Course Summary Homework

(Max useful score: 100 - Available points: 210)

15-382: Collective Intelligence (Spring 2018)

OUT: April 21, 2018, at 1:00am

DUE: May 1, 2018 at 1pm - Available late days: 0

Instructions

Homework Policy

Homework is due on Autolab by the posted deadline. For this final homework you cannot use late days. No credit will be given for homework submitted after the deadline.

If you find solutions in any source other than the material provided, you must mention the source.

Submission

Create a zipped archive including: a PDF file with the answers to the provided questions (they can be handwritten, but in this case you must have / use a "readable" handwriting). The zipped archive should be submitted to Homework on Autolab.

Contents

1 Ant clustering (20 points)

We have studied the use of ant-inspired algorithms for *emergent task allocation*. Let's imagine that the "tasks" are objects spread in the environment that need to be *clustered*. Real ants do actually deal with such a task for instance for cemetery formation. A biological model (Denebourg et al., 1991) of ant behaviors tells that an ant decides to *pick-up* an item which is in its range of action based on the probability

$$P_p = \Big(\frac{\gamma_p}{\gamma_p + \lambda}\Big)^2,$$

where λ is the fraction of items the ant perceives in its neighborhood, and $\gamma_p > 0$. Similarly, each *loaded* ant has a probability of *dropping* the carried object, given by:

$$P_d = \Big(\frac{\lambda}{\gamma_d + \lambda}\Big)^2,$$

provided that the corresponding cell / location is empty. The fraction of items, λ is calculated making use of a short-term memory for each ant.

It is apparent that this clustering model shares strong mathematical and conceptual similarities to the one we have used for ant-based task allocation, the main difference being in the fact that now an ant "drops" the current task under certain probabilistic conditions.

Let's use these ideas to perform *ant-based clustering*. The scenario is the following. The items to cluster are a given set V of N real-valued n-dimensional vectors, $V = \{\boldsymbol{x}_i \in \mathbb{R}^n,\ i = 1, \ldots, N\}$. For the sake of simplicity, all vectors have n dimensional values in some n-dimensional compact set $S^n = [a,b] \times [a,b] \times \ldots [a,b] \subseteq \mathbb{R}^n$ (e.g., all vectors' components have value in $[a = 0, b = 1]$). Vectors are similar or different, based on their *Euclidean distance*: $d(\boldsymbol{x}, \boldsymbol{y}) = \sqrt{\sum_{i=1}^{n}(x_i - y_i)^2}$. The *goal* is to cluster the vectors in a *minimal set of clusters* based on an threshold value $\gamma > 0$ that defines the scale of dissimilarity between two vectors: if two vectors have a distance $d(\boldsymbol{x}, \boldsymbol{y}) > \gamma$, they should not belong to the same cluster. Of course, γ depends on a and b (the distance between two vectors can't be larger than $\sqrt{a^2 + b^2}$), and to which extent we want to have a few clusters (large γ), versus a large number of clusters (small γ). In other words, the value of γ is a *strategic* one.

The ant agents *move* on a 2-dimensional *grid world* of size $m \times m$, where the S vectors are spread randomly (note: the coordinates associated to the vectors have no relation with the coordinates of the grid world).

1. Define and implement an ant-based, *fully distributed*, algorithm for the data clustering scenario based on the above hints (i.e., the real ant model can serve as basic intuition but you need to develop a full algorithm). You have to detail the algorithm, all the formulas that are going to be used for decision-making by the ant agents, and all the parameters of the algorithm (e.g., the local field of view of an ant, the size m of the grid, how the ants move, how many ants, ...).
2. Create reasonably large instance scenarios (with 2D and 3D data vectors, such that it's possible to visualize the data) and study the behavior of the algorithm. Report the data in plots and/or tables (and show the a few sample datasets) versus a few choices of γ. An important metric is the number of clusters that are being created. Moreover, at the end the *intra-cluster* distances (i.e., the distances among all the vectors that have been assigned to the same cluster) should be small, while the inter-clusters distance (i.e., the distance between cluster centroids) should be large (at least compared to intra-cluster ones).
3. Discuss the role of the different parameters of the algorithm and identify the most critical ones.

2 Rock-Paper-Scissors Cellular Automata (20 points)

Read the attached paper: *Simulating a Rock-Scissors-Paper Bacterial Game with a Discrete Cellular Automaton.* that describes the design and implementation of a discrete cellular automata simulating the generation, degradation and diffusion of particles in a two dimensional grid where different colonies of bacteria coexist and interact according to a dynamics analogous to the one induced by a Rock-Paper-Scissors game.

1. (Re)Implement the system. If any aspect is unclear, make your own design choice (and provide a justification / explanation).
2. Perform an experimental analysis of the system similar (or better) than that reported in the paper.
3. Provide an effective visualization of the dynamics of the system (e.g., https://www.youtube.com/watch?v=M4cV0nCIZoc).

4. Can you identify any ESS or Nash equilibria? Perform “some” analysis of the results in this respect.
5. Make a report with your choices, results, and a short critical summary of the paper (i.e., a mini review).

3 Measuring fractal dimension (20 points)

Read the attached paper *An introduction to dimension theory and fractal geometry: fractal dimensions and measures*. You only need to read up to section 5 (page 12). For the Koch curve mentioned at page 5 without further reference, you can find a simple yet complete description on the Wikipedia page https://en.wikipedia.org/wiki/Koch_snowflake.

1. Use the box-counting method to define an approximate measure of the fractal dimension of the Lorenz attractor that we have described in the classes. To generate the necessary data (for a given initial condition), you can use parts of the python code `viz-attractor.py` previously provided for the generation of orbits of chaotic systems. Note that you have to use different values of side length for the “boxes’ , and have to consider the slope that fits the data in the log data representation. Report data and code.

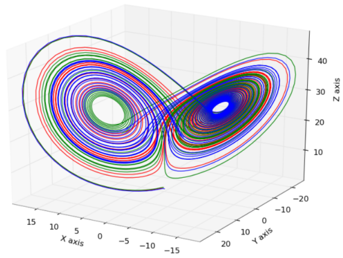

2. Generate the data to approximate a 3D Menger sponge, as described in https://en.wikipedia.org/wiki/Menger_sponge, and use the same approach as in the previous question to compute an approximation of the fractal dimension. Report data and code.

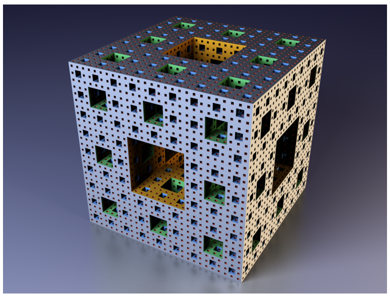

3. The image below charts the development of a rabbit family tree, that grows according to the Fibonacci iterated map. Moving from top to bottom, each point represents a pair of rabbits (Empty dots represent immature rabbit pairs, while filled dots represent mature rabbit pairs capable of breeding). Why does this tree structure has the property of a fractal?

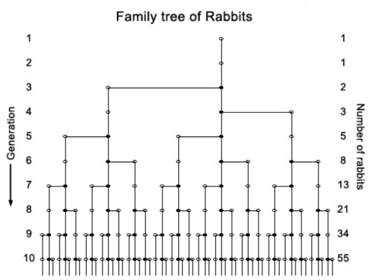

4. Let’s consider now the image of real tree (made of wood!), projected on a 2D plane. Does this tree representation has a fractal structure? Why?

4 Vickrey's second price auctions (20 points)

An item is to be sold at auction according to Vickrey's mechanism: each bidder submits a sealed bid, all the received bids are opened, the object is sold to the highest bidder, but the price is the bid of the second-highest bidder. There are n bidders, and let v_i be the value of the object to bidder i. Each bidder i has a *continuous set of strategies*: it can make any bid b_i in the range $[0, +\infty)$.

1. Define n players game that models the auction scenario, by defining the payoff functions (and assuming that payoff is zero if the bidder doesn't win the auction).

2. Formally show that for a player i the strategy $b_i = v_i$ *weakly dominates* every other strategy (you can assume that the top two bids are never the same). In other words, you should bid exactly what the object is worth to you, that was precisely the purpose of designing second-price auctions!

5 Firm's competition: a continuous game (25 points)

A certain good is produced by just two firms, which we label 1 and 2 (they form a *duopoly*). Each firm tries to maximize its own profit by choosing an appropriate level of production. Firm 1 chooses its level of production *first*; then Firm 2 observes this choice and chooses its own level of production. This is a form of a *leader-follower* Stackelberg game.

Let s be the quantity produced by Firm 1 and let t be the quantity produced by Firm 2. Then the total quantity of the good that is produced is $q = s + t$. The market price p of the good depends on a function ϕ of the available quantity of the good on the market: $p = \phi(q)$. At this price, everything that is produced can be sold. Suppose Firm 1's cost to produce the quantity s of the good is $c_1(s)$, and Firm 2's cost to produce the quantity t of the good is $c_2(t)$.

Let's denote the profits of the two firms by π_1 and π_2, where profit is revenue minus cost, and revenue is price times quantity sold. Each firm aims to maximize its profit.

Let's make the further assumptions:

- Price falls linearly with total production: $p = \alpha - \beta(t + s)$, $\alpha, \beta > 0$.
- Each firm has the same unit cost of production, $c > 0$: $c_1(s) = cs$ and $c_2(t) = ct$.
- $\alpha > c$: the price of the good when very little is produced is greater than the unit cost of production. If this assumption is violated, the good will not be produced.
- Firm 1 chooses its level of production s first. Then Firm 2 observes s and chooses t.
- The production levels s and t can be any real numbers.

1. Model the scenario as a two-player game by defining strategies and payoffs.

2. Use calculus (i.e., partial derivatives, analysis of functions) to find Firms' best response functions (as a function of the parameters of the model) and, in turn, Firm's production levels and profits.

6 Finding Nash equilibria (25 points)

Let's consider the following game in normal form:

		P_2 H	P_2 T
P_1	H	$+1, -1$	$-1, +1$
P_1	T	$-1, +1$	$+1, -1$

We want to find the Nash equilibria. We can use different approaches for the purpose, based on different, yet equivalent definitions of Nash equilibrium that we have given in the classes.

1. A definition of Nash equilibrium which is "useful" for finding Nash equilibria is based on the notion of best response: a pair of strategies (σ_1^*, σ_2^*) is a NE if σ_1^* is a best response to σ_2^* and vice versa. Formally, this means that: $\sigma_1^* \in \arg\max_{\sigma_1 \in \Sigma_1} \pi_1(\sigma_1, \sigma_2^*)$, and $\sigma_2^* \in \arg\max_{\sigma_2 \in \Sigma_2} \pi_2(\sigma_1^*, \sigma_2)$. To use this definition, we find for each player the set of best responses to every possible strategy of the other player. We then look for pairs of strategies that are best responses to each other.

   Use these results and insights to find the Nash equilibria of the given game.

2. Provide a graphical representation of the expected utility function for player 1 and discuss the result found in the previous question in relation to the characteristics of this function.

3. In Homework 5, one of the questions was about proving the *equality of payoffs theorem*: If (σ^*, σ^*) is a mixed Nash equilibrium in a two-player game, then $\pi_1(s, \sigma_2^*) = \pi_1(\sigma_1^*, \sigma_2^*)$, $\forall s \in S_1^*$, where π_1 is the expected payoff of player 1, S is the finite set of pure strategies, and S_1^* is the support of σ_1^* (σ_1^* is the mixed strategy adopted by player 1 at the Nash equilibrium).

   Exploit the theorem for finding the mixed strategy Nash equilibrium of the game.

4. Describe a real-world scenario that could be represented by the given payoff matrix.

7 Asymmetric pairwise contests (25 points)

There are many situations in which the players engaged in a contest can be distinguished (in addition to the fact that they may have different strategies). For instance, in biological problems they may be male and female individuals dividing up the care of their offspring, or they may be larger and smaller males competing over a group of females. In economic contexts, they may be a buyer and a seller, or a firm holding a monopoly and a firm wishing to enter the market. Such difference between individuals can easily lead to an asymmetric payoff matrix, since different player may have different actions available to them, or have different payoffs. Even if the payoff matrix stays symmetric, the presence of individuals playing different *roles* (e.g., male and females), introduces a problem in the notion of evolutionary stability since this requires symmetric games: same actions and symmetric payoffs for the two players in the contest.

A way to overcome the issue is to specify behavior for *all* roles in the contest: use s_1 in role r_1, use s_2 in role r_2, and so on. In other terms, we assume that in a contest an individual can find itself playing role r_i, while it can play role r_j in a later / different contest. By specifying role-conditioned strategies, we can get back a symmetric population game.[1]

Let's consider a variation of the Hawk-Dove game in which two individuals are contesting ownership of a territory that one of them currently controls. This is a typical situation in the animal world, but, unfortunately, also in humans societies. In this scenario, players can have "owner" or "intruder" roles. Therefore, pure strategies are in the form of "play Hawk if owner, play Dove if intruder". Let's represent them with the set $S = HH, HD, DH, DD$, where the first H/D strategy refers to the owner role. For instance, the previous example corresponds to the pure strategy HD.

Let's assume that each pairwise contest involves one player in each role, and that each player has an equal chance of being an owner or an intruder.

1. Using the same notation adopted in the class, where the value of the game is v and the cost is c (for playing Hawk, while there's no cost for playing Dove), construct the payoff table for the asymmetric Hawk-Dove game.

2. Find the two symmetric pure Nash equilibria.

3. Are do they ESSs?

4. Set $v = 4$ and $c = 8$ and show that there is no mixed strategy ESS.

[1] It may seem strange, if not incorrect, to assume that an individual can change its role given that, for instance, any individual is usually either male or female throughout its entire life. However, in genetic terms, the genes that are assumed to control behavior will be passed down to offspring, that may be male or female with equal probability. Therefore, the notion of dynamic evolution can "adjust" the issue at the population level.

8 Cooperators, defectors, and tit-for-tatters (30 points)

In the class we have remarked that a cooperative behavior is not necessarily the equilibrium outcome of a game (conflict) scenario. However, cooperation does actually happen in both human and animal interactions. Evolutionary game theory can help to explain this fact by showing that equilibrium strategies that correspond to cooperation can be achieved through population's dynamics.

At this aim, let's consider the prisoners' dilemma, that is actually a dilemma about cooperating (C) or defecting (D). The considered instance of the game is represented in normal form in the payoff matrix below, where $b > a > 0$.

P_2

P_1		C	D
	C	$b-a, b-a$	$-a, b$
	D	$b, -a$	$0, 0$

In practice, it tells that if cooperate, each player can help the other, conferring a benefit $b > 0$ on the other player at a cost $a > 0$ to himself. In case of mutual defection, none of the two players earns a benefit. Defect is the strictly dominant strategy for both players: if both cooperate, both are better off than if both defect. If we form the replicator equation for the corresponding evolutionary game, we would see that the cooperators die out, and only defectors are left.

1. Justify (using the theorems introduced in the classes) why this is what it will happen.

Thus the benefits of cooperation alone are not sufficient to let cooperation arise in a population. Now let us imagine that when two random players from a population meet, they play the game *twice*, using one of three strategies:

- C: always cooperate;
- D: always defect;
- T (*tit-for-tat*): cooperate the first time; the second time, do what the other player did the first time (retaliation).

In practice, one game interaction comprises *two interactions in a sequence.*

2. Complete the payoff matrix for the game:

P_2

P_1		C	D	T
	C	$2b-2a, 2b-2a$		
	D			$b, -a$
	T		$-a, b$	

Let's assume $b > 2a$. Therefore, the game has two (symmetric) pure strategy Nash equilibria.

3. Find the two equilibria.

In order to simplify the study of the replicator system, let's consider the case $b = 3$ and $a = 1$. This game has three symmetric mixed Nash equilibria: $(D, D), (\frac{1}{2}D + \frac{1}{2}T, \frac{1}{2}D + \frac{1}{2}T)$, $(p_C C + p_T T, p_C C + p_T T)$ with $p_C + p_T = 1$ and $0 \le p_C \le 1/3$, where p_C and p_T are the probability of using, respectively, pure strategy p_C and p_T. The pure Nash equilibrium (T, T) is found for $p_C = 0$.

4. Write the equations for the replicator dynamics for the reduced state (p_C, p_D).

5. Study the resulting dynamical system: find equilibrium points, nullclines, show phase portrait, stability properties of the equilibrium points.

6. Identify the ESS equilibria (and tell why they are ESS).

7. Discuss the "social" meaning of the equilibria and of the evolutionary dynamics of the population.

9 Interaction networks and Cellular Automata (25 points)

The *network topology* that defines how individuals / components of a complex system interact plays a central role in the generation of (useful, interesting) emergent patterns and behaviors at the system level. The attached paper *The role of interaction network in the emergence of diversity of behavior* precisely addresses this issue. In particular, a *cellular automata* is optimized (using a genetic algorithm) to find the rules and the interconnection network that would optimize the solution of the (classical) *density classification task* (given that that $t = 0$ in a binary cellular automata there's a majority of 1s (0s), find the rule that "discovers" this by letting the CA converging to a state with all 1s (0s)).

1. Read the paper and make a summary and a critical review of it (about one page max).
2. Try to replicate *any* of the results shown in the paper using a PSO algorithm as an optimizer instead of the Genetic Algorithm. Justify your choice in terms of which experiments you find more interesting to explore and why.
3. If you are unable to implement the PSO algorithm and physically run the experiments, you can still provide a detailed description of the PSO algorithm that you "would" employ and of the experimental setting. Discuss which difficulties you would expect and what kind of hypothesis you'd like to investigate more in depth (and why).