

This is a repository copy of *Unfreezing autonomous vehicles with game theory, proxemics, and trust*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/196033/>

Version: Published Version

Article:

Camara, Fanta and Fox, Charles (2022) Unfreezing autonomous vehicles with game theory, proxemics, and trust. *Frontiers in Computer Science*. 969194. ISSN: 2624-9898

<https://doi.org/10.3389/fcomp.2022.969194>

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



OPEN ACCESS

EDITED BY

Amir Rasouli,
Huawei Technologies, Canada

REVIEWED BY

Lukas Heuer,
Robert Bosch, Germany
Iuliia Kotseruba,
York University, Canada
Soheil Alizadeh,
Huawei Technologies, Canada

*CORRESPONDENCE

Fanta Camara
fcamara@lincoln.ac.uk

SPECIALTY SECTION

This article was submitted to
Mobile and Ubiquitous Computing,
a section of the journal
Frontiers in Computer Science

RECEIVED 14 June 2022

ACCEPTED 14 September 2022

PUBLISHED 13 October 2022

CITATION

Camara F and Fox C (2022) Unfreezing
autonomous vehicles with game
theory, proxemics, and trust.
Front. Comput. Sci. 4:969194.
doi: 10.3389/fcomp.2022.969194

COPYRIGHT

© 2022 Camara and Fox. This is an
open-access article distributed under
the terms of the [Creative Commons
Attribution License \(CC BY\)](#). The use,
distribution or reproduction in other
forums is permitted, provided the
original author(s) and the copyright
owner(s) are credited and that the
original publication in this journal is
cited, in accordance with accepted
academic practice. No use, distribution
or reproduction is permitted which
does not comply with these terms.

Unfreezing autonomous vehicles with game theory, proxemics, and trust

Fanta Camara^{1,2*} and Charles Fox^{1,2}

¹Institute for Transport Studies (ITS), University of Leeds, Leeds, United Kingdom, ²School of Computer Science, University of Lincoln, Lincoln, United Kingdom

Recent years have witnessed the rapid deployment of robotic systems in public places such as roads, pavements, workplaces and care homes. Robot navigation in environments with static objects is largely solved, but navigating around humans in dynamic environments remains an active research question for autonomous vehicles (AVs). To navigate in human social spaces, self-driving cars and other robots must also show social intelligence. This involves predicting and planning around pedestrians, understanding their personal space, and establishing trust with them. Most current AVs, for legal and safety reasons, consider pedestrians to be obstacles, so these AVs always stop for or replan to drive around them. But this highly safe nature may lead pedestrians to take advantage over them and slow their progress, even to a complete halt. We provide a review of our recent research on predicting and controlling human–AV interactions, which combines game theory, proxemics and trust, and unifies these fields via quantitative, probabilistic models and robot controllers, to solve this “freezing robot” problem.

KEYWORDS

autonomous vehicles, pedestrians, interactions, freezing robot, robotics, game theory, proxemics, trust

1. Introduction

Autonomous vehicles (AVs)—known as intelligent/automated vehicles, autonomous driving systems (ADS), or self-driving cars—are appearing on the roads (Wade, 2018), thanks to huge improvements in localization, mapping, planning and navigation algorithms (Thrun et al., 2005; Cadena et al., 2016) together with large price falls making sensors and compute power widely available (Kato et al., 2015) and also thanks to large public and private investments. For instance, in 2015, the UK government alone invested £100 million in research and development for the deployment of Connected Autonomous Vehicles (CAV) technologies (House of Lords, 2017) and the global market is estimated to be worth £907 billion in 2035 (Catapult, 2017). Self-driving cars can plan routes and control steering, acceleration and braking to follow them, and are promoted as having potential to improve safety and efficiency and reduce pollution and travel time (Wadud et al., 2016; Kim, 2018; Millard-Ball, 2018). But at the same time some researchers question the benefits of AVs, and argue that AV makers may profit from “self-driving data generation,” the attention economy and increase the overall car dependency which may go against the claims for increased road safety (Norton, 2021).

Autonomous vehicles also include smaller delivery vehicles, typically intended to transport goods from shops or transport hubs over the “last mile” to customers’ homes, or for use inside and outside factories, warehouses, hospitals, care homes and other private and public organizations’ facilities (Murphy et al., 2020).

The Society of Automation Engineers (SAE) has created standard definitions for levels of automation in AVs (SAE International, 2019), as shown in Figure 1A, based on how much human driver input is relied upon. For cars, the human driver may be present inside the vehicle or operating it remotely, while for smaller delivery vehicles they must be remote. Many automotive companies currently claim having developed level 3 automated vehicles, the race is now toward the full automation (Rasouli and Tsotsos, 2018). But after decades of development and despite the global enthusiasm around AVs and the big investments, some major challenges still remain (Kirkpatrick, 2022). A full AV revolution would require AVs to share space with, and be challenged by, human pedestrians and drivers. Such humans are much harder to predict and plan for than purely passive environments. Navigating around humans in dynamic environments requires the understanding of human social behavior and remains an active research question (Thomaz et al., 2016).

Pedestrians are complex humans having goals, utilities, and decision making systems. Interactions with them must take these factors into account in order for AVs to plan and control their interactions with them. This is especially important in settings where there is no clear legal priority for space, including unmarked road intersections and crossings, and off-road environments for delivery vehicles such as pavements, corridors and pedestrianized areas (Portouli et al., 2014). Human drivers are trained and able to read the intention and objectives of other road users and to plan interactions with them accordingly (Rasouli and Tsotsos, 2018; Rasouli et al., 2018). We here define “road users” to refer to all agents who are in active control of the position of physical objects which may enter a road. This includes human drivers, autonomous vehicles, cyclists and pedestrians; and excludes passengers of all vehicles. Brooks identifies the need for these higher levels of interaction as “the big problem with self-driving cars” (Brooks, 2017).

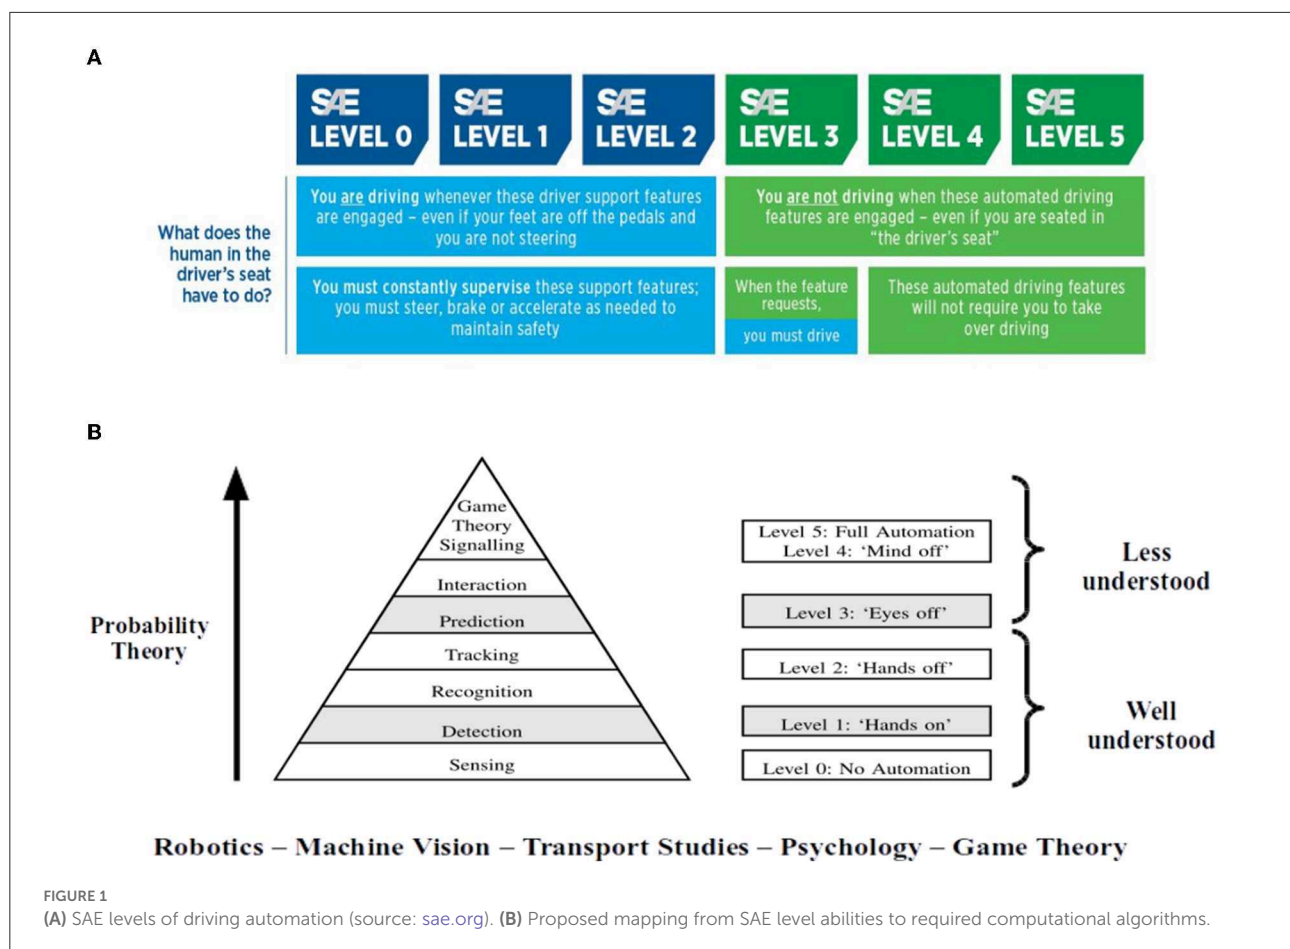
Currently AVs are designed to be as safe as possible, always yielding to other road users in competition with them. Typically this includes the use of a low-level safety system which stops the vehicle if a lidar sensor shows any obstruction within a short safety distance. However, recent studies investigating human interactions with AVs *via* a questionnaire (Deb et al., 2017), video analysis (Madigan et al., 2019) and modeling (Millard-Ball, 2018), have all shown that pedestrians may take advantage of this behavior once they understand it, potentially eventually pushing in front of them for *every* interaction. AVs will therefore make zero progress in busy areas if all pedestrians behave

optimally in this way. This scenario has been named as the “freezing robot problem” (Trautman and Krause, 2010). AVs thus need better prediction and decision-making models, and must find a good balance between stopping for pedestrians when required and encouraging pedestrians to get out of their way so that they can drive to reach their destination as quickly as possible for their passengers on board.

Like human drivers, AVs may therefore need to maintain a credible threat of actually hitting or otherwise causing some smaller negative utility to pedestrians in their planned path, so that they may at least in some cases encourage the pedestrians to get out of their way in order to make progress. Creating and implementing such threats requires understanding of the social behaviors of the human pedestrians and how they can be modified by the vehicle’s own actions. In road-crossing scenarios, a pedestrian and the AV compete against one another for the limited resource of the road space (Šucha, 2014; Parkin et al., 2016; Owens et al., 2018). Such competitions can be modeled with Game Theory (Osborne and Rubinstein, 1994). Game theoretic approaches have been used for decades to model interactions between rational decision-makers, but have not been previously applied to autonomous vehicles or to Psychology research on human proxemics and trust. Results from Game Theory and Psychology studies have yet to be operationalised for autonomous vehicles, our work thus aims to bridge the gap between these separate fields and we propose methods and solutions to bring them to an operational level for fully automated vehicles’ interactions with pedestrians.

This review hence provides a concise overview of our and others’ recent work in this area. The main topics are as follows:

- a summary of key findings from our larger study of the literature of pedestrian models required for autonomous driving, linking low-level models of machine vision detection and tracking with high-level models of human behavior, and an updated review of more recent game theoretic approaches;
- a game-theoretic Sequential Chicken model of pedestrian-AV negotiation, and methods for inference of its parameters from physical and virtual reality experiments;
- methods and findings on pedestrian-vehicle interaction sequences analysis;
- a novel Bayesian method to infer and explain quantitative pedestrian proxemic utility functions;
- the concept of physical trust requirement (PTR) for game theoretic AV interactions, and its recent generalization to human-human and human-robot interactions;
- OpenPodcar, an open source hardware (OSH) AV developed for real-world AV game theoretic research experiments with pedestrians;
- a discussion on the legal and ethical implications of this work as well as possible future directions.



2. Related work

We recently performed a two-part comprehensive review of existing pedestrian models for AVs, which organized them into a taxonomy from lower-level models of sensing, detection and tracking (Camara et al., 2020a) to higher-level models of pedestrian goals, behavior and interaction (Camara et al., 2020b). This work brought together research from many fields including Transport, Robotics, Machine Vision, Data Science, Game Theory and Psychology with levels of required algorithms mapped to SAE level abilities as summarized in Figure 1B. It further suggested how models from multiple levels could be linked to form a single robotics system using Bayesian probability theory, if the more Social Science based methods could be further quantified to operate using this probabilistic language.


At the lower levels of the technology stack, as shown at the top of Figure 2, pedestrian modeling was found to include mature methods from Machine Vision and Robotics for detection and tracking of pedestrians, and for prediction of their short term movements.

Neural network (“deep learning”) based pedestrian detection and recognition has largely replaced classical feature-based

methods, due to price falls in GPUs and new GPU-based heuristics enabling standard neural network algorithms to run at large scales. Open-source implementations of standard neural networks methods includes YOLO (Redmon et al., 2016; Redmon and Farhadi, 2017, 2018), R-CNN (Girshick et al., 2014), Faster R-CNN (Ren et al., 2017) and Mask R-CNN (He et al., 2017). Most methods operate on visual video frames, though lidar or millimeter radar are increasingly used with 3D versions of CNN neural networks (Qi et al., 2017).

Tracking remains challenging however for multiple pedestrians when some of them may be occluded by one another or by environmental obstacles. Open-source implementations of tracking are available including the Bayes Tracking library (Bellotto et al., 2015), SORT tracker (Bewley et al., 2016) and DetTa (Breuers et al., 2018) pipeline. Recurrent and graph neural networks have recently been applied to tracking, again benefiting from the recent availability of GPU computing, outperforming probabilistic methods in many cases (Ravindran et al., 2020; Weng et al., 2020; Wang et al., 2021).

Recognizing pedestrian skeleton pose and head direction is now possible in many cases, again mostly using neural network methods. Open-source systems include Openpose (Wei et al., 2016; Cao et al., 2017, 2019; Simon et al., 2017) for skeleton



estimates, OpenFace (Amos et al., 2016) for head pose and gaze direction recognition, and OpenTrack¹ for head pose tracking.

Recognizing sequences of pedestrian actions as high-level events remains a research area. This includes explicit signaling gestures such as waving, implicit signaling such as body language, and motions suggesting the emotional state or utility function of pedestrians.

At the higher levels shown at the bottom of Figure 2, AVs do not generally attempt to replicate the social intelligence of human drivers as they attempt to model, predict, interact with and communicate with the other human road users that they encounter. This can include understanding factors such as demographics, emotion, and environmental settings which affect probable goals and behaviors.

Our review located existing models of these processes in the Psychology, Transport Studies, and Social Science literature, but found they were not yet quantified and operationalized to the level of detail needed to integrate them into AV controllers. Where multiple pedestrians and occlusions are present, these higher-level understandings may be needed to form priors to fill in the missing parts of the scene, in addition to making more accurate predictions of future behavior.

Higher-level models of pedestrian trajectories go beyond simple short-term linear constant models (Fajen and Warren, 2003; Puydupin-Jamin et al., 2012), and consider the origin and likely destination of the pedestrian (Arechavaleta et al., 2008; Papadopoulos et al., 2013). They may use robotics kinematics models to suggest likely optimal curved paths between the current and goal position. Goal location may be simply based on population statistics, for example most pedestrians stepping into a road are likely to be aiming to cross to the point

immediately opposite (Kruse et al., 1997; Tamura et al., 2012; Dias et al., 2019). Or more detailed models may condition on pedestrian class membership or even individual identity based on historical behaviors. Class memberships may include visible static features of pedestrians—such as those in suits more likely to be walking to an office building entrance—as well as visible dynamic features such as running pedestrians more likely to be heading for a station to catch a train (Holland and Hill, 2007; Kooij et al., 2014; Rasouli and Tsotsos, 2020). Where multiple destinations remain probable, trajectory prediction may need to consider all of them until further information is available (Kitani et al., 2012; Karasev et al., 2016; Koschi et al., 2018; Rehder et al., 2018; Wu et al., 2018; Deo and Trivedi, 2020).

Game Theory is widely used to model decision-making between rational agents, in economics (Morgenstern and Von Neumann, 1953) and in transport network flow simulation (Figliozzi et al., 2008; Kim and Langari, 2014; Na and Cole, 2014; Talebpour et al., 2015; Flad et al., 2017; Tian et al., 2019). It has also been used in multi agent robotics for coordination tasks (Mavrogiannis and Knepper, 2019; Mavrogiannis et al., 2022). There are fundamental differences between these styles, with economics/transport typically taking an offline, data-driven, explanatory approach while robots require an online, single-shot, real-time decision making approach. Fusion of these two research streams appears to be a promising research avenue to understand and control AV-pedestrian interactions, and this idea forms the basis for our own experiments.

2.1. Recent game theory approaches

Recently more game theoretic models and methods have been developed for autonomous vehicle decision-making

¹ <https://github.com/opentrack/opentrack>

with pedestrians, beyond those previously reviewed. Several game theoretic approaches have been used for vehicle–vehicle interactions at intersections such as Schwarting et al. (2019), Tian et al. (2020), and Cleac’h et al. (2021) but we here focus on AV–pedestrian interactions. For example, a game theory model for pedestrian motion and walking behaviors has been developed in Rahmati et al. (2020). This interactive framework for simultaneous decision-making in vehicle–pedestrian or vehicle–vehicle interactions is calibrated with ground truth pedestrian trajectories and its performance is evaluated on predicting the decisions of human agents interacting with other road users. This model was recently extended into a framework for pedestrian–vehicle and pedestrian–pedestrian interactions (Rahmati and Talebpour, 2018). A level- k game theory model for autonomous vehicle controllers has also been recently proposed for unsignalised intersections, based on a discrete time, set of actions and a reward function (Li et al., 2018).

A game theoretic method based on virtual bargaining for road negotiations for priority was proposed in Miszyak et al. (2014) and Chater et al. (2018). The authors argue that autonomous vehicles will need cognitive science knowledge before becoming a reality. Škugor et al. (2020) used the game theory-based model from Chen et al. (2016) for pedestrian–vehicle interactions and extended it with a stochastic component to account for different behaviors and environmental influence. The model only considers single vehicle and pedestrian interactions and is validated *via* large-scale simulations. Michieli and Badia (2018) proposed a game theoretic model to evaluate traffic participants’ safety with the introduction of autonomous vehicles using scenarios involving cyclist–vehicle and pedestrian–vehicle interactions at unmarked intersections. Road users are modeled with sequential actions and environmental factors such as the vehicle speed is taken into account for payoffs and the model was tested in simulation.

Jafary et al. (2018) surveyed autonomous vehicles’ interactions with other traffic participants, including pedestrians and concluded that more research is needed for advanced methodologies and algorithms for a more robust AV decision-making. Zhang et al. (2022) developed a Bayesian game-based decision model for road user interactions including a method to evaluate the opponent’s aggressiveness, and a Turing test is performed to evaluate the human-likeness of the approach.

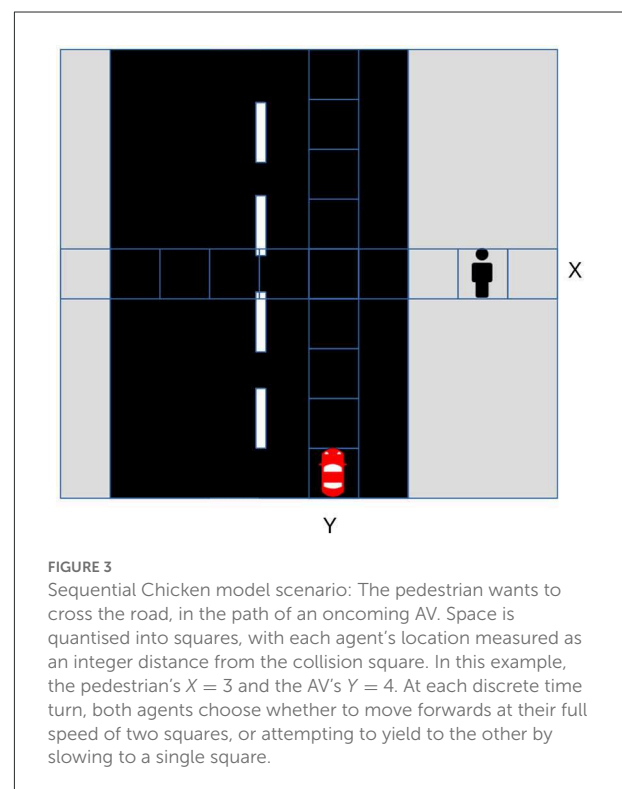
These game theory approaches provide analytic solutions for higher order models of behavior but most of them are currently only evaluated in simulations. Also, game theory models can be computationally expensive and would require tractable solutions for their real-time applications in autonomous driving.

As a first step toward real-time testing, the game theory model used in our own work is called the *Sequential Chicken* and is based upon the well-known “game of chicken”. The model and its extensions are detailed in the following sections.

3. Sequential Chicken game theory model

Game Theory’s core concept is (Nash) equilibrium, defined as a set of optimal actions for the agents when they do not possess information about the other’s choices. Basic game theory assumes two players each select a single action from a finite, and discrete set, at the same time, then each receive a utility as a function of the pair of selected actions. Both players know this utility function in advance of the game. In sequential game theory, several basic games are played in a series, with the choice of game at each step determined by the result of the previous one, and the utilities for the sequence awarded to the players at the end of the final game.

We have created a game theoretic model (Fox et al., 2018) of a pedestrian and an AV negotiating for shared space as a pedestrian considers whether to cross in front of the AV’s path as in Figure 3. This is modeled as a discrete sequential game theory model. Two utility parameters (U_{time} , U_{crash}) are used, which refer to the value of time and the utility of avoiding a collision. The pedestrian X and autonomous vehicle Y approach each other. Space is quantized into squares, and time into discrete turns. At each turn, both agents may select discrete speeds, simultaneously, to be either “slow”—1 square per turn—or



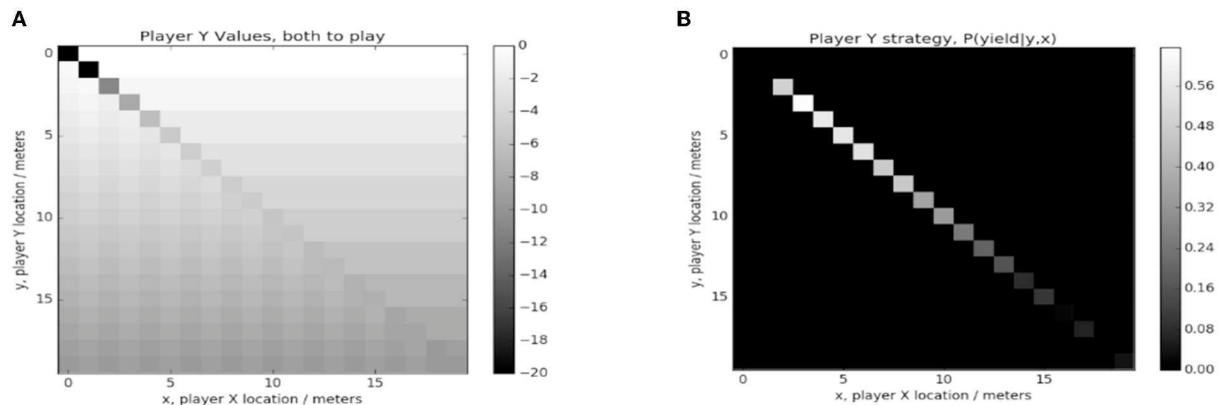


FIGURE 4

(A) Value function for player Y over agent joint distance (x, y) states up to 20 meters of each agent from the collision point. All values are negative as they include components from cost of time needed to reach the destination and the risk of collision occurring. The collision state $x = y = 0$ is shown in the top left and has the worst value. Player Y arriving at their destination occurs along the rest of $y = 0$ and has the best value. (B) Player Y's strategy over agent joint distance (x, y) states up to 20 meters of each agent from the collision point. For each state, Y's best strategy is to yield to Player X (i.e., slow down) with the probability shown. The interesting cases are along the diagonal where a negotiation must occur, and it can be seen that the yield probability ramps up as distance to collision decreases.

“fast”—2 squares per turn. The discretisations are used to make the model easier to solve, however by choosing suitably small squares and short time steps, the model can approximate arbitrarily close to the continuum of possible positions and the continuum of speeds ranging between the quantised slow and fast speeds.

Both agents would like to pass the collision point as quickly as possible as each turn of time has value, U_{time} , which is lost if they slow down. However if neither yield to the other then a collision occurs and they both suffer a much larger negative utility, U_{crash} . At each turn, the available actions for the players are the slow and fast speeds, $a_Y, a_X \in \{1, 2\}$. The model is designed to be simple and solvable, so excludes, for example, the possibilities of steering laterally to avoid the other player, or communicating with them in any way other than by the observable choices of speeds.

The optimal strategies are derivable from sequential game theory together with a novel meta-strategy convergence solution concept, *via* recursion. Sequential Chicken can be viewed as a sequence of one-shot sub-games, whose payoffs are the expected values of new games resulting from the actions, and are solvable by standard game theory.

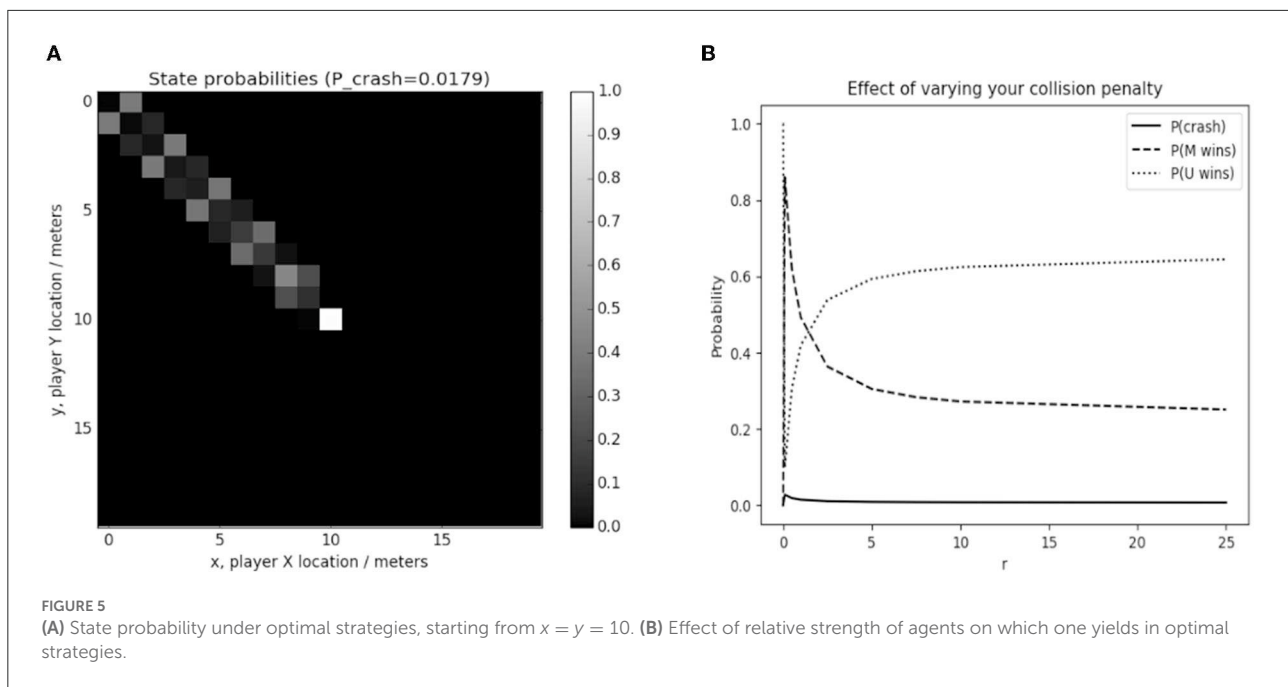
More formally, the discrete locations of the players can be represented by (y, x, t) at discrete turns t and their actions represented by $a_Y, a_X \in \{1, 2\}$ for speed selection. The new state at turn $t + 1$ is given by $(y + a_Y, x + a_X, t + 1)$. Define $v_{y,x,t} = (v_{y,x,t}^Y, v_{y,x,t}^X)$ as the value (expected utility, assuming all players play optimally) of the game for state (y, x, t) . As in standard game theory, the value of each 2×2 payoff matrix can

then be written as,

$$v_{y,x,t} = v \left(\begin{bmatrix} v(y-1, x-1, t+1) & v(y-1, x-2, t+1) \\ v(y-2, x-1, t+1) & v(y-2, x-2, t+1) \end{bmatrix} \right). \quad (1)$$

This is a recursive equation, in which the value $v_{y,x,t}$ at time t is defined in terms of values at later times, such as $v(y-1, x-1, t+1)$. The recursion is guaranteed to terminate as long as both agents are always moving forwards at either of their slow or fast speeds, because this means that they must eventually either reach their destinations or crash, with either result defining the value at that time directly as a numerical utility. The equation can be solved using dynamic programming assuming meta-strategy convergence equilibrium selection. Under some approximations based on the temporal gauge invariance described in Fox et al. (2018), we removed the dependencies on the time t so that only the locations (y, x) are required in computation of $v_{y,x}$ and optimal strategy selection.

Sample results from simulations of optimal agent behavior from this model are shown in the figures: Figure 4A shows the value function for one agent and its optimal strategy in Figure 4B. Figure 5A shows the resulting probability of a collision occurring, when the two agents are identical and start in a symmetric position. The figures show these values as shades, with the (x, y) position of a square in the figure corresponding to a state of the world in which the two agents are x and y steps away from the collision point, respectively. So the state of a collision is the top left square, and the diagonal contains all the symmetric states, i.e., where the two agents are equal distances from the collision point.



These and similar results provide a solution to the freezing robot problem. If the vehicle is programmed to be perfectly safe and always yield to pedestrians, it will freeze and never make any progress in a series of pedestrian interactions. However, if the vehicle is programmed to play the game optimally, then it will sometimes choose to not yield, with a small probability. In an even smaller probability of cases, this leads to an actual collision. The small probability of a collision occurring provides a *credible threat* to the pedestrian, encouraging them to consider yielding, and sufficient for the vast majority of interactions to proceed without collision, but with some of the pedestrians yielding to allow the vehicle to make progress.

In the above results, the utility function is symmetric, assigning equal time and crash costs to both agents. For pedestrians interacting with vehicles, this is typically not the case, as the cost to a pedestrian of being hit is higher than the cost to most vehicles. When the utility functions are made asymmetric to include this effect, the optimal strategies shift so that the weaker agent yields with higher probabilities. Figure 5B shows the effect of varying the ratio, r , of the crash utilities of the two agents, on the probability of who must yield to whom. This suggests for example that purchasing and driving expensive, heavy vehicles such as SUVs can be rational for owners whose value of time is high, as they reduce the owner's perceived risk on the road (Thomas and Walton, 2007), which can reduce their driving time by encouraging other road users to yield in interactions. Unless of course other road users also switch to larger vehicles, creating a higher-order, arms-race style game. This effect appears to occur in some populations, and

game theory suggests applying taxation or other penalties as a Mechanism Design to reduce it.

4. Learning parameters from human experiments

The Sequential Chicken model is not a complete predictive theory of interaction because it contains free parameters describing pedestrian preferences. To make it into a predictive model, experiments are needed to infer values of these parameters. To set the parameters for the Sequential Chicken model, Fox et al. (2018), we performed several experiments with human participants, and inferred the utilities of time and collisions from their behavioral data.

In a first empirical study (Camara et al., 2018d), we measured participants' behavior whilst playing the Sequential Chicken model as a board game. We inferred the model parameters using Gaussian Process regression (Rasmussen and Williams, 2005) and the resulting posterior showed a preference for saving time, U_{time} , rather than avoiding a collision, U_{crash} , in the average subject. This study provided a first understanding of how to perform this type of inference.

In a second type of study (Camara et al., 2018a, 2020c), we developed a novel experimental protocol which tracked human subjects using lidar. Pairs of human subjects were instructed to walk toward and pass one another, negotiating for space as in the AV-pedestrian scenario, as shown in Figure 6. We quantised their positions in space and time (first by having them walk in discrete steps and turns; second by allowing them to

walk naturally and performing the quantisation purely as data processing) in order to consider them as observations from the Sequential Chicken model. The Gaussian process method developed from the board game study was then applied to again infer utility parameters from the new, more realistic data.

These studies showed that participants were mostly playing rationally in accordance with the model under the best-fitting parameters, with only 11% of their actions deviating from optimal behavior. They also showed participants' preference for time saving over prevention of collision. This latter finding was



FIGURE 6
Lidar tracked human subject pair playing continuous Sequential Chicken.

unexpected, but was explained by conditions of the experiment. Being watched in a high safety lab, participants were behaving as if they were in a competition and so their preferences for saving time was rather unrealistic and discarding the high negative utility associated with a real-world collision, which would be more serious outside the lab, whether with another pedestrian, or worse with a vehicle.

As a consequence, we next moved to virtual reality (VR) to further investigate human interaction preferences in two virtual environments, as shown in Figure 7. This was to enable to study of actual realistic collisions between human subjects and vehicles to proceed in a safe way. We developed a virtual game theoretic autonomous vehicle that interacted with human participants (Camara et al., 2019a, 2020d, 2021). These results showed a much more realistic crossing behavior from participants, preferring avoiding collisions with the virtual AV rather than saving time. When presented with different AV behaviors *via* a gradient descent approach, participants preferred an AV that makes its decisions quickly. Finally, we found similar crossing behaviors in both virtual environments, as previously shown in Nuñez Velasco et al. (2019).

Gaussian Process results are shown in the figures. Figure 8A is from the lidar tracked human pairs, and Figure 8B from the VR simulation. These plots show contours of posterior belief distribution over the joint 2D space of possible utilities of time and collision. These contours start out flat, uniforming covering the space, and transform into the curved posteriors as evidence





FIGURE 7
Virtual reality experiment (Camara et al., 2021).

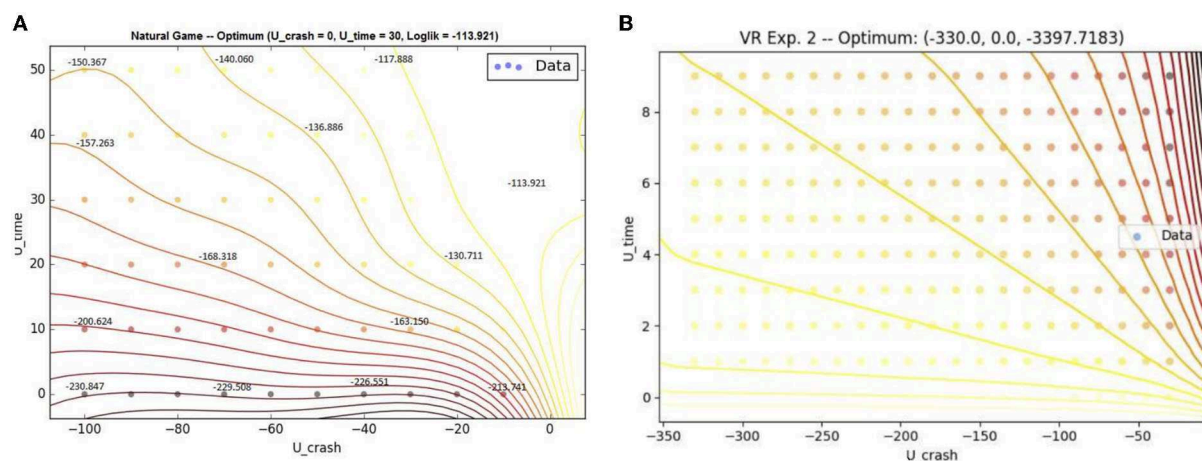


FIGURE 8

(A) Gaussian Process posterior belief over values of time and collision inferred from the physical human pair game experiment (Camara et al., 2018d). The contour lines show heights of the belief function over the 2D space of possible joint values for (U_{crash}, U_{time}) . The posterior appears approximately circumferentially symmetric, with its maxima along a ray from the origin to around $(-10, 50)$. (B) Gaussian Process posterior belief over values of time and collision inferred from the VR experiment (Camara et al., 2021). Circumferential symmetry is again approximately present, with the maxima now found around the horizontal ray from the origin to $(-350, 0)$. This suggests that subjects become more concerned about collisions in VR than in the game.



FIGURE 9
Intersection for the interactions' observation near the University of Leeds (Camara et al., 2018b,c).

from the experiments are Bayes-fused together. In Figure 8A the maxima occur along the ray which equates the negative utility of a crash with around 3 s of time, while in Figure 8B the maxima are close to the ray where the crash utility is worth near infinite time. This shows that the VR simulation is much more like real life than the lab game.

5. Learning interaction sequence patterns from data

The Sequential Chicken model so far assumes that pedestrians are all the same, by inferring single values of parameters to model all pedestrians together. However, pedestrians are not all the same, they have different motivations and risk preferences, which might both affect their game theoretic behavior and appear as externally visible features. If such features can be identified, they could be used to refine the parameter values for the Sequential Chicken model when used to predict and interact with particular individuals.

To learn social behavioral patterns from current pedestrian-vehicle interactions, we recorded a large-scale dataset of real-world human road crossings at the road intersection shown in Figure 9. The obtained pedestrian-vehicle interactions were human annotated as sequences of discrete events, such as looking, stepping and signaling, as listed in Table 1.

Logistic regression, decision tree regression, and motif analysis were then used to learn the most common sub-sequences of actions deployed by human drivers and pedestrians during the interactions. Table 1 shows these 62 temporal and 12 environmental descriptor features, which are predictive about the content and result of the interaction and which could be useful to condition the Sequential Chicken model by informing beliefs about the pedestrian's utility functions (Camara et al., 2018c). Figure 10 then shows the main findings from the

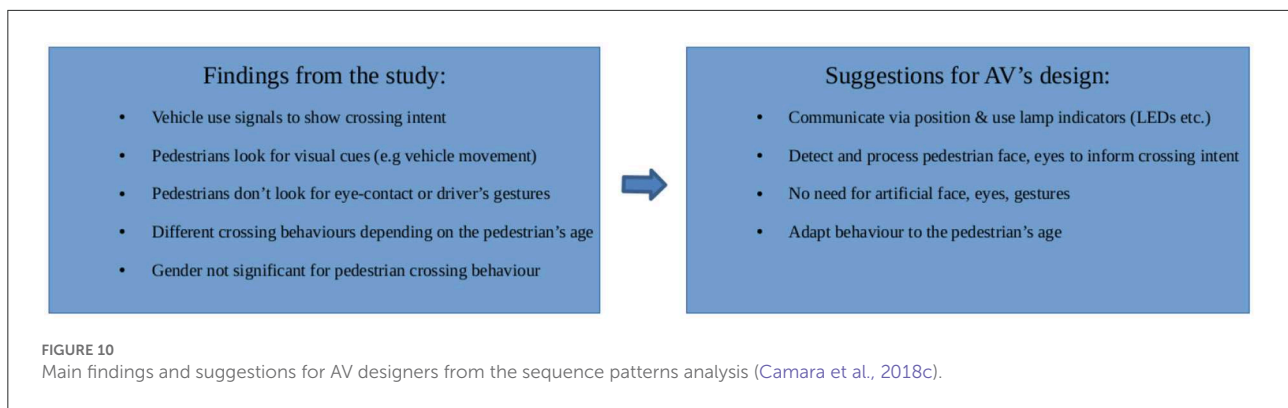
TABLE 1 List of the 74 features selected for the sequence patterns analysis, of which, 62 are temporal events and 12 are environmental descriptor events (Camara et al., 2018b,c).

Features	Features
0: Approaching Phase: Driver/Vehicle Stopped due to traffic	1: Approaching Phase: Vehicle Used signals Turn Indicator
2: Approaching Phase: Driver/Vehicle approached From left	3: Approaching Phase: Pedestrian Movements Slowed down
4: Approaching Phase: Driver/Vehicle Accelerated	5: Approaching Phase: Driver/Vehicle Turned right
6: Approaching Phase: Pedestrian Head Movements Turned left	7: Approaching Phase: Pedestrian Looked at approaching vehicle
8: Approaching Phase: Pedestrian Stepped on road and stopped	9: Approaching Phase: Driver/Vehicle Passed the pedestrian
10: Crossing Phase: Pedestrian Initiated crossing movement	11: Approaching Phase: Driver/Vehicle Decelerated due to traffic
12: Approaching Phase: Driver/Vehicle Kept pace	13: Approaching Phase: Pedestrian Stopped at the edge of the pavement
14: Crossing Phase: Pedestrian Head Movements Turned right	15: Crossing Phase: Pedestrian Looking at Looked at vehicle
16: Crossing Phase: Pedestrian Movements Other (elaborate in notes)	17: Approaching Phase: Pedestrian Movements Kept pace
18: Approaching Phase: Pedestrian Head Movements Turned right	19: Approaching Phase: Driver/Vehicle Interacting vehicle Van
20: Approaching Phase: Driver/Vehicle Turned left	21: Crossing Phase: Pedestrian Stepped back on pavement
22: Crossing Phase: Pedestrian Head Movements Turned left	23: Approaching Phase: Driver/Vehicle Head Turned in the direction of pedestrian
24: Approaching Phase: Pedestrian Movements Did not Stop	25: Crossing Phase: Pedestrian Slowed down/stopped
26: Approaching Phase: Driver/Vehicle Decelerated for observed pedestrian	27: Crossing Phase: Pedestrian Speeded up
28: Approaching Phase: Driver/Vehicle Interacting vehicle Other (elaborate in Notes)	29: Crossing Phase: Pedestrian Looking at other RUs (elaborate in comments)
30: Approaching Phase: Driver/Vehicle Interacting vehicle Bus/Truck	31: Approaching Phase: Driver/Vehicle Stopped due to other pedestrian
32: Approaching Phase: Driver/Vehicle approached from Multiple	33: Approaching Phase: Driver/Vehicle Used signals Flashed Lights
34: Crossing Phase: Pedestrian Raised hand sideways	35: Approaching Phase: Driver/Vehicle Decelerated due to other pedestrians
36: Approaching Phase: Pedestrian Looking at other RUs Others (elaborate in notes)	37: Approaching Phase: Pedestrian Looking at other pedestrians entering the road
38: Crossing Phase: Driver/Vehicle Passed the pedestrian'	39: Crossing Phase: Pedestrian Looked at driver
40: Crossing Phase: Driver/Vehicle Stopped for observed pedestrian	41: Crossing Phase: Driver/Vehicle Head Turned in the direction of pedestrian
42: Crossing Phase: Driver/Vehicle Raised hand in front	43: Crossing Phase: Pedestrian Raised hand in front
44: Crossing Phase: Driver/Vehicle Turned right	45: Approaching Phase: Vehicle Stopped for observed pedestrian
46: Crossing Phase: Driver/Vehicle Accelerated	47: Approaching Phase: Pedestrian Speeded up
48: Crossing Phase: Driver/Vehicle Decelerated for observed pedestrian	49: Crossing Phase: Vehicle Waved hand
50: Crossing Phase: Driver/Vehicle Movement Other (elaborate in notes)	51: Crossing Phase: Driver/Vehicle Turned left
52: Approaching Phase: Pedestrian Hand Movements Other (elaborate in notes)	53: Approaching Phase: Driver Head Turned right
54: Approaching Phase: Driver/Vehicle Movement Other (elaborate in notes)	55: Approaching Phase: Driver/Vehicle Head Movements Other (elaborate in notes)
56: Approaching Phase: Driver/Vehicle Head Turned left	57: Crossing Phase: Pedestrian Waved Hand
58: Crossing Phase: Pedestrian Hand Movements Other (elaborate in notes)	59: Approaching Phase: Driver/Vehicle Waved hand
60: Crossing Phase: Driver/Vehicle Used signals Turn Indicator	61: Crossing Phase: Pedestrian Looking at other pedestrians entering the road
Driver/Vehicle Interacting Vehicle is Single	Driver/Vehicle Interacting Vehicle Coming From right
Weather: Overcast	Weather: Sunny
Weather: Rainy	Group of Pedestrians
Pedestrian: teenager (13-18y)	Pedestrian: young adult (18-30y)
Pedestrian: mid-age adult (30-60y)	Pedestrian: older person (60+ years)
Pedestrian's Distraction	Pedestrian: Gender

sequence analysis and some suggestions for the design and development of AVs.

We then used the dataset to study the filtrations (temporal orderings) in which the annotated features can be revealed to an autonomous vehicle and their amount of information

provided about the interaction result at each time step during the interaction (Camara et al., 2018b). This analysis showed that an AV should continue driving as normal, while it observes around 7–10 features out of the 74 shown in Table 1 from a potentially interacting pedestrian, in order to best inform



its beliefs about their behavior before beginning to make its own game-theoretic actions to interact with them. Using the public Daimler pedestrian dataset (Kooij et al., 2014) we also developed simple heuristic features which can be fused to predict road crossing intent to some extent (Camara et al., 2019b). These predictions could be integrated into Sequential Chicken based AV controllers as priors to improve their predictions and interactions. It should be noted that the detection and recognition of most of the features shown in Table 1, e.g., pedestrian hand gestures, head movements, age, level of distraction, still pose several challenges due to their variabilities and remain active research areas (Camara et al., 2020a,b).

6. Unifying and quantifying proxemics and trust

The Sequential Chicken model showed that if the vehicle's only way to inflict negative utility onto pedestrians is to actually hit them, then it must be programmed to deliberately provoke a crash occasionally in order to make progress. This is clearly an undesirable and unethical solution to the real-world freezing robot problem, and at first appears to be a limitation of the model. However, the equations of the model still work if other, less violent, forms of negative utility are made available for the vehicle to inflict upon members of the public, with higher frequency traded for lower damage.

Possible solutions that are currently under debate might include spraying jets of water at anti-social pedestrians intentionally blocking the AV's path, or humiliating them in public using horns as is often done by human drivers to penalize other road users' anti-social behavior (UK Law Commission, 2019).

An intriguing additional option which we have chosen to explore is to make over-assertive pedestrians feel uncomfortable by invading their proxemic space. The theory of proxemics was introduced by Hall (1966) to describe humans' psychological

sense of comfort or discomfort during physical interactions. Hall proposed four distinct zones: the intimate zone ranges up to 0.45 m, the personal zone from 0.45 to 1.2 m, the social zone from 1.2 to 3.6 m, and the public zone from 3.6 m to infinity (Lambert, 2004). Social robotics experiments have shown that these proxemic zones change in size when humans interact with robots of different heights, appearances, speeds, voices, and also for different HRI activities (Rios-Martinez et al., 2015).

But in order to include proxemics in the Sequential Chicken model, the proxemic utility needs to be given in the form of continuous functions corresponding to the continuous motion of people. A review of the literature showed that there is no available method to infer the continuous proxemic functions. We thus proposed and developed a novel Bayesian method that can infer pedestrian utility functions from observed pedestrian-vehicle interactions' trajectory data.

This method fitted a variety of parametric models, M_i , to the data, optimizing to find the best-fitting parameters for each. The data are the distances between the two agents, X , their speeds, v and v_{ped} , and the results of the interactions, which are binary: either the pedestrian crosses or they yield. Bayesian Information Criteria (BIC) is then used to identify the best fitting model, automatically accounting for the Occam factors due to some models having more parameters than others.

Following that, we recently developed the first mathematical model of proxemics and trust concepts for self-driving cars and pedestrians interactions (Camara and Fox, 2020). It defined the trust zone as the area of the proxemics zones where trust is required i.e., one agent has to rely on the other during the interaction. Based on Lee and See (2004)'s qualitative definition of trust as an attitude in "a situation characterized by uncertainty and vulnerability", we more quantitatively defined a Physical Trust Requirement (PTR) as a Boolean property of the physical state of the world (not of the psychology of the agents) with respect to Agent1 during an interaction, true if and only if Agent1's future utility is affected by an immediate decision made by Agent2. For example, in the case of a car approaching a pedestrian, PTR exists in those states in which the car is driving toward the pedestrian at




FIGURE 11
Autonomous vehicle entering pedestrian's social zone, which can also be viewed and quantified as a trust region (Camara and Fox, 2020, 2022).

speeds such that the car driver is able to brake and avoid a collision, but the pedestrian is unable to get out of the car's path using their own slower walking or running speeds. The driver but not the pedestrian can determine the outcome of the interaction. Trust itself is then a psychological property of the pedestrian, with regard to the car, whose presence means that the pedestrian is willing to enter a PTR state with the car, as the pedestrian assumes that the car will then act in the pedestrian's own interest.

Applying our PTR concept quantitatively to the case of a car and pedestrian approaching each other at a right angle, as in the Sequential Chicken scenario, using physical assumptions for realistic braking distances and pedestrian motion, results in three discrete zones appearing around the pedestrian are shown in Figure 11.

Crash zone is the region close to Agent1, $\{d: 0 < d < d_{crash}\}$,

$$d_{crash} = v_2 t_2 + \frac{v_2^2}{2\mu_2 g}, \quad (2)$$

in which a crash is guaranteed and neither party can prevent it (Lyubenov, 2011), v_2 is Agent2's speed. The first term depends on Agent2's thinking reaction time, t_2 , and the second term represents the physical braking distance if Agent2 is a wheeled agent, μ_2 is the coefficient of friction between Agent2's tires and tarmac, and g is gravity. If Agent2 is a walking agent, we will here assume this second term is omitted as walkers are always in static equilibrium and can stop instantly once a decision is made. Running agents (Kwon and Hodgins, 2010) or finer detailed models of walkers (Patnaik and Umanand, 2015) could use different braking models.

Escape zone is the area where Agent1 is able to choose their own action to avoid collision, without needing to trust Agent2 to behave in any particular way. If w_2 is the width of Agent2, which




FIGURE 12
Distances (d_{crash} and d_{escape}) and zones predicted by the PTR model for different car speeds v at lower speeds. If the vehicle is set to drive at the same speed $v = 1.1 = v_{ped}$ m/s as a pedestrian, then the PTR zones closely match Hall proxemic zones, as shown by the vertical red line (Camara and Fox, 2020).

Agent1 must cross at speed v_1 to pass first, the escape zone is then the set $\{d: d_{escape} < d\}$ with

$$d_{escape} = v_2 t_1 + w_2 \frac{v_2}{v_1}. \quad (3)$$

Trust zone is the region $\{d: d_{crash} < d < d_{escape}\}$ where the PTR is true. Agent2 can here *choose* to slow down to prevent collision, but Agent1 is incapable of making any action to affect this outcome themselves. This occurs when Agent1 cannot get out of Agent2's way in time to avoid collision, but Agent2 is able to slow to prevent the collision if it chooses to yield.

Note that these zones are not symmetric between Agent1 and Agent2. They describe when Agent1 must trust Agent2. Their roles must be swapped and the zones recomputed to

see when Agent2 must trust Agent1. The crash, escape, and trust zones were mapped to Hall's personal, public, and social zones respectively, for Agent1 (Camara and Fox, 2020). The trust/social zone is the region in which physical trust is required. This may be a prerequisite for some types of interactions, with physical trust being useful to enable the content of the interaction. The evidence for this mapping came from the observation that if an autonomous vehicle Agent2 is set to drive at the same speed as a pedestrian Agent1, the model generates Hall's proxemic social zone to within 4% quantitative accuracy of Hall's original empirical sizes, as shown in Figure 12. This unexpected result, found by studying how an AV should interact with pedestrians, may now explain a larger question about how humans interact with each other and with other types of robots. Hall's zone sizes have previously been only empirical observations but the PTR model now explains them generatively and to 4% accuracy for the first time. We then extended this model for more general human-human interactions and HRI by taking different interaction headings into account and found an error down to 1% (Camara and Fox, 2022).

7. OpenPodcar: Open source hardware for social AV research

Laboratory and VR experiments on pedestrians are limited in realism, so to scale models toward the real world, a real autonomous vehicle is needed. Commercial AVs are very expensive, beyond reach of our and most other labs who may wish to replicate and extend our work. Full size, multi-passenger, open source hardware (OSH) cars have also been designed and in some cases built, including the PixBot² and Evo Tabby³. These are very large projects which require tens of thousands of dollars of components and months of build time, which are again outside the budget of most labs. Several small, RC-scale, cars have been completed and built as OSH including F1Tenth⁴, AutoRally (Goldfain et al., 2019), BARC (Gonzales, 2018), MIT Racecar⁵, MuSHR⁶, (Nakamoto and Kobayashi, 2019), and (Vincke et al., 2021). However, their small sizes mean that the utilities for collisions with pedestrians would be smaller than for other vehicles, so they are not a good experimental substitute for them.

So we have developed a new low-cost, autonomous vehicle research platform, targeted at pedestrian-AV research and used to host new versions of the Sequential Chicken model based on proxemic small negative utilities rather than collisions. *OpenPodcar*, shown in Figure 13, is based on an off-the-shelf,




FIGURE 13
OpenPodcar: Open source hardware AV (See preprint Camara et al., 2022).

hard-canopy, mobility scooter donor vehicle. We have released OpenPodcar as open source hardware (OSH, Bonvoisin et al., 2020) together with a full automation open source software (OSS) stack, based on ROS, *gmapping* and *movebase*. This will enable other groups to replicate our complete system and experiments, and to use their own research to extend and contribute to a single shared system, which can evolve over time toward real-world use.

The open hardware release includes step-by-step visual build instructions which enable a typical graduate engineer to build an identical copy of our own build, to within specified tolerances. Open source Arduino (Banzi and Shiloh, 2022), ROS (Cousins et al., 2010), and Gazebo (Koenig and Howard, 2004) software files are included which provide standard interfaces and physical simulation for the vehicle. Open source higher level ROS libraries and configurations are provided to perform pedestrian detection and tracking, SLAM, path planning and control.

The platform is large enough to transport one person at speeds up to 15 km/h, it is designed to be large and fast enough to be useful for both real-world delivery tasks and research into human interaction with the general class of such real-world vehicles, including last-mile people and goods transporters. The build cost is around 7,000 USD from new, or 2,000 USD from used, components.

8. Legal and ethical implications

The Sequential Chicken model has recently been considered in the UK Law Commission's study on Autonomous Vehicles (UK Law Commission, 2019), used to update UK legislation in the area. A key question in this study involving the Sequential Chicken model is "what is the legal status of inflicting negative utilities onto members of the public in public places?". To avoid the freezing robot problem, the model shows that AVs *must* sometimes inflict negative utilities onto pedestrians. This is an unusual situation for a civilian engineered system, as a

² <https://gitlab.com/pixmoving/pixbot>

³ <https://www.openmotors.co/evplatform/>

⁴ <https://f1tenth.org/>

⁵ <https://mit-racecar.github.io/>

⁶ <https://mushr.io/>

fundamental professional responsibility of engineers is to make systems safe for the public and to improve their lives. There is almost no precedent for engineering systems to deliberately cause harm to members of the public—even if doing so helps a larger number of other members of the public as a result.

Actually colliding with pedestrians is obviously undesirable, as for large AVs it may result in injury or death. Such actions might be regarded as pre-meditation on the part of their programmers if they occur due to deliberate programming decisions rather than by programmer error or accident. Causing death by pre-meditated means is generally considered to be the most serious crime—murder—rather than accidental manslaughter.

The Sequential Chicken model shows how frequency and severity of inflicted negative utilities can be traded off by their designers: for example the AV can either deliberately create actual collisions in a tiny set of cases, or it could instead create smaller negative utilities such as proxemic space invasion in a much larger number of cases.

The legal status of inflicting even smaller negative utilities onto members of the public remains unclear. The closest case is perhaps the 2014 controversy of social network Facebook deliberately making some of its users sad by showing them negative news stories, without their consent, as part of a data science experiment. Facebook was widely criticized for this action and it is generally considered to have been unethical if not illegal (Verma, 2014).

Another option for inflicting small negative utilities could be to display images on video billboards, or otherwise draw attention to, pedestrians who are behaving anti-socially toward the vehicle. China's Shenzhen city has recently been reported to use similar methods⁷, while its Hubei province is reported to have deployed spraying water⁸, both to penalize jaywalking.

Issues of personal data and personal customization may arise in future versions of the Sequential Chicken model. The current model assumes that all pedestrians are the same, having the same likely destinations and utility functions. But in reality these may differ, and may be predictable given either the person's unique identity or some weaker class membership, such as their age, gender, or type of clothing. If unique pedestrians can be re-identified then per-pedestrians models can be built; otherwise weaker models can be built from their class membership. Per-pedestrian models will encounter GDPR legislation which requires consent for processing data on individuals; while making predictions based on class membership is a form of bias or prejudice which may encounter equality legislation.

Shenzhen's jaywalking systems (*ibid.*) have also been reported as trialing personal identification *via* face recognition to send cash fines to pedestrians, which is the most pure form of negative utility usually considered in Economics.

The deep learning models reviewed—which are needed for low-level understanding of pedestrians and used as input to the higher level interaction—may be legally and practically problematic in the event of accidents and their subsequent investigations, as their black-box nature usually precludes human understanding and explanation of how their decisions are made (Castelvecchi, 2016). Explainable AI (XAI) (Gunning, 2017) remains a current research area which seeks methods to extract this type of information from neural networks, and would be useful to resolve this issue in future.

9. Discussion and conclusions

The Sequential Chicken model can be used to model and probabilistically predict the behavior of human pedestrians, and to plan interactions with them in order to solve the freezing robot problem. The model shows that if an AV's only available actions are to yield to pedestrians or drive forward, then it must be programmed with a small probability of collision, so as to produce a credible threat which solves the freezing robot problem. However, the same model also shows that the rare large negative utility of a crash can be replaced by more frequent but smaller penalties using human proxemic preferences.

Our proxemics study has provided a generative explanation for the numerical sizes of Hall's empirical zones for the first time, and can now be used to control and adapt the AV's behavior more safely as a small negative utility in the model. Parameters for the model can be inferred empirically from human studies including VR experiments.

There are several limitations with the Sequential Chicken model. For example, it appears to work well when the pedestrian and AV approach at right angles, as when a pedestrian crosses the road in front of an AV, but more complex paths of approach are possible. There is currently no consideration of lanes of a road. In particular, the model does not distinguish between the case of the pedestrian starting from the same side of the road as the car is driving on, vs. starting from the opposite side (c.f. Figure 3). In the latter case, the pedestrian may be making a substantial commitment to crossing by crossing over the *other* lane and arriving in the center of the road before dealing with the AV. Consideration would need to be given to vehicles approaching from both directions in both lanes in busy roads in order to predict the pedestrian's behavior, which may include vehicles having to predict each others' behavior as well as the pedestrian's. Small AVs such as OpenPodcar do not drive on roads at all, but rather share pedestrianized space with pedestrians. We have so far modeled their interactions assuming orthogonal approaches to fit the model, but the trajectories of

7 <https://www.independent.co.uk/tech/china-police-facial-recognition-technology-ai-jaywalkers-fines-text-wechat-weibo-cctv-a8279531.html>

8 <https://abcnews.go.com/International/china-province-spraying-publicly-shaming-jaywalkers-deter/story?id=54607782>

both agents can be here be more complex, including lateral motion such as curving one's path as in Hoogendoorn and Bovy (2004) to let the other pass.

The model currently assumes that all pedestrians have the same parameters, which is unlikely to be true. Rather, different pedestrians have different goals and risk preferences. Advanced driver training for human drivers is largely based around understanding and predicting behavior of other road users as individuals, based on their observable features. We have seen evidence that sequence patterns of higher level events (Camara et al., 2018b,c), if they become similarly detectable with machine vision, can be informative about particular pedestrians' utility functions within the model, which suggests future systems fusing this per-person information with the priors obtained from the VR studies to create more accurate per-person predictions and interactions. Big data now makes it feasible even to learn models of individual other road users encountered multiple times in a local area, potentially creating higher accuracies but also raising ethical questions (Fox et al., 2018). These more detailed models will also introduce uncertainty over parameters which would need to be handled appropriately.

While the above limitations are important for applications to different types of road and pavement cases, the idea of proxemics and trust as a framework for negotiating interactions may be more general than autonomous vehicles, and extend to many other forms of HRI, as proposed in Camara and Fox (2022). For example, for a human working alongside a robot arm, fixed to the ground rather than mobile, the same proxemic trust zones could be defined based on the robot's capabilities, and used to plan interactions. As with the above variations in road scenarios, these extended HRI environments would share the same trust and proxemic concepts, then derive different zone geometries from the specifics of the robots and environments involved.

The OpenPodcar platform will enable other researchers to run this interaction model and extend it with more refined pedestrian models to create more human-like, unfrozen interactions. To operate in real time on OpenPodCar, the Sequential Chicken model requires as input the positions of pedestrians which we have seen can be obtained using mature machine vision and tracking methods.

Future work should thus consider running larger human experiments using the OpenPodcar platform in order to refine the Sequential Chicken model parameters with real-world

interaction settings. We may include more visual features, demographics and environmental data into the model similar to the approach used by Ma et al. (2017) and Rasouli et al. (2019). This work has mainly focused on pedestrian–AV interactions for road-crossing scenarios, but this could be extended to more general dual agent interactions. The legal, societal and ethical implications of this work should also be investigated further.

Author contributions

FC and CF conceived and wrote the manuscript. Both authors contributed to the article and approved the submitted version.

Funding

This work has received funding from EU H2020 interACT: Designing cooperative interaction of automated vehicles with other road users in mixed traffic environments under grant agreement no. 723395 and from InnovateUK grant 5949683 C19-ADVs.

Acknowledgments

The authors would like to thank the Editor and the reviewers for their very useful comments and feedback on this manuscript.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

References

Amos, B., Ludwiczuk, B., and Satyanarayanan, M. (2016). *Openface: A general-purpose face recognition library with mobile applications*. Technical report, CMU-CS-16-118, CMU School of Computer Science. Available online at: <https://cmusatyalab.github.io/openface/>

Arechavaleta, G., Laumond, J.-P., Hicheur, H., and Berthoz, A. (2008). On the nonholonomic nature of human locomotion. *Auton. Robots.* 25, 25–35. doi: 10.1007/s10514-007-9075-2

- Banzi, M., and Shiloh, M. (2022). *Getting Started With Arduino*. Maker Media, Inc.
- Bellotto, N., Dondrup, C., and Hanheide, M. (2015). *Bayestacking: The Bayes tracking library v1.0.5*. Zenodo. Available online at: <https://zenodo.org/record/15825>
- Bewley, A., Ge, Z., Ott, L., Ramos, F., and Upcroft, B. (2016). "Simple online and realtime tracking," in *Proceedings of the IEEE International Conference on Image Processing (ICIP)* (Phoenix, AZ: IEEE), 3464–3468.
- Bonvoisin, J., Molloy, J., Häuer, M., and Wenzel, T. (2020). Standardisation of practices in open source hardware. *arXiv preprint arXiv:2004.07143*. doi: 10.5334/joh.22
- Breuers, S., Beyer, L., Rafi, U., and Leibel, B. (2018). "Detection-tracking for efficient person analysis: the detta pipeline," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (Madrid: IEEE), 48–53.
- Brooks, R. (2017). *The Big Problem With Self-Driving Cars is People and We'll go Out of Our Way to Make the Problem Worse*. IEEE Spectrum.
- Cadena, C., Carlone, L., Carrillo, H., Latif, Y., Scaramuzza, D., Neira, J., et al. (2016). Past, present, and future of simultaneous localization and mapping: towards the robust-perception age. *IEEE Trans. Robot.* 32, 1309–1332. doi: 10.1109/TRO.2016.2624754
- Camara, F., Bellotto, N., Cosar, S., Nathanael, D., Althoff, M., Wu, J., et al. (2020a). Pedestrian models for autonomous driving Part I: low-level models, from sensing to tracking. *IEEE Trans. Intell. Transport. Syst.* 22, 6131–6151. doi: 10.1109/TITS.2020.3006768
- Camara, F., Bellotto, N., Cosar, S., Weber, F., Nathanael, D., Althoff, M., et al. (2020b). Pedestrian models for autonomous driving Part II: high-level models of human behavior. *IEEE Trans. Intell. Transport. Syst.* 22, 5453–5472. doi: 10.1109/TITS.2020.3006767
- Camara, F., Cosar, S., Bellotto, N., Merat, N., and Fox, C. W. (2018a). "Towards pedestrian-AV interaction: method for elucidating pedestrian preferences," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) Workshops* (Madrid).
- Camara, F., Cosar, S., Bellotto, N., Merat, N., and Fox, C. W. (2020c). *Continuous Game Theory Pedestrian Modelling Method for Autonomous Vehicles*. River Publishers. Available online at: <https://www.routledge.com/Human-Factors-in-Intelligent-Vehicles/Olaverri-Monreal-Garcia-Fernandez-Rossetti/p/book/9788770222044>
- Camara, F., Dickinson, P., and Fox, C. (2021). Evaluating pedestrian interaction preferences with a game theoretic autonomous vehicle in virtual reality. *Transport. Res. F Traffic Psychol. Behav.* 78, 410–423. doi: 10.1016/j.trf.2021.02.017
- Camara, F., Dickinson, P., Merat, N., and Fox, C. (2019a). "Towards game theoretic AV controllers: measuring pedestrian behaviour in virtual reality," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) Workshops* (Macau).
- Camara, F., Dickinson, P., Merat, N., and Fox, C. W. (2020d). "Examining pedestrian behaviour in virtual reality," in *Transport Research Arena (TRA) (Conference canceled)*.
- Camara, F., and Fox, C. (2020). Space invaders: Pedestrian proxemic utility functions and trust zones for autonomous vehicle interactions. *Int. J. Soc. Rob.* 13, 1929–1949. doi: 10.1007/s12369-020-00717-x
- Camara, F., and Fox, C. (2022). "Extending quantitative proxemics and trust to HRI," in *Proceedings of the 31st IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. (Best Student Award Paper Finalist and KROS Interdisciplinary Research Award in Social Human-Robot Interaction Finalist) (Naples).
- Camara, F., Giles, O., Madigan, R., Rothmüller, M., Holm Rasmussen, P., Vendelbo-Larsen, S. A., et al. (2018b). "Filtration analysis of pedestrian-vehicle interactions for autonomous vehicles control," in *Proceedings of the International Conference on Intelligent Autonomous Systems (IAS-15) Workshops* (Baden-Baden).
- Camara, F., Giles, O., Madigan, R., Rothmüller, M., Rasmussen, P. H., Vendelbo-Larsen, S. A., et al. (2018c). "Predicting pedestrian road-crossing assertiveness for autonomous vehicle control," in *Proceedings of the IEEE International Conference on Intelligent Transportation Systems (ITSC)* (Maui, HI: IEEE).
- Camara, F., Merat, N., and Fox, C. (2019b). "A heuristic model for pedestrian intention estimation," in *Proceedings of the IEEE International Conference on Intelligent Transportation Systems (ITSC)* (Auckland: IEEE).
- Camara, F., Romano, R., Markkula, G., Madigan, R., Merat, N., and Fox, C. (2018d). "Empirical game theory of pedestrian interaction for autonomous vehicles," in *Measuring Behavior: 11th International Conference on Methods and Techniques in Behavioral Research*. Manchester Metropolitan University (Manchester).
- Camara, F., Waltham, C., Churchill, D. and Fox, C. (2022) OpenPodcar: An open source vehicle for self-driving car research. *arXiv [Preprint]*. arXiv: 2205.04454. Available online at: <https://arxiv.org/abs/2205.04454v1> (accessed May 10, 2022).
- Cao, Z., Hidalgo Martinez, G., Simon, T., Wei, S., and Sheikh, Y. A. (2019). "Openpose: realtime multi-person 2d pose estimation using part affinity fields," in *IEEE Transactions on Pattern Analysis and Machine Intelligence*. Available online at: <https://github.com/CMU-Perceptual-Computing-Lab/openpose>
- Cao, Z., Simon, T., Wei, S., and Sheikh, Y. (2017). "Realtime multi-person 2d pose estimation using part affinity fields," in *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)* (Honolulu, HI: IEEE), 1302–1310.
- Castelvecchi, D. (2016). Can we open the black box of AI? *Nat. News* 538, 20. doi: 10.1038/538020a
- Catapult, T. S. (2017). *Market Forecast for Connected and Autonomous Vehicles*. Available online at: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/642813/15780_TSC_Market_Forecast_for_CAV_Report_FINAL.pdf
- Chater, N., Misyak, J., Watson, D., Griffiths, N., and Mouzakitis, A. (2018). Negotiating the traffic: can cognitive science help make autonomous vehicles a reality? *Trends Cogn. Sci.* 22, 93–95. doi: 10.1016/j.tics.2017.11.008
- Chen, P., Wu, C., and Zhu, S. (2016). Interaction between vehicles and pedestrians at uncontrolled mid-block crosswalks. *Saf. Sci.* 82, 68–76. doi: 10.1016/j.ssci.2015.09.016
- Cleac'h, L., Schwager, M., and Manchester, Z. (2021). Algames: a fast augmented lagrangian solver for constrained dynamic games. *Auton. Rob.* 46, 201–215. doi: 10.1007/s10514-021-10024-7
- Cousins, S., Gerkey, B., Conley, K., and Garage, W. (2010). Sharing software with ROS. *IEEE Rob. Autom. Mag.* 17, 12–14. doi: 10.1109/MRA.2010.936956
- Deb, S., Strawderman, L., Carruth, D. W., DuBien, J., Smith, B., and Garrison, T. M. (2017). Development and validation of a questionnaire to assess pedestrian receptivity toward fully autonomous vehicles. *Transport. Res. C Emerg. Technol.* 84, 178–195. doi: 10.1016/j.trc.2017.08.029
- Deo, N., and Trivedi, M. M. (2020). Trajectory forecasts in unknown environments conditioned on grid-based plans. *arXiv preprint arXiv:2001.00735*. doi: 10.48550/arXiv.2001.00735
- Dias, C., Abdullah, M., Sarvi, M., Lovreglio, R., and Alhajyaseen, W. (2019). Modeling and simulation of pedestrian movement planning around corners. *Sustainability* 11, 5501. doi: 10.3390/su11195501
- Fajen, B. R., and Warren, W. H. (2003). Behavioral dynamics of steering, obstacle avoidance, and route selection. *J. Exp. Psychol.* 343. doi: 10.1037/0096-1523.29.2.343
- Figliozzi, M. A., Mahmassani, H. S., and Jaillet, P. (2008). "Repeated auction games and learning dynamics in electronic logistics marketplaces: complexity, bounded rationality, and regulation through information," in *Managing Complexity: Insights, Concepts, Applications* (Funchal: Springer), 137–175.
- Flad, M., Fröhlich, L., and Hohmann, S. (2017). Cooperative shared control driver assistance systems based on motion primitives and differential games. *IEEE Trans. Hum. Mach. Syst.* 47, 711–722. doi: 10.1109/THMS.2017.2700435
- Fox, C. W., Camara, F., Markkula, G., Romano, R. A., Madigan, R., and Merut, N. (2018). "When should the chicken cross the road? – game theory for autonomous vehicle – human interactions," in *Proceedings of the 4th International Conference on Vehicle Technology and Intelligent Transport Systems*, eds M. Helfert and O. Gusikhin (Funchal: SciTePress), 431–439.
- Girshick, R., Donahue, J., Darrell, T., and Malik, J. (2014). "Rich feature hierarchies for accurate object detection and semantic segmentation," in *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)* (Columbus, OH: IEEE), 580–587.
- Goldfain, B., Drews, P., You, C., Barulic, M., Velez, O., Tsiotras, P., et al. (2019). Autally: an open platform for aggressive autonomous driving. *IEEE Control Syst. Mag.* 39, 26–55. doi: 10.1109/MCS.2018.2876958
- Gonzales, J. (2018). *Planning and Control of Drift Maneuvers with the Berkeley Autonomous Race Car* (Ph.D. thesis). University of California at Berkeley.
- Gunning, D. (2017). "Explainable artificial intelligence (XAI)," in *Defense Advanced Research Projects Agency (DARPA)*.
- Hall, E. T. (1966). *The Hidden Dimension*. Anchor Books.

- He, K., Gkioxari, G., Dollár, P., and Girshick, R. (2017). "Mask R-CNN," in *Proceedings of the IEEE International Conference on Computer Vision (ECCV)* (Venice: IEEE), 2980–2988.
- Holland, C., and Hill, R. (2007). The effect of age, gender and driver status on pedestrians' intentions to cross the road in risky situations. *Accident Anal. Prevent.* 39, 224–237. doi: 10.1016/j.aap.2006.07.003
- Hoogendoorn, S., and Bovy, P. (2004). Pedestrian route-choice and activity scheduling theory and models. *Transport. Res. B Methodol.* 38, 169–190. doi: 10.1016/S0191-2615(03)00007-9
- House of Lords (2017). *Connected and Autonomous Vehicles: The Future?* Available online at: <https://publications.parliament.uk/pa/ld201617/ldselect/ldscitech/115/115.pdf>
- Jafari, B., Rabiei, E., Diaconescu, M. A., Masoomi, H., Fiondella, L., and Mosleh, A. (2018). "A survey on autonomous vehicles interactions with human and other vehicles," in *14th PSAM International Conference on Probabilistic Safety Assessment and Management*.
- Karasev, V., Ayvaci, A., Heisele, B., and Soatto, S. (2016). "Intent-aware long-term prediction of pedestrian motion," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)* (Stockholm: IEEE), 2543–2549.
- Kato, S., Takeuchi, E., Ishiguro, Y., Ninomiya, Y., Takeda, K., and Hamada, T. (2015). An open approach to autonomous vehicles. *IEEE Micro* 35, 60–68. doi: 10.1109/MM.2015.133
- Kim, C., and Langari, R. (2014). Game theory based autonomous vehicles operation. *Int. J. Vehicle Design* 65, 360–383. doi: 10.1504/IJVD.2014.063832
- Kim, T. J. (2018). Automated autonomous vehicles: prospects and impacts on society. *J. Transport. Technol.* 8, 137–150. doi: 10.4236/jts.2018.83008
- Kirkpatrick, K. (2022). Still waiting for self-driving cars. *Commun. ACM* 65, 12–14. doi: 10.1145/3516517
- Kitani, K. M., Ziebart, B. D., Bagnell, J. A., and Hebert, M. (2012). "Activity forecasting," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 201–214.
- Koenig, N., and Howard, A. (2004). "Design and use paradigms for Gazebo, an open-source multi-robot simulator," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (IEEE Cat. No. 04CH37566), Vol. 3 (Sendai: IEEE), 2149–2154.
- Kooij, J. F. P., Schneider, N., Flohr, F., and Gavrila, D. M. (2014). "Context-based pedestrian path prediction," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 618–633.
- Koschi, M., Pek, C., Beikirch, M., and Althoff, M. (2018). "Set-based prediction of pedestrians in urban environments considering formalized traffic rules," in *Proceedings of the IEEE International Conference on Intelligent Transportation Systems (ITSC)* (Maui, HI: IEEE).
- Kruse, E., Gutsche, R., and Wahl, F. M. (1997). "Acquisition of statistical motion patterns in dynamic environments and their application to mobile robot motion planning," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Vol. 2, 712–717.
- Kwon, T., and Hodgins, J. K. (2010). "Control systems for human running using an inverted pendulum model and a reference motion capture sequence," in *Symposium on Computer Animation*, 129–138.
- Lambert, D. (2004). *Body Language*. Harper Collins.
- Lee, J. D., and See, K. A. (2004). Trust in automation: Designing for appropriate reliance. *Hum. Factors* 46, 50–80. doi: 10.1518/hfes.46.1.50.30392
- Li, N., Kolmanovsky, I., Girard, A., and Yildiz, Y. (2018). "Game theoretic modeling of vehicle interactions at unsignalized intersections and application to autonomous vehicle control," in *Annual American Control Conference (ACC)*, 3215–3220.
- Lyubenov, D. (2011). Research of the stopping distance for different road conditions. *Transport. Problems* 6, 119–126. Available online at: http://www.transportproblems.polsl.pl/Archiwum/2011/zeszyt4/2011t6z4_14.pdf
- Ma, W.-C., Huang, D.-A., Lee, N., and Kitani, K. M. (2017). "Forecasting interactive dynamics of pedestrians with fictitious play," in *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)* (Honolulu, HI: IEEE), 4636–4644.
- Madigan, R., Nordhoff, S., Fox, C., Amini, R. E., Louw, T., Wilbrink, M., et al. (2019). Understanding interactions between automated road transport systems and other road users: a video analysis. *Transport. Res. F Traffic Psychol. Behav.* 66, 196–213. doi: 10.1016/j.trf.2019.09.006
- Mavrogiannis, C., Alves-Oliveira, P., Thomason, W., and Knepper, R. A. (2022). Social momentum: design and evaluation of a framework for socially competent robot navigation. *ACM Trans. Hum. Rob. Interact.* 11, 1–37. doi: 10.1145/3495244
- Mavrogiannis, C. I., and Knepper, R. A. (2019). Multi-agent path topology in support of socially competent navigation planning. *Int. J. Rob. Res.* 38, 338–356. doi: 10.1177/0278364918781016
- Michieli, U., and Badia, L. (2018). "Game theoretic analysis of road user safety scenarios involving autonomous vehicles," in *IEEE 29th Annual International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)* (Bologna: IEEE), 1377–1381.
- Millard-Ball, A. (2018). Pedestrians, autonomous vehicles, and cities. *J. Planning Educ. Res.* 38, 6–12. doi: 10.1177/0739456X16675674
- Misyak, J. B., Melkonyan, T., Zeitoun, H., and Chater, N. (2014). Unwritten rules: virtual bargaining underpins social interaction, culture, and society. *Trends Cogn. Sci.* 18, 512–519. doi: 10.1016/j.tics.2014.05.010
- Morgenstern, O., and Von Neumann, J. (1953). *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press.
- Murphy, R. R., Gandudi, V. B. M., and Adams, J. (2020). *Robots are Playing Many Roles in the Coronavirus Crisis-and Offering Lessons for Future Disasters*. The Conversation.com.
- Na, X., and Cole, D. J. (2014). Game-theoretic modeling of the steering interaction between a human driver and a vehicle collision avoidance controller. *IEEE Trans. Hum. Mach. Syst.* 45, 25–38. doi: 10.1109/THMS.2014.2363124
- Nakamoto, N., and Kobayashi, H. (2019). "Development of an open-source educational and research platform for autonomous cars," in *IECON-45th Annual Conference of the IEEE Industrial Electronics Society, Vol. 1* (Lisbon: IEEE), 6871–6876.
- Norton, P. (2021). *Autonorama: The Illusory Promise of High-tech Driving*. Island Press.
- Núñez Velasco, J. P., Farah, H., van Arem, B., and Hagenzieker, M. P. (2019). Studying pedestrians' crossing behavior when interacting with automated vehicles using virtual reality. *Transport. Res. F Traffic psychol. Behav.* 66, 1–14. doi: 10.1016/j.trf.2019.08.015
- Osborne, M. J., and Rubinstein, A. (1994). *A Course in Game Theory*. MIT Press.
- Owens, J. M., Greene-Roesel, R., Habibovic, A., Head, L., and Apricio, A. (2018). "Reducing conflict between vulnerable road users and automated vehicles," in *Road Vehicle Automation 4* (Springer), 69–75.
- Papadopoulos, A. V., Bascetta, L., and Ferretti, G. (2013). "Generation of human walking paths," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 1676–1681.
- Parkin, J., Clark, B., Clayton, W., Ricci, M., and Parkhurst, G. (2016). *Understanding interactions between autonomous vehicles and other road users: a literature review*. Technical report. University of the West of England, Bristol. Available online at: <http://eprints.uwe.ac.uk/29153>
- Patnaik, L., and Umanand, L. (2015). Physical constraints, fundamental limits, and optimal locus of operating points for an inverted pendulum based actuated dynamic walker. *Bioinspirat. Biomimet.* 10, 064001. doi: 10.1088/1748-3190/10/6/064001
- Portouli, E., Nathanael, D., and Marmaras, N. (2014). Drivers' communicative interactions: on-road observations and modelling for integration in future automation systems. *Ergonomics* 57, 1795–1805. doi: 10.1080/00140139.2014.952349
- Puydupin-Jamin, A., Johnson, M., and Bretl, T. (2012). "A convex approach to inverse optimal control and its application to modeling human locomotion," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)* (Saint Paul, MN: IEEE), 531–536.
- Qi, C. R., Yi, L., Su, H., and Guibas, L. J. (2017). "Pointnet++: deep hierarchical feature learning on point sets in a metric space," in *Advances in Neural Information Processing Systems*, Vol. 30.
- Rahmati, Y., and Talebpour, A. (2018). Learning-based game theoretical framework for modeling pedestrian motion. *Phys. Rev. E* 98, 032312. doi: 10.1103/PhysRevE.98.032312
- Rahmati, Y., Talebpour, A., Mittal, A., and Fishelson, J. (2020). Game theory-based framework for modeling human-vehicle interactions on the road. *Transp. Res. Rec.* 2674, 701–713. doi: 10.1177/0361198120931513
- Rasmussen, C. E., and Williams, C. K. I. (2005). *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press.
- Rasouli, A., Kotseruba, I., and Tsotsos, J. K. (2018). Understanding pedestrian behavior in complex traffic scenes. *IEEE Trans. Intell. Vehicles* 3, 61–70. doi: 10.1109/TIV.2017.2788193
- Rasouli, A., Kotseruba, I., and Tsotsos, J. K. (2019). "Pedestrian action anticipation using contextual feature fusion in stacked RNNs," in *British Machine Vision Conference (BMVC)*. Available online at: <https://www.bmvc2019.org/wp-content/uploads/papers/0283-paper.pdf>

- Rasouli, A., and Tsotsos, J. K. (2018). Joint attention in driver-pedestrian interaction: from theory to practice. *arXiv[Preprint].arXiv:1802.02522*. doi: 10.48550/arXiv.1802.02522
- Rasouli, A., and Tsotsos, J. K. (2020). Autonomous vehicles that interact with pedestrians: a survey of theory and practice. *IEEE Trans. Intell. Transport. Syst.* 21, 900–918. doi: 10.1109/ITITS.2019.2901817
- Ravindran, R., Santora, M. J., and Jamali, M. M. (2020). Multi-object detection and tracking, based on dnn, for autonomous vehicles: a review. *IEEE Sens. J.* 21, 5668–5677. doi: 10.1109/JSEN.2020.3041615
- Redmon, J., Divvala, S., Girshick, R., and Farhadi, A. (2016). “You only look once: unified, real-time object detection,” in *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)* (Las Vegas, NV: IEEE), 779–788.
- Redmon, J., and Farhadi, A. (2017). “Yolo9000: better, faster, stronger,” in *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)* (Honolulu, HI: IEEE), 7263–7271.
- Redmon, J., and Farhadi, A. (2018). YoloV3: an incremental improvement. *arXiv preprint arXiv:1804.02767*. doi: 10.48550/arXiv.1804.02767
- Rehder, E., Wirth, F., Lauer, M., and Stiller, C. (2018). “Pedestrian prediction by planning using deep neural networks,” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)* (Brisbane, QLD: IEEE), 5903–5908.
- Ren, S., He, K., Girshick, R., and Sun, J. (2017). Faster r-cnn: Towards real-time object detection with region proposal networks. *IEEE Trans. Pattern Anal. Mach. Intell.* 39, 1137–1149. doi: 10.1109/TPAMI.2016.2577031
- Rios-Martinez, J., Spalanzani, A., and Laugier, C. (2015). From proxemics theory to socially-aware navigation: a survey. *Int. J. Soc. Rob.* 7, 137–153. doi: 10.1007/s12369-014-0251-1
- SAE International (2019). *SAE J3016 Levels of Driving Automation*. Available online at: <https://www.sae.org/news/2019/01/sae-updates-j3016-automated-driving-graphic>
- Schwartz, W., Pierson, A., Alonso-Mora, J., Karaman, S., and Rus, D. (2019). Social behavior for autonomous vehicles. *Proc. Natl. Acad. Sci. U.S.A.* 116, 24972–24978. doi: 10.1073/pnas.1820676116
- Simon, T., Joo, H., Matthews, I., and Sheikh, Y. (2017). “Hand keypoint detection in single images using multiview bootstrapping,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (Honolulu, HI: IEEE), 1145–1153.
- Škugor, B., Topić, J., Deur, J. S., Ivanović, V., and Tseng, E. (2020). “Analysis of a game theory-based model of vehicle-pedestrian interaction at uncontrolled crosswalks,” in *2020 International Conference on Smart Systems and Technologies (SST)*, 73–81.
- Sucha, M. (2014). “Road users’ strategies and communication: driver-pedestrian interaction,” in *Transport Research Arena (TRA)*, Vol. 1.
- Talebpoor, A., Mahmassani, H. S., and Hamdar, S. H. (2015). Modeling lane-changing behavior in a connected environment: a game theory approach. *Transport. Res. C Emerg. Technol.* 59, 216–232. doi: 10.1016/j.trc.2015.07.007
- Tamura, Y., Le, P. D., Hitomi, K., Chandrasiri, N. P., Bando, T., Yamashita, A., et al. (2012). “Development of pedestrian behavior model taking account of intention,” in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (Vilamoura-Algarve: IEEE), 382–387.
- Thomas, J. A., and Walton, D. (2007). Measuring perceived risk: self-reported and actual hand positions of suv and car drivers. *Transport. Res. F* 10, 201–207. doi: 10.1016/j.trf.2006.10.001
- Thomaz, A., Hoffman, G., and Cakmak, M. (2016). Computational human-robot interaction. *Found. Trends Robot* 4, 105–223. doi: 10.1561/9781680832099
- Thrun, S., Burgard, W., and Fox, D. (2005). *Probabilistic Robotics*. MIT Press.
- Tian, R., Li, N., Kolmanovsky, I., Yildiz, Y., and Girard, A. R. (2020). “Game-theoretic modeling of traffic in unsignalized intersection network for autonomous vehicle control verification and validation,” in *IEEE Transactions on Intelligent Transportation Systems*.
- Tian, Z., Gao, X., Su, S., Qiu, J., Du, X., and Guizani, M. (2019). Evaluating reputation management schemes of internet of vehicles based on evolutionary game theory. *IEEE Trans. Vehicular Technol.* 68, 5971–5980. doi: 10.1109/TVT.2019.2910217
- Trautman, P., and Krause, A. (2010). “Unfreezing the robot: Navigation in dense, interacting crowds,” in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (Taipei: IEEE), 797–803.
- UK Law Commission (2019). *Automated Vehicles: A Joint Preliminary Consultation Paper*. Available online at: https://s3-eu-west-2.amazonaws.com/lawcom-prod-storage-11jsxou24uy7q/uploads/2018/11/6.5066_LC_AV-Consultation-Paper-5-November_061118_WEB-1.pdf
- Verma, I. M. (2014). Editorial expression of concern: Experimental evidence of massive scale emotional contagion through social networks. *Proc. Natl. Acad. Sci. U.S.A.* 111, 10779. doi: 10.1073/pnas.1412583111
- Vincke, B., Rodriguez Florez, S., and Aubert, P. (2021). An open-source scale model platform for teaching autonomous vehicle technologies. *Sensors* 21, 3850. doi: 10.3390/s21113850
- Wade, M. (2018). *Silicon Valley Is Winning the Race to Build the First Driverless Cars*. Available online at: <https://theconversation.com/silicon-valley-is-winning-the-race-to-build-the-first-driverless-cars-91949>
- Wadud, Z., MacKenzie, D., and Leiby, P. (2016). Help or hindrance? the travel, energy and carbon impact of highly automated vehicles. *Transport. Res. A* 86, 1–18. doi: 10.1016/j.tra.2015.12.001
- Wang, Y., Kitani, K., and Weng, X. (2021). “Joint object detection and multi-object tracking with graph neural networks,” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)* (Xi’an: IEEE), 13708–13715.
- Wei, S.-E., Ramakrishna, V., Kanade, T., and Sheikh, Y. (2016). “Convolutional pose machines,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (Las Vegas, NV: IEEE), 4724–4732.
- Weng, X., Wang, Y., Man, Y., and Kitani, K. M. (2020). “Gnn3dmot: graph neural network for 3d multi-object tracking with 2d-3d multi-feature learning,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (Seattle, WA: IEEE), 6499–6508.
- Wu, J., Ruenz, J., and Althoff, M. (2018). “Probabilistic map-based pedestrian motion prediction taking traffic participants into consideration,” in *Proceedings of the IEEE Intelligent Vehicles Symposium (IV)* (Changshu: IEEE), 1285–1292.
- Zhang, Y., Hang, P., Huang, C., and Lv, C. (2022). Human-like interactive behavior generation for autonomous vehicles: a Bayesian game-theoretic approach with Turing test. *Adv. Intell. Syst.* 4, 2100211. doi: 10.1007/978-981-16-8048-9