A. Albarelli, A. Torsello, S. Rota Bulò, and M. Pelillo. Matching as a non-cooperative game. *ICCV 2009*.

Grouping

M. Pavan and M. Pelillo. Dominant sets and hierarchical clustering. *ICCV 2003*.

M. Pavan and M. Pelillo. Efficient out-of-sample extension of dominant-set clusters. *NIPS 2004*.

A. Torsello, S. Rota Bulò, and M. Pelillo. Beyond partitions: Allowing overlapping groups in pairwise clustering. *ICPR 2008*.

S. Rota Bulò and M. Pelillo. A game-theoretic approach to hypergraph clustering. *NIPS 2009*.

Applications

R. Hamid, A. Johnson, S. Batta, A. Bobick, C. Isbell, G. Coleman. Detection and explanation of anomalous activities: Representing activities as bags of event n-grams. *CVPR 2005*.

L. Li, X. Zhang, W. Hu, W. Li, and P. Zhu Soccer video shot classificaton based on color characterization using dominant sets clustering. *PCM 2009*.

A. Albarelli, E. Rodolà, A. Cavallarin, and A. Torsello. Robust figure extraction on textured background: A game-theoretic approach. *ICPR 2010*.

X. Yang, H. Liu, and L. J. Latecki. Contour-based object detection as dominant set computation. *ACCV 2010*.