

A Socially Aware Reinforcement Learning Agent for The Single Track Road Problem

Ido Shapira and Amos Azaria

Computer Science Department, Ariel University, Israel

Abstract

We present the single track road problem. In this problem two agents face each-other at opposite positions of a road that can only have one agent pass at a time. We focus on the scenario in which one agent is human, while the other is an autonomous agent. We run experiments with human subjects in a simple grid domain, which simulates the single track road problem. We show that when data is limited, building an accurate human model is very challenging, and that a reinforcement learning agent, which is based on this data, does not perform well in practice. However, we show that an agent that tries to maximize a linear combination of the human's utility and its own utility, achieves a high score, and significantly outperforms other baselines, including an agent that tries to maximize only its own utility.

Introduction

While humans can cope with new situations quite easily, even state-of-the-art algorithms struggle with new situations that they haven't been trained on. Unfortunately, when it comes to autonomous vehicles the results may be devastating. One example for an uncommon, yet important scenario for autonomous vehicles is the problem of a single track road. In this problem two vehicles in opposite directions must cross a narrow road, which is not wide enough to allow both vehicles to pass at the same time. Therefore, one vehicle must deviate from the road and let the other vehicle cross. However, if both vehicles deviate to the margins, they might both return to the road, and may either end-up deviating to the margins again or even colliding with each-other. Despite only a small portion of the roads being single track roads, autonomous vehicles must be able to function properly in these types of roads. Furthermore, some more common situations resemble the single track road problem, for example, if cars park where they shouldn't and block one of the lanes or if one lane is blocked for any other reason (e.g., a falling tree), the traffic in both ways must operate with a single lane.

In this paper we model the single track road problem as a sequential two player game on a two row grid (see Figure 1). The upper row represents a road that allows both players to advance. However, the lower row can only be used



Figure 1: The initial state of the single road game board. The red circle is controlled by the human player and the blue circle is controlled by the autonomous agent. Both players must reach the opposite side of the board without colliding. The players may travel freely on the upper row, but they cannot advance when located on the lower row.

for allowing the other player to pass, as the players cannot advance when placed in the lower row. We find several equilibria of the game, which should determine how a perfectly rational agent should behave in such a game. However, people tend to deviate from what is considered rational behavior, since they are influenced by different effects including anchoring, inconsistency of utility and a lack of understanding of other agent's behavior (Tversky and Kahneman 1981; Ariely, Loewenstein, and Prelec 2003; Camerer 2003). Indeed, as we later show, while some people tend to follow the game theoretic solution, many others do not follow it, and behave unexpectedly. Due to non-perfectly rational behavior of humans, algorithmic approaches that assume rational behavior tend to perform poorly with humans (Bitan et al. 2013; Azaria et al. 2015; Nay and Vorobeychik 2016).

Therefore, a common approach for developing an agent that can proficiently interact with humans is composed of several stages (Azaria et al. 2011; Nguyen et al. 2013; Rosenfeld et al. 2015). The first stage includes the collection of a data-set of humans interacting in the environment. Next, based on the collected data-set a human behavior model is developed, usually by applying machine-learning techniques. Finally, the human model is used by the agent to determine the actions that are the most beneficial for it. In this paper we attempt to follow this common practice for the single track road game. Therefore, we collect human data in this game and use it to compose a human model. Then, we model the agent's problem as a reinforcement learning environment by an MDP with the human model being a part of the environment. Finally, we use value iteration, a dynamic programming based method, to find the supposedly optimal action for the agent. We note that the solution provided by

value iteration is guaranteed to be optimal under the assumption that the MDP models the environment perfectly, which includes the human model.

However, composing a human behavior model based on a relatively small data-set may be inaccurate, as people are many times unpredictable and different humans tend to behave very differently from one another, despite a game being relatively simple (Shvartzon et al. 2016; Azaria, Richardson, and Rosenfeld 2016). Therefore, we introduce a novel method for solving an MDP that is based on a non-perfect human behavior model. Namely, instead of solving the MDP with an attempt to maximize the agent’s outcome, we propose to maximize a linear combination of the agent’s outcome and the human’s outcome. We expect that optimizing toward a linear combination will be beneficial for the agent, since the humans are likely to try and optimize their own utility function, so they are likely to deviate from the human model in a way that will indeed maximize their utility function. By optimizing toward a linear combination, the agent acts as if it already accounts for these deviations and is therefore more likely to adapt to them. We provide a formula for determining the proposed linear combination, which is based on the similarity of the agents’ utility functions.

We introduce our Socially Aware Reinforcement Learning agent (SARL), an agent that attempts to maximize the linear combination of the two utility functions, using our proposed formula. We show that SARL significantly outperforms all other baselines in the single track road game when interacting with humans, in terms of the agent’s final outcome. Somewhat less surprisingly, humans interacting with SARL also achieve the highest outcome. Therefore, SARL not only perform better with respect to its own outcome, but also with respect to social welfare.

To summarize, the contributions of this paper are three-fold:

1. We present the single track road problem, model it as a sequential game, and present the equilibria of the game. We show that people do not follow strategies that are in equilibrium.
2. We model the problem as an MDP in which the human’s actions are modeled as a part of the environment. The model uses data from humans interacting with simple agents to determine the probability of the human taking each action at a given state.
3. We present SARL, a socially aware reinforcement learner, that uses a linear combination of the rewards of both agents. We provide a formula for finding the parameter to be used in this linear combination. Finally, we show that our method significantly outperforms all other baselines.

Related Work

Trajectory prediction of surrounding vehicles and pedestrians is very important for the development of autonomous vehicles, as such knowledge can prevent accidents. Indeed, trajectory prediction is challenging due to the unexpected nature of human behavior. Therefore, many works attempt

to find a sufficient solution to overcome this challenge (Leon and Gavrilescu 2021).

According to Houenou et al. (Houenou et al. 2013) trajectory prediction can be based on a deterministic method that selects the current maneuver from a predefined set using kinematic measurements and road geometry detection. The authors state that their model cannot be applied to very low speed scenarios and therefore is not applicable to our scenario.

Deo and Trivedi (Deo and Trivedi 2018) estimate a probability distribution of future positions of a vehicle conditioned on its track history and the track histories of vehicles around it, at a certain time. Using this information, they select one of six possible maneuvers that have been defined. They use the publicly available NGSIM US-101 and I-80 highway datasets for their experiments. Their model relies purely on vehicle tracks to infer maneuver classes and ignores the lanes and the map.

Ding et al. (Ding, Chen, and Shen 2019) use a *recurrent neural network* for composing an observation encoding. Based on this encoding, they propose a *Vehicle Behavior Interaction Network* (VBIN) to capture the social effect of another agent on the prediction target, based on their maneuver features and relative dynamics (e.g., relative positions and velocities). VBIN is an end-to-end trainable framework and is suitable for dynamic driving scenarios where the dynamics of the agents affect their importance in social interactions. They use data collected from highways US-101 and I-80 as used in other work (Deo and Trivedi 2018); since it deals only with highway roads with a large number of agents, it is not applicable for our setting.

Kim et al. (Kim et al. 2017) propose a deep learning approach for trajectory prediction based on a *Long Short Term Memory* (LSTM). Their model is used to analyze the temporal behavior and predict the future coordinates of the surrounding vehicles. Based on the coordinates and velocities of the surrounding vehicles, the vehicle’s future location is produced after a certain short amount of time. However, the experiments were conducted using data collected from highway driving, which is again not suitable to our case.

Chandra et al. (Chandra et al. 2020) present an approach for trajectory prediction in urban traffic scenarios using a two-stream graph-LSTM network. The first stream predicts the trajectories, while the second stream predicts the behavior (i.e. overspeeding, underspeeding, or neutral) of road-agents. It is based on the vehicle coordinates and a weighted dynamic geometric graph (DGG) that represents the relative proximity among road agents. They also present a rule-based behavior prediction algorithm to forecast whether a road agent is overspeeding (aggressive), underspeeding (conservative), or neutral, based on the traffic behavior classification from the psychology literature. They evaluate their approach on the Argoverse, Lyft, Apolloscape, and NGSIM datasets and highlight the benefits over prior trajectory prediction methods.

Elhenawy et al. (Elhenawy et al. 2015) introduce a real time game-theory-based algorithm that is inspired by the *chicken-game* for controlling autonomous vehicle movements at uncontrolled intersections. They assume that all

vehicles communicate to a central management center in the intersection to report their speed, location and direction. The intersection management center uses the information from all vehicles approaching the intersection and decides which action each vehicle will take. They further assume that vehicles obey the *Nash-equilibrium* solution of the game and will take the action received from the management center. Unfortunately, these assumptions are very strong and cannot be applied to our setting.

Camara et al. (Camara et al. 2018) suggest a more realistic game-theory model based on the *sequential chicken-game*. The model assumes both agents share the same parameters U_{crash} and U_{time} , both know this is the case, and both play optimally from their state. It assumes that no lateral motion is permitted, and that there is no communication between the agents other than seeing each other's positions. The sequential chicken-game can be viewed as a sequence of one-shot (sub-)games, which can be solved similarly. The sub-game at time t can be written as a standard game theory matrix, which can be solved using recursion, game theory, and equilibrium selection to give values and optimal strategies at every state. While they handle the case of a junction by finding a Nash equilibrium and assuming that humans obey it, we deal with the single track road and give not only a game-theory analysis but also provide a novel Reinforcement Learning solution that does not require assumptions about humans and Nash equilibria.

There have been several previous works attempting to model human behavior in normal form games (Wright and Leyton-Brown 2010, 2014). Wright and Leyton-Brown (Wright and Leyton-Brown 2010) collected the results of multiple experiments from normal form games studied in the literature, and showed how the human action distribution can be modeled with high accuracy. However, our problem is clearly more complex and cannot be modeled as a simple normal form game.

Azaria et al. (Azaria et al. 2012, 2016) introduce SAP, a social agent for advice provision. They show that humans tend to ignore advice provided by a selfish agent. Therefore, they suggest using some linear combination of the user's and the agent's preferences. The exact ratio is determined by simulating human behavior and selecting the ratio that achieve the highest performance for the agent in simulation. Therefore, both SAP and our work attempt to maximize agent performance and consider a linear combination of both the user and the agent, however, the environment and settings are completely different, as SAP is an agent for advice provision, and we use a grid environment. In addition, the purpose of the linear combination used by SAP is to address the issue of human trust, while in our work, it is used to mitigate the uncertainty we have in our human model. Furthermore, we propose a formula for obtaining our proposed ratio, rather than running a simulation for obtaining that value.

The Single Track Road Game

We now provide a formal definition for the single track road game, which is the main focus of this paper. Two agents A and B are placed on a $2 \times n$ grid at both ends on the upper

row, where agent A is positioned at the upper right corner, with coordinates $(1, n)$, and agent B is positioned at the upper left corner, with coordinates $(1, 1)$, each agent's goal is to maximize $u(W)$, their future outcome where W refers to the agent. Each agent's goal is to reach the other side in a minimal number of steps, and without colliding with the other agent. The set of actions available for each agent depends on its location. In the upper row each agent can perform the following actions:

- *Advance*: move toward the other side.
- *Stay*: remain in current position.
- *Down*, move to the bottom row.

In the bottom row each agent can perform one of the following actions:

- *Stay*: remain in current position.
- *Up*: return to the top row.

Both agents take actions synchronously, and do not observe the other's action before they take their own action. We define the reward function as follows:

- *Collusion*: if both agents collide, each agent loses 100 points, and the game ends.
- *Arrived at destination*: an agent that arrives at its destination receives a reward of 30 points. The game ends only for the agent that has reached its destination, i.e., the second agent continues to play until it reaches its destination, in which case it will receive a reward of 30 points as well.
- *Time loss*: any agent that is still in the game (did not reach its destination or collided with the other agent) loses 1 point each time-step.

Game Theoretical Analysis

In this section we present the game-theory analysis for the single track road problem. Let $x(W)$ be the x coordinate (column) of agent W and let $y(W)$ be its y coordinate (row). Let $d(A, B) = x(A) - x(B)$. Note that if agent B has passed agent A , $d(A, B)$ will be negative.

Theorem. *For two agents A, B in the $2 \times n$ grid of the single track road game. The following strategies are in a sub-game perfect Nash equilibrium:*

- *Agent A uses the following strategy:*
 - If $y(A) = 1$ (it is in the upper row) it takes action *Advance*.
 - If $y(A) = 2$ (it is in the lower row) it takes action *Up*.
- *Agent B uses the following strategy:*
 - If $y(B) = 1$ (the agent is in the upper row):
 - * If $d(A, B) \geq 3$ or $d(A, B) < 0$, it takes action *Advance*.
 - * If $y(A) = 1$ and $d(A, B) = 1$ it takes action *Down*.
 - * If $y(A) = 1$ and $d(A, B) = 2$, it may either take action *Stay* or *Down* (or any mixed strategy of the two).
 - If $y(B) = 2$ (the agent is in the lower row):
 - * If $d(A, B) \leq 0$ it takes action *Up*.

- * If $y(A) = 1$ and $d(A, B) = 1$ it takes action *Stay*.
- * If $y(A) = 1$ and $d(A, B) \geq 4$ it takes action *Up*.
- * If $y(A) = 2$ and $d(A, B) \geq 3$ it takes action *Up*.
- * Otherwise, it may either take action *Stay* or *Up* (or any mixed strategy of the two).

Proof. The proof handles each of the agents separately and shows that no agent should deviate from its determined strategy under the assumption that the other agent remains with its strategy. This is true also for any sub-game. Given agent B 's strategy, agent A should not deviate, as deviation will either cause it longer to reach its destination (resulting in a lower reward), or to collide with agent B (if it decides to take action *Down* when agent B is directly below it), resulting in a much lower reward. Similarly, given agent A 's strategy, agent B should not deviate, due to the following:

- If $y(B) = 1$ (the agent is in the upper row):
 - If $y(A) = 1$ and $d(A, B) = 1$, under the assumption that A would *Advance*, taking an action other than *Down* would lead to a collision, which will result in a very low reward.
 - If $d(A, B) \geq 3$ or $d(A, B) < 0$, so either agent A is very far or it has already passed agent B . Therefore, there is no risk of collision, and deviating and taking action *Down* or *Stay* will result in arriving later at the destination, which will result at a lower reward.
 - If $y(A) = 1$ and $d(A, B) = 2$, deviating and taking action *Advance* would result in a collision. Therefore, agent B should take either action *Down* or *Stay* (or any mixed strategy of the two).
- If $y(B) = 2$ (the agent is in the lower row):
 - If $d(A, B) < 0$, there is no risk of a collision since agent A already passed agent B . Therefore, deviating and playing *Stay* delays B 's arrival at the destination.
 - If $y(A) = 1$ and $d(A, B) = 1$, playing action *Up* (instead of *Stay*) will lead to a collision, resulting in a lower reward.
 - If $y(A) = 1$ and $d(A, B) \geq 4$, since there is no risk of collision, taking action *Up* will yield the greatest reward, and any other action will cause it to reach the destination later.
 - If $y(A) = 2$ and $d(A, B) \geq 3$, similarly, any action other than *Up* will cause a delay in arriving at the destination.
 - Otherwise, agent B can choose whether to take action *Stay* or *Up* because there is no risk of a collision and it will not affect the arrival time. We note that if it takes action *Up* and agent A follows its strategy, agent B 's next action will be *Down*.

□

Clearly, due to the symmetry of the game, agents A and B may switch policies and the resulting set of strategies will be in equilibrium. However, since both sets of policies and equilibria are symmetrical, we cannot predetermine which equilibrium to select. Furthermore, as we will show in the

experiments, human agents, in most cases, do not follow any of the above strategies (see section).

Socially Aware Reinforcement Learning (SARL)

To solve the single track road problem, we introduce the Socially Aware Reinforcement Learning agent (SARL). For SARL, we model the problem as an MDP, in which the human is a part of the environment. A state is composed of the current location of both agents, as well as the location of both agents in the previous timestep, which serves as a model of the velocity. In order to model human behavior, we use a data-set of humans interacting in the single track road problem, and for each state we compute the fraction of humans that were in that state and took each of the possible actions. We apply the Laplace rule of succession (Zabell 1989). Given state s , let

$$A_s = \{a \in A : \text{apermitted from } s\}$$

where A is the set of all actions. For $a \in A_s$ let $|a_s|$ be the number of times in the dataset that action a was performed from state s and let $n_s = \sum_{a \in A_s} |a_s|$ be the total number of actions performed from s . We assume that the probability that the human will take action a at state s is

$$\mathbb{P}(a|s) = \frac{|a| + 1}{n_s + |A_s|}.$$

We use value iteration, an MDP planning algorithm, which is based on dynamic programming, for solving the MDP (Sutton and Barto 2018). The value iteration guarantees to find the optimal policy, under the assumption that the MDP model, which includes the human model, is accurate. Clearly, policy iteration would have yield the same policy, and model free reinforcement learning methods, such as q-learning (using a simulation), should also converge to the same policy, if allowed enough running time. However, since our human model is inaccurate, instead of using the common approach for solving the MDP by trying to maximize the agent's outcome directly, SARL uses a linear combination of its own outcome and the human's outcome. It is important to note that SARL is still selfish, it considers the human's outcome only because this is its way to maximize its own outcome. It is interesting to note that it has been shown in the field of psychology that people who consider other people's goals and show empathy, feel better with themselves and are more likely to reach their own goals (Carey, Tai, and Griffiths 2021). Furthermore, reciprocation and cooperation may result in the human returning a favor.

However, since our human model is based on a limited data-set size, we propose to incorporate our knowledge related to the human reward function into the optimization problem. Yet, instead of trying to use the human reward function as a part of the human model or the transition function of the MDP, we propose to add it to the objective function of the agent. To that end, we define the parameter β , a value between 0 and 1, that quantifies the degree to which the agent considers its own outcome and the human's outcome. Namely, the agent, A , instead of optimizing towards

$u(A)$, optimizes towards $\beta u(A) + (1-\beta)u(B)$. We note that when $\beta = 1$ the agent optimizes towards its own outcome. A β value of 0.5 denotes that the agent tries to optimize the social outcome (i.e., $0.5u(A) + 0.5u(B)$), which is identical to optimizing simply towards $u(A) + u(B)$, and when $\beta = 0$ the agent only considers the human’s utility function.

In general, for a two player game, in which one of the players is human, and given a data-set we propose a formula for computing the β value as follows. Let the vectors R_A and R_B , of length n , denote the final outcomes in the data-set for players A and B, respectively for each episode. The following formula provides the proposed β value to be used by SARL:

$$\beta = \frac{1 - \text{correl}(R_A, R_B)}{2}.$$

$\text{correl}(R_A, R_B)$ is the correlation between R_A and R_B , which is computed by:

$$\text{correl}(X, Y) = \frac{\sum_{x_i \in X, y_i \in Y} (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{x_i \in X} (x_i - \bar{x})^2 \sum_{y_i \in Y} (y_i - \bar{y})^2}}.$$

For example, in a zero-sum game, the correlation between the rewards of both players is -1 ; therefore, the agent will ignore the human’s outcome and only maximize its own. On the other hand, when the rewards of both players are independent, the correlation will be 0. Therefore, β will be 0.5, that is, the agent will consider both rewards equally. In a game that is more cooperative the human’s utility function is not as different from the agent’s, and therefore the value of β is lower.

Problem Specification

We use a 2×6 grid to model the single-road game problem, and the reward functions described in Section (see Figure 1). We set the discount factor, γ , to 0.999, so that the overall return is very close to the sum of the rewards.

We define a state as a pair (i, j) in which i is a position of the autonomous agent, and j is a position of the human agent. We refer to this state representation as a state *without* velocity. We also use a more complex representation of a state by considering also the previous locations of both players; this representation is referred to as a state with velocity. That is, a state is a tuple of two pairs $((i, j), (l, k))$, where the first coordinate of each pair corresponds to the position on the board of the autonomous agent, and the second coordinate corresponds to the position of a human agent. The first pair, (i, j) , is the current state of the two agents, and the second pair, (l, k) , is their previous state. Accounting for both players’ velocity allows the composition of a more accurate human model, which is a part of the transition function.

Since the game ends only when both players reach their destination (or collide with each other), in order to model the problem as an MDP with an objective function that also considers the human’s reward function we had to slightly modify the reward function. This is because a standard MDP only considers the agent’s point of view while the human is considered a part of the environment. Given β and a state

s , let $\text{remainingSteps}(s, X)$ be the number of remaining steps from state s for agent X , where X is either the autonomous agent or human player.

- If the two agents collide, the reward is

$$\beta \cdot (-100) + (1 - \beta) \cdot (-100).$$

- If the autonomous agent has not yet reached the destination but the human agent has, the reward is

$$\beta \cdot (30 - \text{remainingSteps}(s, \text{autonomous})) + (1 - \beta) \cdot 30.$$

- If the autonomous agent has reached the destination and the human agent has not, the reward is

$$\beta \cdot 30 + (1 - \beta) \cdot (30 - \text{remainingSteps}(s, \text{human})).$$

- For any other step, the reward is

$$\beta \cdot (-1) + (1 - \beta) \cdot (-1).$$

Experimental Design

In order to compose the data-set and evaluate SARL’s performance, we recruited 470 participants from Mechanical Turk (Paolacci, Chandler, and Ipeirotis 2010) to play the single road game. The participants first read the game instructions and were then required to answer three short and simple questions, to ensure that they had read and understood the instructions. The participants then played the game only once. Upon completion (either by reaching the other side, or if colliding with the other agent), the participants provided demographic information including whether they have a valid driving license, an expired driving license or no driving license. In addition, the participants were asked to state how much they agreed with each of the following five statements:

1. The agent played aggressively.
2. The agent played generously.
3. The agent played wisely.
4. The agent was predictable.
5. I felt the agent was a computer.

We used a seven point Likert-like scale (Joshi et al. 2015) for these statements, ranging from strongly disagree (1) to strongly agree (7).

446 participants completed the game and answered the survey. We used the following 4 different baseline agents for the data gathering phase.

1. *Careful*: an agent that adheres to the strategy of agent B in Theorem . That is, it tries to move left, but tries to avoid colliding with the other agent as well, so if moving left may risk colliding with the other agent it stays in place. If staying in place also risks colliding with the other agent, it moves down.
2. *Aggressive*: an agent that adheres to the strategy of agent A in Theorem . That is, the agent always moves left.
3. *Semi-aggressive*: an agent that moves left unless the other agent is already there, in which case it stays in place until the other agent moves out of its way.
4. *Random*: an agent that moves randomly.

Results

In this section we present a comparison of all agents mentioned above and show that SARL significantly outperforms all other agents. In addition, we consider the following agents:

1. *Non-Velocity VI*: runs a value iteration on the MDP without velocity using the appropriate human model.
2. *Velocity VI*: runs a value iteration on the MDP with velocity using the appropriate human model.
3. *Equal Social VI*: uses value iteration to maximize the sum of the agent’s and the human’s utilities (i.e., used a β value of 0.5).

The agent’s score is calculated by averaging all its scores in each game it plays. We begin by comparing the performance of each of the agents. Figure 2 presents a comparison between the performance of all baseline agents, Velocity and Non-Velocity Value Iteration ($\beta = 1$), Equal Social Value Iteration ($\beta = 0.5$) and SARL ($\beta = 0.13$). As depicted by

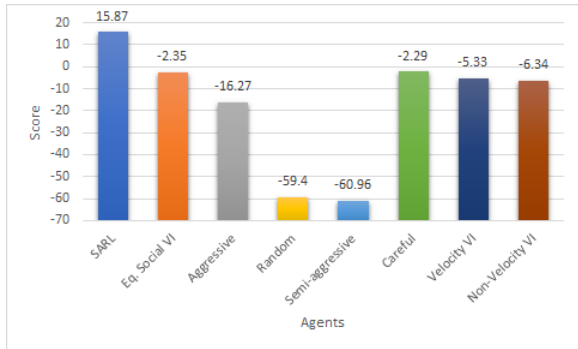


Figure 2: A comparison between the performance of all baseline agents, Velocity and Non-Velocity Value Iteration ($\beta = 1$), Equal Social Value Iteration ($\beta = 0.5$) and SARL ($\beta = 0.13$).

Figure 2, SARL significantly outperforms all other agents ($p < 0.01$) in terms of the agent’s performance. Furthermore, it is the only agent that achieved a positive average reward. We also note that the agent that uses the state representation with velocity obtained slightly better results than the agent that used the non-velocity state representation, though these differences are not statistically significant. We now turn to evaluate the human’s score when playing with each of the agents. Although the agents are designed to be selfish, clearly, it is more beneficial if also the human player would result with a better score. Table 1 presents the performance of each of the agents along with the performance of the humans playing against them.

As shown in Table 1, SARL also significantly outperforms all other agents ($p < 0.01$) in terms of the human’s performance.

In addition, we tested the performance of a velocity value iteration agent with the β value set to 0. That is, an agent that only considers the human reward. Interestingly, such an agent simply moves down and remains there forever. This

Table 1: A comparison between the performance of each of the agents along with the human player who played against each of them.

	Avg. agent’s score	Avg. human’s score	Avg. social welfare
Careful	-2.29	-0.86	-3.15
Aggressive	-16.27	-18.40	-34.67
Semi-aggressive	-60.97	-62.11	-123.08
Random	-59.40	-57.62	-117.02
Non-Velocity VI	-6.34	-9.03	-15.37
Velocity VI	-5.33	-6.03	-11.36
Eq. Social VI	-2.35	-4.09	-6.44
SARL	15.87	17.12	32.99

is because this way it does not disturb the human. Unfortunately, such an agent achieves a final outcome of $-\infty$ (or $-\frac{1}{1-\gamma}$) because it can never reach its destination, since the when the human’s reaches her goal, the agent is directly beneath her.

Next, we evaluate the prediction of the *policy evaluation* algorithm, using both forms of state representations (i.e., with and without velocity). Table 2 presents the prediction compared with the actual score of every agent. As can be seen in the table, the prediction that uses a state representation with velocity outperforms the prediction that uses a state representation without velocity. However, both predictions performed badly, and imply that our human model is not accurate, as an accurate human model would have resulted with an accurate prediction. This demonstrates that it is not enough to rely on the dataset, and strengthens the need for the socially aware approach, which also considers the human’s rewards.

Table 2: The accuracy of the prediction of a policy evaluation algorithm using a model with velocity and a model without velocity.

	True score	Prediction with velocity (error)	Prediction without velocity (error)
Careful	-2.29	-14.41 (12.12)	-4.86 (2.57)
Aggressive	-16.27	-6.21 (10.6)	1.14 (17.41)
Semi-aggressive	-60.97	-56.47 (4.5)	-47.81 (13.16)
Non-Velocity VI	-6.34	0.51 (6.85)	13.63 (19.97)
Velocity VI	-5.33	14.47 (20.02)	N/A
Social VI	-2.35	12.34 (14.69)	N/A
SARL	15.87	7.55 (8.32)	N/A

We now turn to analyze the survey results for each agent (see Table 3). Each value in the table is the average of all the scores of the measured values: Aggressively, Computer, Generously, Wisely and Predictable. Note that the lower the ‘Aggressively’ and ‘Computer’ parameters, the better the performance. On the other hand, the higher the ‘Generously’, ‘Wisely’ and ‘Predictable’ parameters, the better the performance. As can be seen in Table 3, SARL obtained the

Table 3: Survey results of all agents

	Aggress.	Comp.	Generous	Wise	Pred.
Careful	3.94	5.70	4.23	4.92	4.28
Aggressive	5.04	5.83	3.28	4.59	4.97
Semi-aggressive	4.57	5.73	3.21	4.33	4.52
Random	3.51	5.64	4.01	3.72	3.57
Non-Velocity VI	4.88	6.20	3.27	4.65	4.82
Velocity VI	4.82	6.01	4.20	4.72	4.76
Social value iteration	4.78	5.60	3.69	4.92	4.98
SARL	3.30	5.58	5.14	5.01	4.00

best results compared to the other agents among all parameters except its score on Predictable. These results entail that SARL demonstrates a clear improvement over all the other agents.

Next, we compare the performance of the humans according to their demographic information. We did not find any significant differences between male and female players, with female participants obtaining an average of -25.98 and male participants an average of -25.66 . Similarly, education level did not seem to have any impact on the performance of the participants. Most interestingly, participants with a driving license that has expired obtained a much lower average score (-60.15) than those with a valid driving license (-24.77) and those without a driving license. Although these differences appear to be statistically significant using a one-tail t-test ($p < 0.05$), this result requires deeper investigation, as the number of participants whose driving license has expired is only 13. Furthermore, an ANOVA test (Rogan and Keselman 1977) does not show that these differences are statistically significant.

Finally, we present the number of human participants who followed a strategy that could be in a Nash equilibrium. As can be seen in Figure 3, only a small portion of the participants followed one of the two strategies that could be in equilibrium: the ‘Careful’ strategy or the ‘Aggressive’ strategy. Clearly, most of the participants did not follow a strategy that could be in a Nash equilibrium.




Figure 3: The number and percentage of human participants who followed a strategy that could be in a Nash equilibrium as well as the number and percentage of them who did not follow any strategy in equilibrium. The error bars present the 95% confidence interval.

Conclusions and Future Work

In this paper we present the single track road problem. In this problem two agents face each-other at opposite positions of a road that can only have one agent pass at a time. We focused on the scenario in which one agent is human, while the other is an autonomous agent. We ran experiments with human subjects in a simple grid domain, which simulates the single track road problem. We showed that when data is limited, building an accurate human model is very challenging, and that a reinforcement learning agent, which was based on this data, did not perform well in practice. However, we showed that a social agent, i.e., an agent that tried to maximize a linear combination of the human’s utility and its own utility, achieved a high score, and significantly outperformed other baselines, including an agent that tried to maximize only its own utility. We provided a formula to compute what we believe to be a good choice for the β parameter, i.e., the ratio between the human’s and the agent’s utility when attempting to maximize the agent’s utility. In addition, we showed that the human achieved highest utility when interacting with SARL, a value that was significantly higher than when interacting with any other baseline agent. Furthermore, SARL was perceived by humans as less aggressive, more generous and wiser than all other baselines.

In future work we intend to show that SARL performs well also when considering other, possibly very different, settings. One option for such a setting is a setting with a continuous state space as well as a continuous action space. We further intend to utilise the idea of using a social agent approach, learned in the grid game environment, and to adapt SARL to a simulated autonomous vehicle environment with human drivers controlling simulated vehicles. Once we perform well in the simulated environment, we expect to run SARL in a real single-track-road scenario, with an autonomous vehicle and human drivers. We hope to show that SARL will perform well in the real-world environment, and that a social agent approach will be useful in practice. Another direction for future work is to focus on situations in which the human reward function is not available a priori. Such a situation would challenge the use of SARL, as it uses the human reward function for computing its objective function. One appealing option may be to use inverse reinforcement learning (Ng, Russell et al. 2000) to first learn the human’s reward function, and then, to use this function to compute the optimal policy for SARL.

Acknowledgment

This research was supported in part by the Ministry of Science, Technology & Space, Israel.

References

- Ariely, D.; Loewenstein, G.; and Prelec, D. 2003. “Coherent arbitrariness”: Stable demand curves without stable preferences. *The Quarterly Journal of Economics*, 118(1): 73–106.
- Azaria, A.; Gal, Y.; Kraus, S.; and Goldman, C. V. 2016. Strategic advice provision in repeated human-agent interactions. *Autonomous Agents and Multi-Agent Systems*, 30(1): 4–29.
- Azaria, A.; Rabinovich, Z.; Goldman, C. V.; and Kraus, S. 2015. Strategic information disclosure to people with multiple alternatives. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 64.
- Azaria, A.; Rabinovich, Z.; Kraus, S.; Goldman, C.; and Gal, Y. 2012. Strategic advice provision in repeated human-agent interactions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 26.
- Azaria, A.; Rabinovich, Z.; Kraus, S.; and Goldman, C. V. 2011. Giving advice to people in path selection problems.
- Azaria, A.; Richardson, A.; and Rosenfeld, A. 2016. Autonomous agents and human cultures in the trust–revenge game. *Autonomous Agents and Multi-Agent Systems*, 30(3): 486–505.
- Bitan, M.; Gal, Y.; Kraus, S.; Dokow, E.; and Azaria, A. 2013. Social rankings in human-computer committees. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 27.
- Camara, F.; Romano, R.; Markkula, G.; Madigan, R.; Merat, N.; and Fox, C. 2018. Empirical game theory of pedestrian interaction for autonomous vehicles.
- Camerer, C. F. 2003. *Behavioral Game Theory. Experiments in Strategic Interaction*, chapter 2, 43–118. Princeton University Press.
- Carey, T. A.; Tai, S. J.; and Griffiths, R. 2021. *Deconstructing Health Inequity: A Perceptual Control Theory Perspective*. Springer Nature.
- Chandra, R.; Guan, T.; Panuganti, S.; Mittal, T.; Bhat-tacharya, U.; Bera, A.; and Manocha, D. 2020. Forecasting Trajectory and Behavior of Road-Agents Using Spectral Clustering in Graph-LSTMs. *IEEE Robotics Autom. Lett.*, 5(3): 4882–4890.
- Deo, N.; and Trivedi, M. M. 2018. Convolutional Social Pooling for Vehicle Trajectory Prediction. In *2018 IEEE Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2018, Salt Lake City, UT, USA, June 18-22, 2018*, 1468–1476. IEEE Computer Society.
- Ding, W.; Chen, J.; and Shen, S. 2019. Predicting Vehicle Behaviors Over An Extended Horizon Using Behavior Interaction Network. *CoRR*, abs/1903.00848.
- Elhenawy, M.; Elbery, A.; Hassan, A.; and Rakha, H. 2015. An Intersection Game-Theory-Based Traffic Control Algorithm in a Connected Vehicle Environment. 343–347.
- Houenou, A.; Bonnifait, P.; Cherfaoui, V.; and Yao, W. 2013. Vehicle trajectory prediction based on motion model and maneuver recognition. In *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems, Tokyo, Japan, November 3-7, 2013*, 4363–4369. IEEE.
- Joshi, A.; Kale, S.; Chandel, S.; and Pal, D. K. 2015. Likert scale: Explored and explained. *British Journal of Applied Science & Technology*, 7(4): 396.
- Kim, B.; Kang, C. M.; Lee, S.; Chae, H.; Kim, J.; Chung, C. C.; and Choi, J. W. 2017. Probabilistic Vehicle Trajectory Prediction over Occupancy Grid Map via Recurrent Neural Network. *CoRR*, abs/1704.07049.
- Leon, F.; and Gavrilescu, M. 2021. A Review of Tracking and Trajectory Prediction Methods for Autonomous Driving. *Mathematics*, 9(6): 660.
- Nay, J. J.; and Vorobeychik, Y. 2016. Predicting human cooperation. *PloS one*, 11(5): e0155656.
- Ng, A. Y.; Russell, S. J.; et al. 2000. Algorithms for inverse reinforcement learning. 1: 2.
- Nguyen, T.; Yang, R.; Azaria, A.; Kraus, S.; and Tambe, M. 2013. Analyzing the effectiveness of adversary modeling in security games. 27(1).
- Paolacci, G.; Chandler, J.; and Ipeirotis, P. G. 2010. Running experiments on amazon mechanical turk. *Judgment and Decision making*, 5(5): 411–419.
- Rogan, J. C.; and Keselman, H. 1977. Is the ANOVA F-test robust to variance heterogeneity when sample sizes are equal?: An investigation via a coefficient of variation. *American Educational Research Journal*, 14(4): 493–498.
- Rosenfeld, A.; Azaria, A.; Kraus, S.; Goldman, C. V.; and Tsimhoni, O. 2015. Adaptive advice in automobile climate control systems.
- Shvartzon, A.; Azaria, A.; Kraus, S.; Goldman, C. V.; Meyer, J.; and Tsimhoni, O. 2016. Personalized alert agent for optimal user performance.
- Sutton, R. S.; and Barto, A. G. 2018. *Reinforcement learning: An introduction*. MIT press.
- Tversky, A.; and Kahneman, D. 1981. The Framing of Decisions and the Psychology of Choice. *Science*, 211(4481): 453–458.
- Wright, J. R.; and Leyton-Brown, K. 2010. Beyond equilibrium: Predicting human behavior in normal-form games.
- Wright, J. R.; and Leyton-Brown, K. 2014. Level-0 meta-models for predicting human behavior in games. 857–874.
- Zabell, S. L. 1989. The rule of succession. *Erkenntnis*, 31(2): 283–321.