Assignment 2: Cooperative Games

CS886 - Fall 2021

Out: 5 Nov 2021
Due: 25 Nov 2021

In this assignment, you will investigate the experimental multi-player cooperative game called "Trust-ya" Huang et al. (2021), and will design an automated player for this game. **Submission Material:** A single page write-up describing your approach for the design of an automated Trust-ya player. **Submit on LEARN to the A2 dropbox**. You are also expected to play the game (yourself) on Nov 11th (see below). There is a **cash reward**!

Trust-ya

Trust-ya is a game in which 3 or more players aim to maximize their personal stash of coins drawn from a large central pile. Although they are in competition for these coins, players must cooperate to draw as many as possible from the central pile before the game ends.

Question I (6 pts)

Steps:

a. Read the paper on Trust-ya Huang et al. (2021) – see the course webpage or get it at `https://ieee-cog.org/2021/assets/papers/paper_294.pdf`

b. On November 11th during class time (230pm-400pm)

   (a) meet online
   `https://uwaterloo.zoom.us/j/94076554629?pwd=WHB3RzBCQ1ZscTFxY0srT25UTm5NUT09`
   (Meeting ID: 940 7655 4629; Passcode: WsXK0b)

   (b) You will be split into three groups of 8 people randomly, and each group will play two games of Trust-ya. **No communication is allowed with other players except through the game interface.**[1] Players in Trust-ya are given anonymous names, so communicating outside the game will be difficult in any case. Your objective will be to end up with as much as possible in your "Savings Pile" in each game. In Trust-ya, the objective is individual: you are trying to get as much as possible for yourself only!

---

[1]The Zoom link will remain on but all video and audio will be muted. Any player experiencing technical problems can unmute to alert other the moderator.

(c) the game has some slight differences with that in the paper:

  i. Each game will last for an pre-determined but unknown (to you) number of rounds. It will be more than 10 but less than 30. Each team will play the same number of rounds in total in each game, however, to ensure the chances of winning are the same. The central pile is infinite for all games.
  ii. The winnings (equation 1 in paper) are the same, with the card values for Jack, Queen, and King being 11, 50, and 100, respectively.
  iii. The loss for a face card without p-card (equation 2 in paper) is slightly different. The first term, instead of (savings_pile+round_pile-investment), is just (savings_pile+round_pile). The second term, instead of (investment / round_pile), is (investment / (round_pile + investment)).
  iv. The cards in the deck are the standard for a 52-card deck. Cards are always drawn with replacement.
  v. p-cards are kept if a player has sufficient investors (for games with 8 players):
     - Jack: At least one investor (not including self)
     - Queen: At least two investors (not including self)
     - King: At least four investors (not including self)

     otherwise they are lost, but may be re-purchased
  vi. a player's savings pile cannot go negative

(d) We will randomly select eight participants from all three groups, weighted by each player's sum of savings over both games. Thus, if Bob and Alice played Trust-ya and Bob ended up with 20 total (sum of his two savings piles), while Alice ended up with 180 total, then a random number between 1 and 10 would be drawn and Bob would be selected if it was a 1, while Alice would be selected if it was any other number.

(e) The eight selected participants will then play a single, final game, and a **Grand and Glorious Winner** will be selected randomly in the same manner based on the eight player's savings pile *for the single final game only.* The **Grand and Glorious Winner** will be awarded a **fabulous prize** of (a **$50 Amazon Gift Card**). Note that all 8 players advancing from the first round *could come from the same group of 8 players.*

(f) As you are playing these games, think about how you would design a bot to play the game (see next step below)

c. Design an agent to play Trust-ya. Your agent must implement the following functions, but you will not write these in code, only by summarizing them in your 400-word write-up (see below).

  (a) `move()`: First state in the match state machine, when all players must decide whether to give or take coins. This function returns the identity of another player to whom to give (this can be your own player, in the case of taking).

(b) `invest()`: Deciding how much money to invest in this round from the amount you have received in your round pile. This returns any amount from 0 to the total of your round pile.

(c) `share_winnings()`: Decide how much money should be shared each investor. This returns a list of amounts for each other player including yourself, which sum to the total amount you won in this round.

(d) `purchase()`: Make purchases. Returns a list of emojis/pcards to purchase. Your savings pile is reduced by the amount it costs, and you can't buy more than what you have in your savings pile.

d. You can assume you have access to various game elements you may need, basically anything that is available to you when you play the game in person (names, emoji of other players, your player's round pile, your emoji and p-cards, your clients (people who donated to you), what you know about what others have won, whether a new round has started, which card was drawn, or whether the game has ended, etc...). You can also use any elements of private state you want so long as you explain what these are (e.g. you may want to store the history of your interactions so you can refer back to it later, or you might want to store a profile of some kind for each other player).

e. Write a maximum 400-word summary in plain language describing your agent.

What your agent needs to do:

As in the Trust-ya paper Huang et al. (2021), the objective of a good agent is to help people be more cooperative. That is, we would like to design agents to play this game such that the outcome of such an agent playing *with a group of humans who are self-interested* will be **both equitable and valuable**. As in Figure 2 of the paper, we want outcomes in the top left corner, where the final distribution is simultaneously as large and as fair as possible. Therefore, the best possible outcome is one where the central pile is completely distributed by the end of the game, and each player (including the agent itself!) gets the same amount. Your summary should explain why your bot will help "steer" the group towards such solutions. This is a difficult question to answer as you don't know what the human group will look like, not even what cultural group they are from! You can consider the bots we wrote about in the paper, but remember these agents *were only tested in homogeneous groups* -i.e., all agents were the same. While this may give you some insights into your agent, it may not generalize to heterogeneous groups of humans. This is a short summary, so **be bold** - don't be afraid to make wild claims about your agent!

Reference:

Huang, J., Jung, J., Budnarain, N., McGregor, B., and Hoey, J. (2021). Trust-ya: design of a multiplayer game for the study of small group processes. In *Proc. IEEE Conference on Games*.

Question II (4 pts)

What is the decision-theoretically optimal first move in *Trust-ya* as defined above?[2] There are two possible moves: take 2 for oneself, which yields 2 with probability 1, but may also yield more (payoff U_t) if someone else gives to you and you invest; and give 2 to someone else, which yields some payoff U_g. U_g and U_t are dependent both on equation (1) in the paper, and on how you think the other players are gong to approach their first moves. Therefore, you must have a model of how the other players are going to play their first move. You must define this model. You can assume you are playing a game with only one round (your payoff is whatever you get after the first round).

If you make approximations to solve this problem, clearly state what they are in your calculation. For example, you might want to assume that all other players will be exactly the same as you. Note that making this assumptions will likely invalidate the usage of this optimal policy as a reasonable thing to do in practice in the game.

Make your best attempt to solve this problem. I have worked out an approximate solution which is not difficult to find, but does require making some assumptions. I am therefore looking for novel and creative ways of approaching this problem, as well as optimal solutions. I will reward both such attempts equally.

What to hand in: `{Studentid}_CS886_A2.pdf` in a single PDF file

- 400-word summary of your agent in Question I,
- an answer to Question II.

You will be graded on this assignment as follows (out of 10):

2pts played at least 2 games of Trust-ya on Nov 11th

4pts 400-word summary in plain language of your agent's workings. Your write up must specify what each of the 4 functions above do (move, share_winnings, purchase and invest), and how they fit together. **Your write-up must justify why you think your method will be the best (as described above) at playing the game.**

4pts for either a creative and novel proposal for how to compute the optimal first move in Trust-ya, or a (close-to) optimal estimate of the expected value of the first move, and what the optimal move is.

[2] The decision-theoretic optimal is computed by computing the expected value of each action. Figure out how many outcomes there are for each of your actions, then figure out how likely each outcome is, and how much it will be worth to you. Multiply these two together and sum over all outcomes to figure out your expectation in value for each action.