Open-minded imitation in vaccination games and heuristic algorithms

Winfried Just
Department of Mathematics, Ohio University

Based on joint work with Ying Xin, Montana State University
and David Gerberry, Xavier University

Association for Computing Machinery
Ohio University
October 24, 2018

A conversation about games

Bob: What mathematical tools do you use in your research on human reactions to the spread of infectious diseases?

WJ: Game theory.

Alice: That sounds like a lot of fun!!!

Bob: (Laughs.) Yeah. But now, seriously?

WJ: Seriously.

Bob: How can you play games with epidemics, which can literally be matters of life and death?

WJ: Let's look at games from a mathematical perspective.

Alice: I play a lot of games! I can tell you all about games!!

Bob: Alice!

What did Mom tell you about conversations between grown-ups?

AI about games

WJ: Relax, Bob! Alice can tell us about games, and I will translate it for you into more abstract, or if you will, grown-up language.

Al: I play rock-paper-scissors, chess, poker, soltaire, monopoly, Rubik's cube ...

WJ: And you like playing all of them?

Al: Solitaire—not so much. And Rubik's cube even less.

WJ: And why not?

Al: Because what I really like is playing with other people. Now that I think about it, Solitaire and Rubik's cube perhaps aren't really games at all.

WJ: Would you then say that a "real" game involves interactions between two or more players?

Al: Yeah, this is what I meant.

WJ: We mathematicians can use such games as models for all sorts of real-world situations that involve interactions between people.

Payoffs

WJ: And what else do you like about games?

AI: Winning! I really like winning!! But I don't like losing.

WJ: Are winning or losing the only possible outcomes of a game?

AI: No. In poker or monopoly, you can win or lose by a lot.

WJ: So could we then always think of the outcome of a game for each player as a real numbers that represent that player's payoff?

AI: Maybe. ... But if the player loses some money, like in poker, wouldn't that be a cost?

WJ: We can treat costs as negative payoffs. And we can treat outcomes like "win," "lose," or "draw" as payoffs 1, 0, 0.5, for example.

Maximizing payoffs

WJ: What are the players after?

Al: Each player wants to make her payoff as large as possible.

Bob: Alice! Watch your language! What did Mom and I tell you?

Al: I meant to say, "her or his payoff." Would you mathematicians say: "Each player wants to maximize his or her payoff?"

WJ: Exactly. But your original phrase is more interesting. You said make. How do you "make" your payoff as large as possible? Can you explain it with an example?

Al: With the Rubik cube, for example, from each scrambled-up starting configuration I make the right sequence of moves to get to the winning configuration where each side has only one color. There is an algorithm for this. I even taught it to my computer!!

Strategies

WJ: In game theory, such an algorithm would be called a strategy. We use this word even when the algorithm prescribes moves that are not the best or the right ones for a given configuration.

When you play solitaire, do you also follow a strategy that makes your payoff as large as possible?

AI: Yes, but I don't always win, in the sense of putting all the cards on ordered stacks.
It all depends on how the cards are shuffled.

WJ: Yes, I know. The success of your strategy in this game depends on some random events.

Expected payoffs

WJ: Isn't it the case though that in some rounds of solitaire that you lost by random chance you would have won with a different strategy? So that, in a way, you regret your moves?

AI: Yes, this happens sometimes.

WJ: How then would your strategy maximize your payoff?

AI: I didn't say my strategy gives always the highest payoff.

WJ: So what did you mean then?

AI: If I play many, many rounds of solitaire I will have fewer such regrets than with any other strategy.

WJ: So you were talking about the average or expected payoff that you get when you average the payoffs for many rounds.

AI: Exactly! We should have said earlier that the goal of each player is to maximize her or his expected payoff.

On to the really fun games

WJ: You said that Rubik's cube and solitaire are not as much fun as "real" games, that is, games with other players.

Al: In real games my expected payoff, does not only depend on my moves, but also on the moves of all other players. The really fun thing about games is that each player tries to outsmart all the others so as to get maximum payoff for her- or himself.

Bob: That doesn't sound like a nice goal to strive for. Wouldn't it be better if players cooperated so as to ensure maximum overall expected payoff for the group?

WJ: We will talk about that later today.

But first let's go from "moves" to "strategies."

Are the really fun games really that much fun?

WJ: When you play chess, don't you also follow a strategy?

AI: Yes I do!

WJ: And wouldn't you assume that your opponent does the same?

AI: Sure. At least, if my opponent is rational and plays so as to maximize her or his expected payoff.

WJ: Wouldn't then the game boil down to pitting your strategy against the one of your opponent, and the better strategy would win or at least ensure a draw?

AI: Yeah, sort of. At least when the players are smart and don't make mistakes.

WJ: But then, once the strategies are chosen, we know the resulting payoffs, and the actual sequence of moves is no more interesting than when you play the Rubik cube by recipe.

Choosing strategies

Al: (With a lump in her throat.) But what happened to the fun?

WJ: The whole fun is in the player's choosing their strategies.

Al: How could that be the fun part? Wouldn't this have to happen before the players even start playing and make any moves?

WJ: Yes. . . . But can you think of a fun game that is practically over once the players choose their strategies?

Al: Rock-paper-scissors!! This is a lot of fun. You see, when I play it with my friend Ronnie, I choose paper, because he always chooses rock, when I play it with my friend Sequi, I choose rock, because she will choose scissors, and when I play it with my friend Paul, I choose scissors, because he prefers paper.

This way I outsmart them all!

WJ: All of them??

How about Maxi?

Al: Except for Maxi, who is really, really smart.

I think she must be able to read my mind, at least most of the time. But I cannot figure out hers.

And her real name isn't Maxi, but very weird.

WJ: Is it Maximinia?

Al: How did you guess that??? Do you know her?

WJ: Not personally. But in an abstract sort of way, I do.

Al: You mathematicians with your abstractions!

Maxi is a wonderful person and my bestest friend!!!

But how come I never can beat her in any game??

WJ: I will show you how she does it.

Let's go back to chess first.

Alice and Maxi play chess

WJ: Think of Maxi (white) choosing between 4 possible chess strategies SM_1, SM_2, SM_3, SM_4 and you (black) between 4 strategies SA_1, SA_2, SA_3, SA_4. When each of you follows her chosen strategy, the payoff vectors will be as in the table below.

Table: $(1,0)$—Maxi wins; $(0,1)$—Alice wins, $(0.5,0.5)$—draw.

	SA_1	SA_2	SA_3	SA_4
SM_1	$(0,1)$	$(0,1)$	$(0,1)$	$(0.5,0.5)$
SM_2	$(0,1)$	$(0.5,0.5)$	$(1,0)$	$(0.5,0.5)$
SM_3	$(1,0)$	$(1,0)$	$(0.5,0.5)$	$(0.5,0.5)$
SM_4	$(0,1)$	$(1,0)$	$(0.5,0.5)$	$(0.5,0.5)$

WJ: Which strategy would you pick?

Which strategy would you pick?

Table: $(1,0)$—Maxi wins; $(0,1)$—Alice wins, $(0.5,0.5)$—draw.

	SA_1	SA_2	SA_3	SA_4
SM_1	$(0,1)$	$(0,1)$	$(0,1)$	$(0.5,0.5)$
SM_2	$(0,1)$	$(0.5,0.5)$	$(1,0)$	$(0.5,0.5)$
SM_3	$(1,0)$	$(1,0)$	$(0.5,0.5)$	$(0.5,0.5)$
SM_4	$(0,1)$	$(1,0)$	$(0.5,0.5)$	$(0.5,0.5)$

Al: I would pick strategy SA_1 because it beats three of the four strategies of Maxi and maximizes my expected payoff.

WJ: But wouldn't your payoff depend on what Maxi does?

Al: Oh, yeah! She is really smart and can read my mind. So she would pick SM_3, and then I would lose and regret my choice. I really don't like losing, you know.

Which strategy would you pick, then?

Table: $(1,0)$—Maxi wins; $(0,1)$—Alice wins, $(0.5,0.5)$—draw.

	SA_1	SA_2	SA_3	SA_4
SM_1	$(0,1)$	$(0,1)$	$(0,1)$	$(0.5,0.5)$
SM_2	$(0,1)$	$(0.5,0.5)$	$(1,0)$	$(0.5,0.5)$
SM_3	$(1,0)$	$(1,0)$	$(0.5,0.5)$	$(0.5,0.5)$
SM_4	$(0,1)$	$(1,0)$	$(0.5,0.5)$	$(0.5,0.5)$

Al: OK, I will pick strategy SA_3 because then Maxi cannot beat me with strategy SM_3, and I can still win if she plays strategy SM_1. But that would be really dumb of her.

WJ: Wouldn't then Maxi regret her choice of SM_3 and play SM_2 instead?

Al: Yes, of course! She is sooo smart and can read my mind. And she too likes winning. So I would lose again.

Which strategy should you pick, then?

Table: (1, 0)—Maxi wins; (0, 1)—Alice wins, (0.5, 0.5)—draw.

	SA_1	SA_2	SA_3	SA_4
SM_1	(0, 1)	(0, 1)	(0, 1)	(0.5, 0.5)
SM_2	(0, 1)	(0.5, 0.5)	(1, 0)	(0.5, 0.5)
SM_3	(1, 0)	(1, 0)	(0.5, 0.5)	(0.5, 0.5)
SM_4	(0, 1)	(1, 0)	(0.5, 0.5)	(0.5, 0.5)

AI: Oh, I see!! I need to pick strategy SA_4! Then I will have no regrets when Maxi picks SA_3, and Maxi will have no regrets either in this case, because she could do no better with any other of her strategies against my SA_4.

WJ: A choice of strategies where no player has any regrets about his or her choice given the choices of all other players is called Nash equilibrium.

AI: Just to make sure: By "no regrets" you mean that no player could achieve a higher expected payoff by switching to another strategy?

WJ: Exactly.

AI: But tell me, can Maxi really read minds?

Can Maxi really read minds?

WJ: Yes, at least minds of very, very smart people like Alice and herself. She assumes that all players are perfectly rational and as smart as she is and will only pick strategies from a Nash equilibrium. Then so does she.

Al: How can she reason all that out?

Table: (1, 0)—Maxi wins; (0, 1)—Alice wins, (0.5, 0.5)—draw.

	SA_1	SA_2	SA_3	SA_4
SM_1	(0, 1)	(0, 1)	(0, 1)	(0.5, 0.5)
SM_2	(0, 1)	(0.5, 0.5)	(1, 0)	(0.5, 0.5)
SM_3	(1, 0)	(1, 0)	(0.5, 0.5)	**(0.5, 0.5)**
SM_4	(0, 1)	(1, 0)	(0.5, 0.5)	(0.5, 0.5)

WJ: Maximinia may simply look at her minimum payoff in each row of the payoff matrix and then choose a strategy for which this minimum in the corresponding row is maximal.

Al: Ah! That's how you knew her! But if there is no Nash equlibrium?

WJ: Every game has at least one Nash equilibrium.

How about rock-paper-scissors?

Al: But that's not true!! Not when I play rock-paper-scissors with Maxi! Look at the payoff matrix:

Table: (1, 0)—Maxi wins; (0, 1)—Alice wins, (0.5, 0.5)—draw.

	RA	*PA*	*SA*
RM	(0.5, 0.5)	(0, 1)	(1, 0)
PM	(1, 0)	(0.5, 0.5)	(0, 1)
SM	(0, 1)	(1, 0)	(0.5, 0.5)

Al: You see, Maxi's minimum for each row is zero.

WJ: So Maxi doesn't have a strategy in your table that she couldn't possibly regret,

Al: and neither have I. So there is no Nash equilibrium.

WJ: Not if we consider only the strategies in this table.

Al: But by the definition of the rock-paper-scissors game, as you mathematicians would say, the table lists all possible strategies!!

Mixed strategies

WJ: You said that you can outsmart Ronnie, Sequi, and Paul, because you know their favorite strategies.

But what if a player rolls a fair die and plays R if it comes up 1 or 2, S if it comes up 3 or 4, and P if it comes up 5 or 6?

This is an example of a mixed strategy. The table on the previous slide lists only the so-called pure strategies.

Al: So then the expected payoff of **any** strategy against D is 0.5!

WJ: Would Maxi perhaps play D?

Al: Oh no!! Maxi is way too smart for that. She can read my mind. If I play S, she would outsmart me, play R, and win! And if I play P . . .

WJ: Couldn't you prevent her from outsmarting you by playing D yourself?

Al: Yes!!! She can read my mind, but not the die's mind!!

Nash equilibrium for rock-paper-scissors

WJ: Look at the payoff matrix when we include D:

Table: (Maxi's payoff, Alice's payoff).

	RA	PA	SA	DA
RM	$(0.5, 0.5)$	$(0, 1)$	$(1, 0)$	$(0.5, 0.5)$
PM	$(1, 0)$	$(0.5, 0.5)$	$(0, 1)$	$(0.5, 0.5)$
SM	$(0, 1)$	$(1, 0)$	$(0.5, 0.5)$	$(0.5, 0.5)$
DM	$(0.5, 0.5)$	$(0.5, 0.5)$	$(0.5, 0.5)$	$(0.5, 0.5)$

AI: So when we both play D, neither of us has any regrets,

WJ: and (D, D) is a mixed strategy Nash equilibrium.

Alice wants to know more

Alice: This sounds so cool!!! Can you show me how you use this in your research?

WJ: Not now. In a few years, if you decide to study math . . .

Al: But I want to become a game designer!!

WJ: And for that a math degree is very useful.

Bob: (Alarmed) Listen Alice: Mom and I love you very much and want you to have a great career in business or medicine.

Alice and I need to go now. Thank you for the enlightening conversation!

Al: Oh-ok! Bye for now, but let's get together after your talk and play some games!

WJ: Deal!

Another conversation: Two people met by accident ...

WJ: Hi! I'm Vinny!

YX: Hi! I'm Ying Xin, a Ph.D. student in mathematics.

WJ: Oh no! I'm terrible at math!

YX: I find that hard to believe.

WJ: But so glad we met. I have a problem that you as a mathematician might be able to help me with.

YX: Would be happy to.

WJ: You see, I am trying to decide whether or not to get a flu shot this year, and I'm trying to make a rational decision. How would you mathematicians approach this problem?

Costs

YX: Well, there is a cost of vaccination.

WJ: You mean, like, waiting in line, getting poked with a needle ... and there might be nasty side effects. But perhaps I will not have to wait in line or suffer any side effects?

YX: Let's consider the average cost, and let c_v denote it.

WJ: Catching the flu is a lot nastier though than getting a flu shot.

YX: You are saying that the (average) cost of infection c_i is a lot larger than c_v, that is, $c_i \gg c_v > 0$.

Probability of infection and vaccine efficacy

YX: Do you always catch the flu when you don't get vaccinated?

WJ: No, last year I remained unvaccinated and did not catch the flu.

YX: Let x denote the probability that an unvaccinated person will catch the flu. This probability will depend on the vaccination coverage V, so that $0 \leq x(V) < 1$.

YX: And is it true that a person who does get vaccinated never catches the flu?

WJ: No!!! Two years ago I did get vaccinated. And then I caught the flu nevertheless. That was really bad.

YX: I am sorry to hear this. So we might need to consider another parameter rE, called the efficacy of the vaccine. For an ideal vaccine, we would have $rE = 1$. In general, let us assume here that the probability of a vaccinated person catching the flu is $(1 - rE)x$.

Expected costs for vaccinators and nonvaccinators

YX: Now you can calculate your expected costs when you vaccinate and when you don't vaccinate.

WJ: If I don't vaccinate, my expected cost will be

$C_u = C_u(V) = c_i x(V)$,

and if I vaccinate my expected cost will be

$C_v = C_v(V) = c_v + c_i(1 - rE)x(V)$.

So since $c_i \gg c_v$, the cost for not vaccinating will be higher and everybody should vaccinate!

Herd immunity

YX: Not necessarily. When rE is not too small, then there exists a vaccination coverage $V_{hit} < 1$, called the "herd immunity threshold," such that for all $V \geq V_{hit}$ we have $x(V) = 0$.

WJ: Great! So then it would suffice to vaccinate a proportion of $V_{hit} < 1$ of the population to provide perfect protection for all. We could then save the cost of vaccinating a proportion of $1 - V_{hit}$ of the population.

YX: Exactly! So it would not be necessary or optimal for **everybody** to get vaccinated.

But who should get vaccinated?

WJ: But who should and who shouldn't get vaccinated?

And should I or shouldn't I? This is exactly my dilemma.

YX: If the government were to draw up a list ...

WJ: You must be kidding ... How could you trust **them** with a problem of minimizing costs??

I don't want no government making health care decisions for me.

YX: Who should decide then?

WJ: We, the people. Like you and me. By making rational decisions as individuals, we will arrive at the vaccination coverage V_{hit} that's best for everybody.

Vaccination games

YX: How would this work?

WJ: Can't we treat this situation as a game where each player has two pure strategies: To vaccinate, or not to vaccinate. We can then treat the costs as negative payoffs.

YX: You have outlined mathematical models called vaccination games.

WJ: So when each player makes a rational decisions, the population will then reach a mixed-strategy Nash equilibrium with vaccination coverage V_{Nash} where the payoffs $-C_v(V_{Nash})$ and $-C_u(V_{Nash})$ are equal, where no player has any regrets whatsoever about their strategies no incentive to switch to another strategy, right?

Nash equilibria *vs.* societal optimum

YX: Yes, this is what models of the vaccination game predict.

WJ: Beautiful! So perfectly rational people will arrive at a no-regrets-whatsoever situation with optimal vaccination coverage as in your herd immunity threshold by just making rational decisions on how to randomize their individual vaccination choices. No government meddling required!

YX: You are assuming here that the vaccination coverage at Nash equilibrium is optimal and is equal to the herd immunity threshold. But this is not true.

WJ: Now give me a break: Doesn't "optimal" mean the same thing as "no regrets whatsoever?"

Is $V_{Nash} = V_{hit}$?

YX: We are talking about different types of regrets. At Nash equilibrium, nobody has any regrets about their **individual** decision. At V_{hit}, we have no regrets about the cost to society as a whole.

WJ: So how could what's best for each of us individually not be best for all us?

YX: It's possible.

What would be the rational choice for you, and thus for everybody, when $V = V_{hit}$ so that $x(V) = 0$?

WJ: Then $C_u(V) = c_i x(V) = 0$ for an unvaccinated person, and $C_v(V) = c_v + c_i(1 - rE)x(V) = c_v$ for a vaccinated person. Thus $C_v(V_{hit}) > C_u(V_{hit}) = 0$. So $V_{Nash} < V_{hit}$.

What's to be done about it?

YX: In other words, individually optimal decisions lead to a suboptimal outcome for the whole society.

WJ: Bummer! Anything you mathematicians can do about it?

YX: That would take an effort of the whole society. As mathematicians, we can only carefully study whether our models are realistic and make accurate predictions.

First, we need to carefully check our assumptions. The prediction of a Nash equilibrium is based on the idea that everybody makes perfectly rational decisions.

WJ: Are you saying that since most people aren't all that smart, there is some hope?

YX: I would not say it this way. But as a society we could help people in making more beneficial decisions.

How do real people make decisions?

WJ: OK, but what I meant was this: People like me wouldn't even know how to make the calculation for the Nash equilibrium.

YX: So how do you usually arrive at your vaccination decisions?

WJ: You didn't notice? I might ask an expert, like you.

YX: Well, thank you, but ...

WJ: And if I hadn't met you by accident, I would ask my friend George how things went for him last year. If what he did worked reasonably well, I might do the same this year.

YX: So you might then imitate George's strategy.

Imitation of good decisions

WJ: You can call it this way. It think if people were to imitate good decisions of other people, that would lead to better outcomes for the society as a whole.

YX: This conjecture has been widely studied.

WJ: So what have these studies found?

YX: The literature reports that when $c_i > 2c_v$ the population will always arrive at a vaccination coverage that is even lower than V_{Nash}, with an even higher cost to the overall population.

WJ: Bummer again!

YX: But this may be an artifact of how the process of imitation is conceptualized.

WJ: What do you mean?

Imitating George

YX: I will tell you, but first explain to me how, exactly, you imitate your friend's George's strategy.

WJ: Well, most of the time, I would just do what I did last year. But once in a while, I would ask George what his cost was last year. If it was lower or at least in the same ballpark, I would most likely switch to his strategy. But if his cost was a lot higher than mine, I will most likely stick to my own previous strategy.

YX: You said "most likely." So: Not always?

WJ: Yeah. He or I might have just lucked out with not getting vaccinated.

Fermi functions

In the literature the probability p_{switch} of switching is usually modeled by a so-called Fermi function:

$$p_{switch} = \frac{1}{1 + e^{-\beta(C(\text{your strategy}) - C(\text{other}))}},$$

with $\beta > 0$. When β gets larger, this becomes closer to a best-response function.

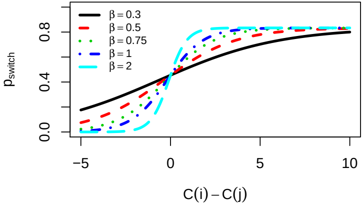

Are Fermi functions realistic?

$$p_{switch} = \frac{1}{1 + e^{-\beta(C(\text{your strategy}) - C(\text{other}))}}$$

WJ: But wait! When my strategy has a higher cost, then this model predicts that I would switch to the other with probability > 0.5. This isn't what I do. Most of the time I just stick with my strategy for the previous year.

YX: My Ph.D. advisor Prof. Just noticed the same thing.

WJ: I know this guy! He is your advisor? I could tell you stories ...

YX: Maybe not now ...

WJ: OK. Some other time. So what happened next?

Some empirical research in the psychological literature supports more flexible functional forms of the switching probabilities. We generalized the Fermi function by introducing a parameter $\alpha \geq 1$ so that:

$$
\begin{aligned}
p_{switch} &= \frac{1}{\alpha + e^{-\beta(C(\text{your strategy}) - C(\text{other}))}} \\
&= \frac{\alpha^{-1}}{1 + \alpha^{-1} e^{-\beta(C(your\ \ strategy) - C(other))}}.
\end{aligned}
$$

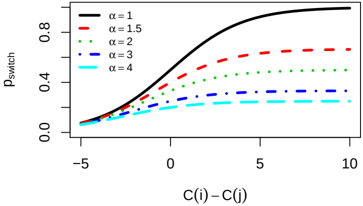

Generalized Fermi functions, continued

WJ: Would that be closer to my way of imitating?

YX: You can think of α^{-1} as the probability of making your decision by imitation in a given year. If α is large, then p_{switch} would always be close to 0.

WJ: But if I do consider imitating somebody else, then my switching probability is close to 1, unless that other person did really poorly.

YX: So you would be open-minded about trying out the other's strategy.

WJ: You can call it this way.

YX: Recall that:

$$p_{switch} = \left(\frac{1}{\alpha}\right)\left(\frac{1}{1+\alpha^{-1}e^{-\beta(C(your\ \ strategy)-C(other))}}\right).$$

The following figure shows how the second fraction, which represents the conditional switching probability, depends on α.

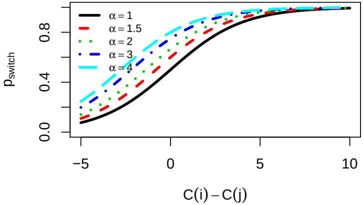

Vaccination games with generalized Fermi functions

WJ: But would that α make any difference in vaccination games?

YX: Yes. We found that when $\alpha = 1$, as in the previously published papers, then the population converges to a vaccination coverage $V^* < V_{Nash}$, with expected overall cost $C(V^*) > C(V_{Nash})$.
However, for sufficiently large values of α we found many parameter settings where $V_{hit} \approx V^* > V_{Nash}$, with $C(V_{hit}) \approx C(V^*) < C(V_{Nash})$.

WJ: So the degree of open-mindedness of the imitation, if you will, does really make a big difference then?

YX: Yes it does. We confirmed this both with simulations and by proving rigorous theorems.

Time to say good-bye

WJ: Theorems? Er ...

Thoroughly enjoyed our conversation, but gotta go now.

YX: Same here. Thoroughly enjoyed our conversation, but need to go now on a long trip to Montana.

WJ: Wow! To do mountain climbing?

YX: No, work on a postdoc project.

WJ: Well, good luck with whatever you are going to do over there!

But now tell me: Should I get that flu shot or not?

XY: You should keep an open mind about it.

On the characters of these plays

Bob and Alice are purely fictional characters.

Neither of them is based on any one actual person.

My former Ph.D. student Ying Xin is both a protagonist and a co-author of an earlier version of the second conversation.

The character of WJ . . .

Alice: We know. Please get on with the third part of your talk so that you can get out of here and play with me!

A variant of evolutionary computation

Consider the following outline of an evolutionary algorithm:

- Initialize the population with agents.
- Until stopping criterium holds, repeat:
  1. Evaluate fitness $f(i)$ of each agent i.
  2. For each agent i, randomly choose another agent j. Replace agent i with agent j with probability p_{switch} that depends on the fitnesses of agents i and agent j. Make all the required comparisons simultaneously, before any actual replacement takes place.
  3. Modify agents by mutations and/or crossover (optional).

Is there an official name for this type of evolutionary algorithms?

Some choices p_{switch} in Step 2

1. Evaluate fitness $f(i)$ of each agent i.
2. For each agent i, randomly choose another agent j. Replace agent i with agent j with probability p_{switch} that depends on the fitnesses of agents i and agent j. Make all the required comparisons simultaneously, before any actual replacement takes place.
3. Modify agents by mutations and/or crossover (optional).

- Switch if, and only if, $f(j) > f(i)$ ("best response").
- Make p_{switch} proportional to $f(j) - f(i) + c$ for some constant c ("Darwinian." Feasible if there are bounds on $f(i)$).
- Make $p_{switch} = \frac{1}{\alpha + e^{-\beta(f(j)-f(i))}}$ for some $\alpha \geq 1, \beta > 0$.

Where have we seen this before?

- Initialize the population with agents that vaccinate or don't.
- Until stopping criterium holds, repeat:
  1. Evaluate fitness $f(i) = -C(i)$ of each agent i, where $C(i)$ is the cost for this agent in the previous flu season.
  2. For each agent i, randomly choose another agent j. Replace agent i with agent j with probability $p_{switch} = \frac{1}{\alpha + e^{-\beta(f(j)-f(i))}}$ for some $\alpha \geq 1, \beta > 0$.
     Make all the required comparisons simultaneously, before any actual replacement takes place.
  3. Skip modification of agents.

Essentially, our results show that for sufficiently high values of the "open-mindedness" parameter α, this algorithm "computes" near-optimal vaccination coverage V^*.

In contrast, if we use a "Darwinian" version of p_{switch}, then the algorithm "computes" V_{Nash}.

Open-mindedness in evolutionary computation?

Open problem: For what kind of problems would the following general type of algorithm outperform comparable ones when α is suitably large?

- Initialize the population with agents.
- Until stopping criterium holds, repeat:
  1. Evaluate fitness $f(i)$ of each agent i.
  2. For each agent i, randomly choose another agent j. Replace agent i with agent j with probability $p_{switch} = \frac{1}{\alpha + e^{-\beta(f(j)-f(i))}}$ for some $\alpha \geq 1, \beta > 0$. Make all the required comparisons simultaneously, before any actual replacement takes place.
  3. Modify agents by mutations and/or crossover (optional).

Why does this seem a promising question?

In evolutionary computation, finding an optimal balance between exploration and exploitation is of paramount importance. "Open-mindedness" can be interpreted in such a way that provided an agent i considers switching at all, the conditional switching probability will remain close to 1 even if $f(i)$ is slightly larger than $f(j)$. Thus high values of the parameter α will shift this balance towards exploration. This may work well for certain types of fitness landscapes.

In contrast, the "best response" version of p_{switch} would shift this balance towards exploitation. One might think of versions of the algorithms that have mixtures of these two versions of p_{switch} whose proportions change over time.

High values of α also will slow down approach to the equilibrium. At least in our model, this can be remedied by introducing a scaling factor that does not qualitatively alter the fitness landscape, in particular, that preserves location of the equilibria.

A couple of references

The results presented her are proved in:

Y. Xin, D. Gerberry, and W. Just (2018); Open-minded imitation can achieve near-optimal vaccination coverage. *arXiv:1808.08789*
`https://arxiv.org/abs/1808.08789`

Our model is based on the model of:

F. Fu, D. I. Rosenbloom, L. Wang and M. A. Nowak (2011); Imitation dynamics of vaccination behaviour on social networks. *Proc. R. Soc. B* **278** 42–49 doi:10.1098/rspb.2010.1107