Game Theory and Evolution

02-251 Spring 2019

Much of this material is derived from the book *Evolutionary Dynamics* by Martin A. Nowak.

1. Evolutionary Games

When looking at genetic algorithms, our fitness was a function only of the individual. It didn't take into account the other individuals in the environment. This is unrealistic.

Evolution can be a competitive or cooperative process. An bird evolves a beak that can reach particular insects, insects evolve a poison to avoid being eaten by the bird, etc. Cells in a cancer tumor evolve to avoid particular treatments. A virus evolves to avoid a vaccine. A bacteria evolves to not be susceptible to antibiotics. Organisms are evolving various *strategies* in games of competition and cooperation.

This leads to the use of game theory to try to model observed evolutionary phenomena. As with our model of evolution based on NAND gates, these game-theoretic models are idealized and omit many aspects of real evolution. By doing this, we can hope to gain some understanding about fundamental evolutionary principles.

"Game theory" was introduced by John von Neumann, a mathematician (and computer scientist before that was a thing). He was one of the architects of the first computers. He also coined the terms "transcription" and "translation".

1.1 Two-player games

- A two-player is given by a payoff grid P where the rows and columns represent *actions* that the players can take.
- In each cell of the grid, there are two numbers r_{ij}, c_{ij}.
- There is a "row" player and a "column" player (hence the name "two player")
- The row player chooses a row i, the column player chooses a column j, and the row player wins r_{ij} and the column player wins c_{ij}. If these numbers are negative, the player loses money (i.e. $r_{ij} = -3$ means the row player loses 3 dollars if the row player picks i and the column player picks j).
- Example: A prisoner's dilemma. Two colluding criminals have been caught by the police and taken to separate interrogation rooms. Each has to decide whether to "Confess" or to "Keep Quiet" ("Silence"). If you confess, but your co-conspirator does not, the police will go easy on you. Both if you both keep quiet, your punishment will be even less.
- Payoff matrix:

	Silence	Confess
Silence	-2,-2	-10,0
Confess	0,-10	-7,-7

  Read this as the "reward" that the player gets when playing the various strategies against the strategies of the other player. The player is trying to maximize their reward. (Here all the rewards are 0 or negative).
- Sometimes we will look only at the payoff matrix of the row player. For the prisoner's dilemma above that is:

	Silence	Confess
Silence	-2	-10
Confess	0	-7

- Consider a general row-player payoff matrix:

	A	B
A	a	b
B	c	d

  - *Dominating strategies.* What if $a > c$ and $b > d$? Then we should always play strategy "A". Similarly, if $c > a$ and $d > b$, we should always play "B".
  - If $a > c$ and $b < d$, you should try to mimmic the strategy of the other player. E.g.:

	A	B
A	10	3
B	7	10

  - If $c > a$ and $b > d$, then you should play the opposite strategy of the other player, E.g.:

	A	B
A	7	10
B	10	3

1.2 Nash equilibria

- John Nash is a famous CMU alumnus.
- A set of strategies for the players of a game is a Nash Equilibrium if deviation from the selected strategies by any one player causes the payoff of that player to get worse.
- Nash equilibria are stable in the sense that no one person has incentive to deviate from the plan.
  1. Is Silence / Silence a Nash equilibrium? No: If I switch to Confess while the other player stays with Silence, I do better.
  2. Is Silence / Confess a Nash equilibrium? again no: if the other player is confessing, I should switch to confess (going from -10 to -7).
  3. What about Confess / Silence? Nope, because the *other* player has an incentive to switch to Confess (this is the symmetric case).
  4. How about Confess / Confess? Yep: If I switch to Silence, I go from value -7 to -10. Same is true if the other player switches. Hence: Confess / Confess is a Nash Equilibrium.

1.3 Mixed strategies

- If a player is forced to pick a single row or column, they may never win. An example here is Rock/Paper/Scissors: if you always pick Rock, you're not going to win for very long:

	Rock	Paper	Scissors
Rock	0,0	-1,1	1,-1
Paper	1,-1	0,0	-1,1
Scissors	-1,1	1,-1	0,0

- These leads to the idea of *mixed strategies*: players choose an action randomly according to some probability distribution. This is what you naturally do in rock, paper, scissors!
- The connection to evolution: we think of a heterogenous population as implementing a mixed strategy. Note: in a population, each *individual* has a non-mixed, fixed action (i.e. "sleep during the day"), but the population as a whole implements a mixed strategy if individuals choose different strategies.

- If there are two actions A and B (i.e. "sleep during day" or "sleep during the night"), and x_A fraction of the population does action A and x_B fraction of the population does action B, then this is a mixed strategy with probability vector (x_A, x_B). The *player* is the entire population!
- Nash's famous theorem: If players are allowed to play mixed strategies, there is always a Nash equilibrium in a finite game.

1.4 Evolutionary Game Theory

- Following Chapter 4 of *Evolutionary Dynamics* by Martin A. Nowak.
- Assume a population of two different species (types, classes, alleles), A and B, with abundances x_A and x_B respectively ($0 \leq x_A, x_B \leq 1$).
- Assume pairs of members of the population meet each other randomly and play a game with row payoff matrix:

	A	B
A	a	b
B	c	d

  and assume the column player's payoff matrix is the same.
- We assume that the population is well mixed, so the probability that any individual meets an A-type is x_A (and the prob they meet a B-type is x_B). We can think of this as the "population" is playing the game with the mixed strategy of (x_A, x_B): the random choice part comes from which two individuals randomly meet.
- We set the *fitness* of the species to be the expected payoff from these games:

$$f_A = ax_A + bx_B \qquad \text{Type A meets an A with prob. } x_A \text{ and wins } a\text{, etc.} \tag{1}$$
$$f_B = cx_A + dx_B \qquad \text{Type B meets an A with prob. } x_A \text{ and wins } c\text{, etc.} \tag{2}$$

  In other words, the fitness of an A-type, given a population structure x_A, x_B, is $ax_A + bx_B$.
- The *average* fitness of the population is:

$$\phi = x_A f_A + x_B f_B \tag{3}$$

- Recall that $\dot{x}$ is shorthand for dx/dt. Suppose we have the following model of selection:

$$\dot{x}_A = x_A[f_A(\vec{x}) - \phi] \tag{4}$$
$$\dot{x}_B = x_B[f_B(\vec{x}) - \phi] \tag{5}$$

  The intuition: If type A is more fit than the average fitness ($f_A(\vec{x}) - \phi > 0$), the number of that type will increase, otherwise it decreases. You can think of the $[f_A(\vec{x}) - \phi]$ term as a score that is related to the chance an individual of type A will reproduce, which is why it is multiplied by x_A.
- Considering $x_B = 1 - x_A$, letting $x = x_A$, and substituting ϕ into (4), we have

$$\dot{x} = x[f_A(x) - xf_A(x) - (1-x)f_B(x)] \tag{6}$$
$$= x[(1-x)f_A(x) - (1-x)f_B(x)] \tag{7}$$
$$= x(1-x)[f_A(x) - f_B(x)] \tag{8}$$

- Substituting in (1) and (2) to $f_A(x) - f_B(x)$, we get:

$$f_A(x) - f_B(x) = ax + b(1-x) - cx - d(1-x) = ((a-c) - (b-d))x + b - d \tag{9}$$

- Behavior:
  - The steady state of the population are when (9) is 0. This happens when $x = (d-b)/(a-b-c+d)$, so long as this quantity is in $(0, 1)$.
  - If $a > c$ and $b > d$, then for any population $(x, 1-x)$ type A is going to be more fit, (9) will be positive, and the amount of type A will increase.
  - If $a < c$ and $b < d$, the opposite happens and A will decrease (and B will increase).
  - If If $a > c$ and $b < d$, then you want to mimimic the other player, so you will end up at the entire population being either A or B (depending on which was more common to start with).
  - If $a < c$ and $b > d$, then to beat an A you should be a B and vice versa $\rightarrow$ a stable equilibrium at $x = (d-b)/(a-b-c+d)$.

2. Spatial Games

2.1 A spatially distributed population

- Following Chapter 9 in Nowak.
- In real biological systems, the evolving individuals have some physical proximity to one another. One of the first people to consider this from a computational perspective was Alan Turing, who investigated computational systems that could give rise to complex patterns such as zebra strips or animal spots. (His approach was not based on game theory so we won't discuss it here.)
- In a spatial model, nearby individuals compete, and winning strategies (aka phenotypes) begin to take over.
- Idealized model of this:
  1. Individuals arranged in a 2D grid.
  2. The neighbors $N(i)$ of individual i are the 8 cells that share a border or a corner with i:

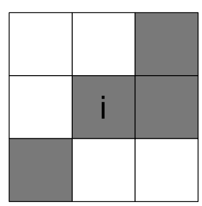

  3. Each cell (individual) has a *strategy* (from, say, a finite collection of possible strategies). The strategy might be "Install Windows" or "Install Mac OS X". It might be "Join Facebook" or "Join Twitter" (assuming you can do only one or the other). It might be "Adopt bright feathers" or "Adopt subdued feathers", etc.
  4. Now each game is played by two individual members of the population rather than between populations. So each cell is a player that plays lots of two-person games (one with each of its neighbors).
  5. We assume each payoff matrix for each cell is the same for each game and the same for both players.
  6. At each time step, each cell plays a game with each of its neighbors using its current strategy. The total payoff for each individual (cell) is the sum of the payoffs of each of those games:

$$\text{score}(i) = \sum_{j \in N(i)} \text{payoff}(\text{strategy}(i), \text{strategy}(j)) \tag{10}$$

7. For the next time step, a cell changes its strategy to the strategy of the cell with the highest score in its neighborhood. You can either think of this as an individual i consciously adopting the strategy or as the highest scoring (aka most fit) individual reproducing and outcompeting weaker phenotypes.

8. You can consider other geometries besides the 2D grid above by changing the $N(i)$ relation. E.g. (1) individuals have a location p_i and $N(i) = \{p_j \mid \text{dist}(p_j, p_i) \leq \theta \text{ and } j \neq i\}$, or (2) neighbors exclude the diagonal neighbors, or (3) individuals are on a different lattice (hexagonal, triangular, etc.). The neighbors could also be on a general graph. See for example:

   – Nicole Immorlica, Jon Kleinberg, Mohammad Mahdian, Tom Wexler. The role of compatibility in the diffusion of technologies in social networks, EC 2007 and references therein.

2.2 Invasion in Spatial Prisoner's Dilemma

- Central question: when can strategy A take over or at least survive among a sea of alternative strategies? E.g. within a tumor, can a few drug-resistant cells survive?
- Actually used for predicting technology adoption when the technology has "network effects": which word processor file format should I use? Which file-sharing protocol or encryption protocol should I use?
- Consider the game:

	C	D
C	$1, 1$	$0, b$
D	$b, 0$	ϵ, ϵ

  where $b > 1$ and ϵ is very small. This is a Prisoner's dilemma-type game: if they both cooperate (C), the two players do reasonably well (each gets a payoff of 1), but each has an incentive (because $b > 1$) to switch to the other strategy of defect (D).
- We set $\epsilon = 0$ for simplicity.
- The score of both strategies is $\approx$ the number of cooperators in the neighborhood (but strategy D gets b points for every cooperator, while strategy C gets 1 point for every cooperator).
- Cooperators grow in strength by being together in a block, while grouping defectors together weakens a defector.
- In fact, an isolated cooperator will disappear after a single step.
- When does C or D take over? Can they coexist?

2.3 Behavior for different values of b

- Defector in a sea of cooperators: will expand into a block of 9 defectors, and then will either: (a) continue expanding, (b) stay as a block of 9, (c) immediately shrink back to an isolated defector, or (d) shrink back to an isolated defector after 2 time steps (via a cross). (Figure from Nowak.)

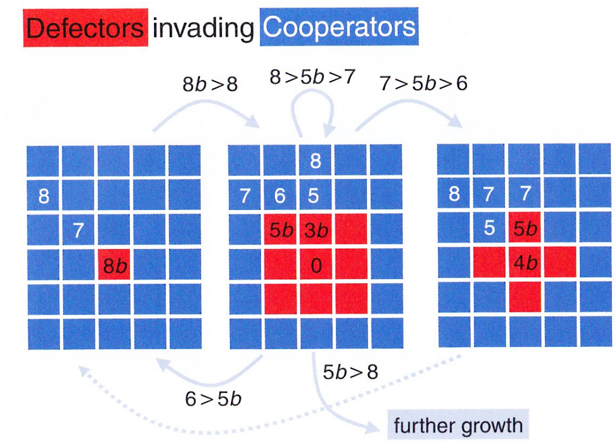

- Block of 9 cooperators in a sea of defectors: will either (a) expand uniformly, (b) expand along horizontal and vertical lines, (c) stay as a block of 9, or (d) disappear.
- In the tiny parameter region of $25/15 > b > 24/15$, you get "kaleidoscopes" and complex behavior where the frequency of cooperators fluctuates over time. This is because an expanding defector region is forced to expand only from the corners.
- In the right parameter region, a mass of cooperators can expand even in a sea of defectors. That mass can be either created by a large single block or from the collision of two "walker" structures:

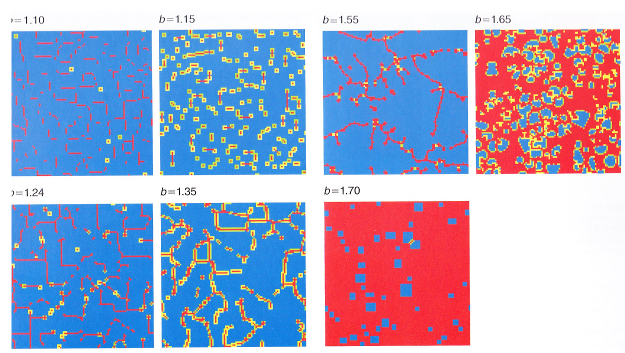

3. Other Biological Application of these Ideas*

- Sandholm. Steering Evolution Strategically: Computational Game Theory and Opponent Exploitation for Treatment Planning, Drug Design, and Synthetic Biology `https://www.cs.cmu.edu/~sandholm/game%20solving%20for%20medical.aaai15SMT.pdf`
- Orlando PA, Gatenby RA, Brown JS. Cancer treatment as a game: integrating evolutionary game theory into the optimal control of chemotherapy. Phys Biol. 2012 Dec;9(6):065007 `https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3653600/`
- Daniel Nichol, et al., Steering evolution with sequential therapy to prevent the emergence of bacterial antibiotic resistance. PLoS Comput Biol, 11(9):e1004493, 2015. `https://doi.org/10.1371/journal.pcbi.1004493`