UNIVERSITY OF LEEDS

This is a repository copy of *Continuous Game Theory Pedestrian Modelling Method for Autonomous Vehicles*.

White Rose Research Online URL for this paper:
https://eprints.whiterose.ac.uk/162717/

Version: Accepted Version

Book Section:

Camara, F orcid.org/0000-0002-2655-1228, Cosar, S, Bellotto, N et al. (2 more authors) (2020) Continuous Game Theory Pedestrian Modelling Method for Autonomous Vehicles. In: Olaverri-Monreal, C, García-Fernández, F and Rossetti, RJF, (eds.) Human Factors in Intelligent Vehicles. River Publishers . ISBN 978-8770222044

© 2020 River Publishers. This is an author produced version of a book chapter published in Human Factors in Intelligent Vehicles. Uploaded with permission from the publisher.

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

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.

White Rose university consortium
Universities of Leeds, Sheffield & York

eprints@whiterose.ac.uk
https://eprints.whiterose.ac.uk/

Continuous Game Theory Pedestrian Modelling Method for Autonomous Vehicles

Fanta Camara[1,2], Serhan Cosar[3], Nicola Bellotto[2],
Natasha Merat[1] and Charles W. Fox[1,2,4]

[1] Institute for Transport Studies, University of Leeds, UK
[2] School of Computer Science, University of Lincoln, UK
[3] Institute of Engineering Sciences, De Montfort University, UK
[4] Ibex Automation Ltd., UK

Abstract. Autonomous Vehicles (AVs) must interact with other road users. They must understand and adapt to complex pedestrian behaviour, especially during crossings where priority is not clearly defined. This includes feedback effects such as modelling a pedestrian's likely behaviours resulting from changes in the AVs behaviour. For example, whether a pedestrian will yield if the AV accelerates, and vice versa. To enable such automated interactions, it is necessary for the AV to possess a statistical model of the pedestrian's responses to its own actions. A previous work demonstrated a proof-of-concept method to fit parameters to a simplified model based on data from a highly artificial discrete laboratory task with human subjects. The method was based on lidar-based person tracking, game theory, and Gaussian process analysis. The present study extends this method to enable analysis of more realistic *continuous* human experimental data. It shows for the first time how game-theoretic predictive parameters can be fit into pedestrians natural and continuous motion during road-crossings, and how predictions can be made about their interactions with AV controllers in similar real-world settings.

Keywords: Pedestrian Crossing Behaviour; Autonomous Vehicles; Interactions; Game Theory

1 Introduction

Understanding pedestrian behaviour is now of upmost importance for Autonomous Vehicles (AVs) [5]. The potential future deployment of AVs is currently creating much enthusiasm [4][43], as such vehicles would make transportation more efficient [22]. Huge improvements have been made on robotic localisation and mapping problems using simultaneous localisation and mapping (SLAM) algorithms [6][38], together with new, cheap sensors, computation technologies, free and open-source software implementations [20] [42]. 'Self-driving' cars can now localise themselves and navigate by planning and controlling their routes on some roads, promising a future society with a better mobility system with less accidents and traffic in cities [22].

But before any fully self-driving revolution happens, AVs must share space with and will be challenged by human drivers and pedestrians, who are much harder to model and act upon than passive environments. Full self-driving must include this ability as well as the now-mature localisation, planning and routing technologies. Decades of research on human interaction in Transport Psychology and Human Factors has not yet been translated into robotic control systems, and many questions are still unanswered.

In most current ‘self-driving’ systems, for safety and legal reasons, pedestrians are considered as obstacles, such that the vehicle always stops for them. But recent real-world AV studies have shown that pedestrians may then take advantage of this predictable behaviour [27] [25] [5], pushing in front of them for priority eventually in *every* negotiation, so that the vehicles then make no progress. This has become known as the Freezing Robot Problem (FRP)[39].

Real human driving is massively more complex than simply mapping, localising and path planning. It is considered an art form by advanced practitioners such as members of the Institute for Advanced Motorists and other advanced drivers such as high-speed police and ambulance drivers [17]. In their training, these practitioners emphasise the human psychological processes involved in reading and predicting the behaviours of other road users as the most important skill of human drivers. Can you tell if a pedestrian is assertive enough to risk stepping out in front of you from their body language, their facial expressions, even their clothes and demographics? Road users have different utility functions, ranging from timid pedestrians likely to give way to all oncoming traffic, though to business-people late for a meeting or patients for an urgent medical appointment becoming much more assertive and risk-taking. Drivers must also consider the psychological effects of their own actions. Speeding up and slowing down are not just ways to control one’s own progress, but also send information about our own personality and risk preferences to pedestrians engaged in such negotiations for priority, along with other possible signals including lateral road positing, and more conventional signals such as flashing indicator lights and headlights, and driver face and arm expressions.

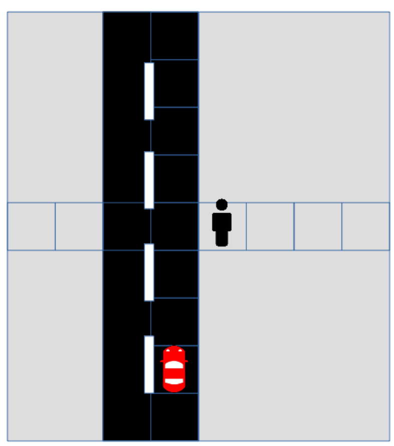

Fig. 1. Two agents negotiating for priority at an intersection

To progress towards automation of such understandings, Fox et al. [18] proposed and solved a simple game-theoretical mathematical model of the unsigned road-crossing scenarios represented in Figs. 1 and 7. This model, based on the famous game of 'chicken', is called 'sequential chicken'. In this model, two agents – which may be pedestrians and/or vehicles – compete for space at an unsigned intersection, using only their positions to signal information to one another. Time, space and actions are discretised and it is assumed that both players have equal utility functions and know this to be the case. The model leaves open free parameters specifying the utility function for human players. Camara et al. [11] then asked human subjects to play sequential chicken as a board game, and developed a statistical method to fit parameters to the mathematical model to describe and predict their behaviours. In [9], the same authors extended this experiment to the case of human subjects playing a physical version of the board game, moving their bodies between discrete squares on and near the road at discrete time turns, integrating their positions into the sequential chicken model via lidar sensors, support vector machines, and Bayesian tracking.

Contributions: The present chapter is a methods study which presents a new, full stack approach to measuring and modelling natural, continuous time and continuous space pedestrian interactions. It shows how to infer pedestrian preferences for time delays and collisions from their body motions as tracked by lidar. Inferred parameters could then be used in AV controllers during pedestrian interactions. First, pedestrian tracking is used to estimate the trajectories of the agents involved in semi-structured human–human interactions while playing the sequential chicken model. Second, optimal strategies are computed using the game theory model in [18]. Lastly, parameters of the interactions are inferred by comparison to optimal strategies, using Gaussian process regression over the parameter space. This study is intended to illustrate a proof-of-concept of this full-stack *method*: more detailed and controlled experiments will be needed to obtain robust parameters results and to learn about variations in parameters between different classes of pedestrians. The demonstrated method could also be used to model and measure pedestrian/pedestrian, human–driver–vehicle/AV, and human–driver/human–driver and AV/AV interactions as well as the primarily intended pedestrian/AV case.

This work is part of the EU H2020 interACT project[1] with a consortium of European partners investigating on the future deployment of AVs in mixed traffic environments with human drivers, cyclists and pedestrians. The overall aims of the project are to understand the behaviour of other road users, and how AVs could interact with them in a safe and efficient manner, and to propose new external Human–Machine Interface (eHMI) solutions that could facilitate the communication between AVs and people.

2 Related work

2.1 Pedestrian crossing behaviour

A review on different approaches for pedestrian behaviour modelling is provided in [8]. Methods of pedestrian behaviour analysis are often performed via video recording, semi-structured interviews and VR recording. Previous studies on pedestrian crossing behaviour can be found in [19][29][32]. For example, Gorrini et al. [19] analysed video

[1] `https://www.interact-roadautomation.eu/`

data of interaction between pedestrians and vehicles at an unsignalized intersection using semi-automatic tracking. Their study showed that pedestrian crossing behaviour can be divided into three phases: approaching (stable speed), appraising (deceleration due to evaluation of speed and distance of oncoming vehicles) and crossing (acceleration). Papadimitriou et al. [29] compared observed and declared behaviour of pedestrians at different crossing areas, as a method to assess pedestrian risk-taking while crossing. They found that their observed behaviour is in accordance with their declared behaviours from a questionnaire survey and they report that female and male participants have similar crossing behaviour. Many studies such as [35] were focused on the evaluation of speed, TTC (Time To Collision), gap acceptance and communication means (e.g., eye contact and motion pattern) of the road users. Some other studies (e.g., [15]) have suggested that for autonomous vehicles, some apparently intuitive human communication styles might not be necessary for interactions with pedestrians. Dey and Terken [15] showed that facial communication cues such as eye contact do not play a major role in pedestrian crossing behaviour, and that the motion pattern and behaviour of vehicles are more important. The field study in [34] showed similar results with an ‘unmanned’ vehicle, suggesting that the same results could be found with autonomous vehicles. Risto et al. [33] showed that vehicle movement is sufficient for indicating the intention of drivers and presented some motion patterns of road users such as advancing, slowing early and stopping short.

2.2 Game theory

Game theory offers a framework for modelling conflict and cooperation between rational decision-makers. It was developed in the 1940s by von Neumann and Morgenstern [28]. Its core concept is (Nash) *equilibrium* which is the pair of strategies (probability distributions over actions to be played) such that none of the players would change their strategy if they knew the other’s strategy. Previous studies in Transport Studies and highway design have applied game theory to several driver behaviour modelling tasks, as reviewed in [16]. Kim et al. [21] developed a mixed-motive game theory model for deciding the strategy chosen by two AVs equipped with adaptive cruise control (ACC). Meng et al. [26] also used game theory for modelling AV lane changing maneuvers. Rakha et al. [30] proposed a game theory approach for intersection conflicts management with reactive agents (the automated vehicles) equipped with ACC systems and a manager agent is used to decide the optimal strategy that increases the overall performance of all the agents. This approach prevents crashes from occurring and it also minimises the time delay in the intersection. Similar to our work, Ma et al. [**?**] computed Nash equilibria using Fictitious Play. Their method differs from ours in that not only their model takes into account pedestrians’ position from a single image but also used some visual features from their appearance as part of the utility function to improve trajectory prediction. Adkins [1] presented an algorithm for intersection management involving up to four self-driving cars communicating with each other. Two motion choices are available for each player (move forward or stop) and an optimised solution using game theory to solve the discrete intersection problem is presented. Turnwald et al. [40] proposed a non-cooperative game theoretic approach to human collision avoidance. Their method differs from ours in that they used a motion capture system to record human motions, a Bootstrap algorithm to compute the confidence intervals and applied a Dynamic Time Warping (DTW) algorithm to measure similarity between the trajectories. Variants of the game of chicken were proposed in [13][27] [31] to solve conflicts between agents at intersections. A cellular automata approach was implemented in [31] and [13] for agents' interactions while [27] focused on the interaction between an AV and a pedestrian.

When multiple equilibria are present in games, standard game theory does not specify how the players should choose the best one. In the above studies, no method is proposed for the players to select which equilibrium to use. Typically this is because Transport Studies seeks to describe macroscopic flows of traffic rather than prescribe actions for individual vehicles, and considers that *any* possible equilibrium is a good description of observed data. For example in [27], the choice for the best solution depends on 'local social norms' which assumes that drivers should have prior knowledge of local customs. Unusually, [18] proposed a novel approach for optimal strategy prescription, called *meta-strategy convergence*. This method begins by choosing an equal-weighted mixture of strategies from all rational equilibria (after removing dominated and asymmetric equilibria where possible). The resulting strategies do not in general form an equilibrium themselves, but by applying fictitious play until convergence, a single equilibrium is obtained upon which it is argued that two rational players should agree without communication. Most of the game theory models reviewed in this section outperform non-game theoretic predictive models [13][24][30][41].

2.3 Pedestrian tracking

Pedestrian tracking plays an important role in many commercial applications but it is still a challenge for computer vision systems because of the multiple uncertainties (e.g., occlusions) due to complex environments [7]. Tracking of pedestrians requires the estimation of non-linear, non-Gaussian problems due to human motion, pedestrian scales, and posture changes. Monte Carlo methods such as particle filtered-based approaches draw a set of samples assigned to a target and perform the data association for multiple targets using probabilistic techniques such as Nearest Neighbor (NN), Multiple Hypothesis Tracking (MHT), JPDAF and PHD-filter [2][7]. Pedestrian tracking is composed of two steps: (*i*) a prediction step to determine the expected position and motion state and (*ii*) an update step to refine the prediction using sensor observations. Tracking has been previously combined with game theory for multi-robot system coordination problems. For instance, Skrzypczyk et al. [37] used non-cooperative games to control a team of mobile robots for a target tracking. When multiple equilibria are present, an arbiter based on the min-max method is used to fairly distribute costs among robots. Li et al. [23] applied cooperative game theory to improve tracking performance for a group of robots, allowing communication between the robots in order to minimise tracking costs and maximise the interests of the overall system of robots. Yan et al. [46] proposed a cooperative non-zero sum game approach for the problem of multi-target tracking for a multi-robot system in dynamic environment.

3 Methods

The present study demonstrates a method to fit parameters of the sequential chicken model to *continuous* human behaviour collected from controlled laboratory pedestrian–pedestrian interactions. The laboratory environment is designed to enable the simplest possible mapping of continuous physical human motions onto the model. Studying pedestrian–pedestrian interactions in place of pedestrian–AV interactions allows us to collect twice as much pedestrian data, and not require us to bias the experiment by involving an AV programmed with its own preferences.

3.1 Human experiment

Eighteen human volunteer subjects (University of Lincoln Computer Science staff and students) were divided into nine pairs, one designated as player Y and the other as player X. Each pair was asked to play a physical version of the sequential chicken game on a plus-maze shaped playing area drawn on an indoor floor as shown in Fig. 2. Player Y was starting from $y = 6$ m and player X from $x = 6$ m such that they were both starting 6 m away from the intersection. Players were instructed that their objective was to pass the intersection as soon as possible, 'as if they were trying to reach their office entrance in a busy pedestrian area', on hearing the command 'go' to begin, given to both players at the same time. Each pair performed five interactions, i.e. 'games'. If both players walk at the same speed, then they collide with each other. Otherwise, one of them must yield to allow the other to pass the intersection point before them. Sometimes, both players try to yield at the same time, which does not break the symmetry, forcing them to continue negotiating one or more times. Players' motions were recorded using a Velodyne 3D lidar. Figure 3 shows an example of the lidar output during the games.

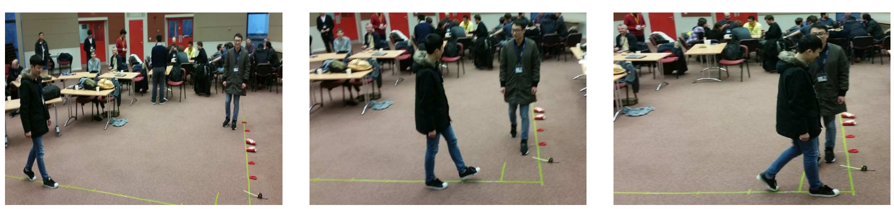

Fig. 2. Two participants playing the game of chicken during the experiment

3.2 Pedestrian detection and tracking

Pedestrian positions and velocities are provided by a robust Bayesian multi-target tracking systems based on 3D lidar detections[47], suitable for real-time, long-range tracking of multiple people in dynamic scenarios. Non-overlapping clusters of adjacent points are extracted based on their 3D Euclidean distance. An adaptive threshold accounts for

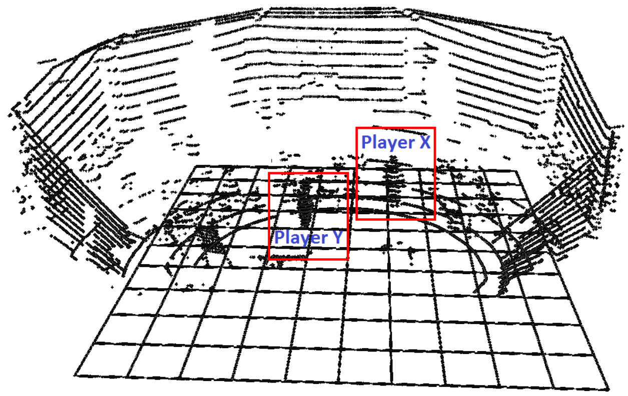

Fig. 3. 3D LIDAR output

the variation in shape and size of the human body in 3D lidar point clouds, which is a function of the person's distance from the sensor. Finally, clusters too large or too small to be humans are discarded by the detector, which outputs the distance and bearing of the cluster's centroid projected on the floor. The information from the detector is processed by a multi-target tracker, including an efficient implementation of Unscented Kalman Filter (UKF) and NN data association to deal with multiple detections simultaneously [3]. The tracker estimates the 2D coordinates and velocities of each pedestrian using a standard prediction-update recursive algorithm. The prediction step is based on the following constant velocity model,

$$\begin{cases} x_k = x_{k-1} + \Delta t\ \dot{x}_{k-1} \\ \dot{x}_k = \dot{x}_{k-1} \\ y_k = y_{k-1} + \Delta t\ \dot{y}_{k-1} \\ \dot{y}_k = \dot{y}_{k-1} \end{cases} \tag{1}$$

where x_k and y_k are the Cartesian coordinates of the target at time t_k, $\dot{x}_k$ and $\dot{y}_k$ are the respective velocities, and $\Delta t = t_k - t_{k-1}$. (The symbols x, y, t in this section are re-used to name different things than in the game theory model sections.) The update step of the estimation uses a 2D polar observation model to represent the position of a detected cluster,

$$\begin{cases} \phi_k &= \tan^{-1}(y_k/x_k) \\ \gamma_k &= \sqrt{x_k^2 + y_k^2} \end{cases} \tag{2}$$

where ϕ_k and γ_k are, respectively, the bearing and the distance of the cluster's centroid with respect to the sensor. More details can be found in [3,47].

Figures 4 to 6 show the filtering process for pedestrian tracks. Like all detection and tracking methods, the system sometimes produces false positives and false negatives. To remove false positives, tracks were filtered to exclude those including any locations outside the plus-maze area, as shown in Fig. 5. Due to occasional false positives with tracks, and false negatives missing tracks, filtering resulted in a collection of 14 games, from 6 different pairs of players, having good and complete tracks for both players together, that are used in the rest of the analysis.

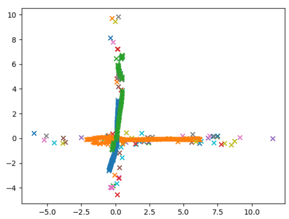

Fig. 4. Unfiltered tracks

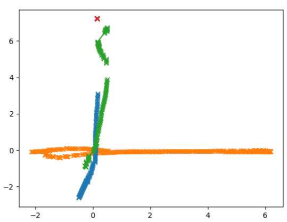

Fig. 5. Filtered tracks

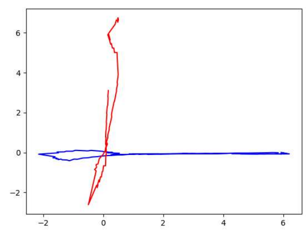

Fig. 6. Tracks assigned to players

3.3 Sequential chicken model

In sequential chicken, two agents called Y and X are driving straight towards each other at right angles as in Fig. 1, such that they will collide unless one of them yields to the other. The sequential chicken model operates on discrete *space* as in Fig. 7; discrete *times* ('turns') during which the agents can adjust their discrete *speeds*, simultaneously selecting between speeds of either 1 square per turn or 2 squares per turn, at each turn. Both agents want to pass the intersection as soon as possible to avoid travel delays, but if they collide, they are both bigger losers as they both receive a negative utility U_{crash}. Otherwise if the players pass the intersection, each receives a time delay penalty $-TU_{\text{time}}$, where T is the time from the start of the game and U_{time} represents the value of saving one turn of travel time. The model assumes that the two players choose their actions (speeds) $a_Y, a_X \in \{1, 2\}$ simultaneously, then implement them simultaneously, at each of several discrete-time turns. There is no lateral motion (positioning within the lanes of the roads) or communication between the agents other than via their visible positions. The game is symmetric, as both players are assumed to know that they have the same utility functions ($U_{\text{crash}}, U_{\text{time}}$), hence they both have the same optimal strategies. These optimal strategies are derivable from game theory together with meta-strategy convergence, via recursion [18]. 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.

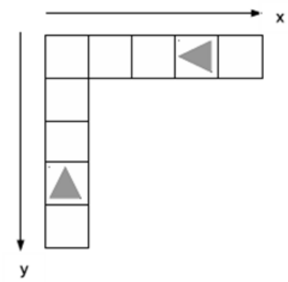

Fig. 7. Sequential Chicken Game

Discretised locations of the players can be represented by (y,x,t) at discretised turn t and their discretised actions $a_Y, a_X \in \{1,2\}$ for speed selection. Similar to the approach used in [10], discretisations are obtained from the continuous data by quantizing continuous position into about 0.1 m locations every 0.09 s turn, by averaging over all locations during that interval; and quantizing actions into SLOW or FAST between each pair of quantised locations according to whether the location change is greater or lower than a 1 m/s threshold.

The new state at turn $t+1$ is given by $(y+a_Y, x+a_X, t+1)$. Define $v_{y,x,t} = (v^Y_{y,x,t}, v^X_{y,x,t})$ 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,

$$v_{y,x,t} = v(\begin{bmatrix} v(y-1,x-1,t+1) & v(y-1,x-2,t+1) \\ v(y-2,x-1,t+1) & v(y-2,x-2,t+1) \end{bmatrix}), \tag{3}$$

which can be solved using dynamic programming assuming meta-strategy convergence equilibrium selection. Under some approximations based on the temporal gauge invariance described in [18], we may remove the dependencies on the time t in our implementation so that only the locations (y,x) are required in computation of $v_{y,x}$ and optimal strategy selection.

In the sequential chicken model, if the two players play optimally, then there must exist a non-zero probability for a collision to occur. Intuitively, if we consider an AV to be one player that always yields, it will make no progress as the other player will always take advantage over it, hence there must be some threats of collision [18].

3.4 Gaussian process parameter posterior analysis

We use Gaussian processes regression [45] to fit the posterior belief over the behavioural parameters of interest, $\theta = (U_{\text{crash}}, U_{\text{time}})$ from the observed data, D. Under the sequential chicken model, M, these are,

$$P(\theta|M,D) = \frac{P(D|\theta,M)P(\theta|M)}{\sum_{\theta'} P(D|\theta',M)P(\theta'|M)}. \tag{4}$$

We assume a flat prior over θ so that,

$$P(\theta|M,D) \propto P(D|\theta,M), \tag{5}$$

which is the data likelihood, given by,

$$P(D|\theta,M) = \prod_{\text{game}} \prod_{\text{turn}} P(d_Y^{\text{game,turn}}|y,x,\theta,M')P(d_X^{\text{game,turn}}|y,x,\theta,M'), \tag{6}$$

where $d_{\text{player}}^{\text{game,turn}}$ are the observed action choices, and y and x are the observed player locations at each *turn* of each *game*. Here M' is a noisy version of the optimal sequential chicken model M, which plays actions from M with probability $(1-s)$ and maximum entropy random actions (0.5 probability of each speed) with probability s. This modification is necessary to allow the model to fit data where human players have made deviations from optimal strategies which would otherwise occur in the data with probability zero. Real humans are unlikely to be perfectly optimal at anytime as they may make mistakes of perception and decision-making. This is a common method to weaken psychological models to allow non-zero probabilities for such mistakes if present.

For a given value of θ, we may compute the optimal strategy for the game by dynamic programming as in Algorithm 1. Optimal strategies are in general probabilistic, and prescribe the $P(d_Y^{\text{game,turn}}|y,x,\theta,M), P(d_{\text{X}}^{\text{game,turn}}|y,x,\theta,M)$ terms to compute the above data likelihood. We then use a Gaussian process with a Radial Basis Function (RBF) kernel to smooth the likelihood function over all values of θ beyond a sample whose values are computed explicitly. In practice, this is performed in the log domain to avoid numerical computation problems with small probabilities. The resulting Gaussian process is then read as the (un-normalized, log) posterior belief over the behavioural parameters $\theta = \{U_{\text{time}}, U_{\text{crash}}\}$ of interest.

Algorithm 1 Optimal solution computation

for U_{crash} in range($U_{crash_{min}}$, $U_{crash_{max}}$) **do**
    for U_{time} in range($U_{time_{min}}$, $U_{time_{max}}$) **do**
        S $\leftarrow$ strategy matrix($NY \times NX \times 2$) for P(player X chooses speed $2|y,x$)
        loglik = 0
        for each game in data **do**
            for each turn in game **do**
                loglik $= \prod_{game} \prod_{turn} (1-s)P(d_Y^{game,turn}|y,x,\theta,M)P(d_X^{game,turn}|y,x,\theta,M) + s(\frac{1}{2})$
            end for
        end for
        Store loglik(U_{crash}, U_{time})
    end for
end for
maxloglik $\leftarrow$ max of loglik(U_{crash}, U_{time})

4 Results

After applying Gaussian process regression and optimising s to maximise the likelihood at the Maximum A Posteriori (MAP) point of θ, the posterior distribution over $\theta = \{U_{\text{crash}}, U_{\text{time}}\}$ is shown in Fig. 8. The MAP estimate of the parameters is then around $U_{\text{crash}} = -220$, $U_{\text{time}} = 465$, at $s = 0.11$. The $-44 : 93 \simeq 1 : 2$ ratio in the utilities means that assuming the noisy model M' the subjects value about a 1/2 turn time delay equally to a crash, and the s value means that the subjects make mistakes from optimal behaviour in 11% of actions. Significance of the results can be seen by inspection of the thin standard deviation widths of 1D slices through the 2D posterior as in Fig. 9. We can only see a small deviation when U_{crash} is too small or too large.

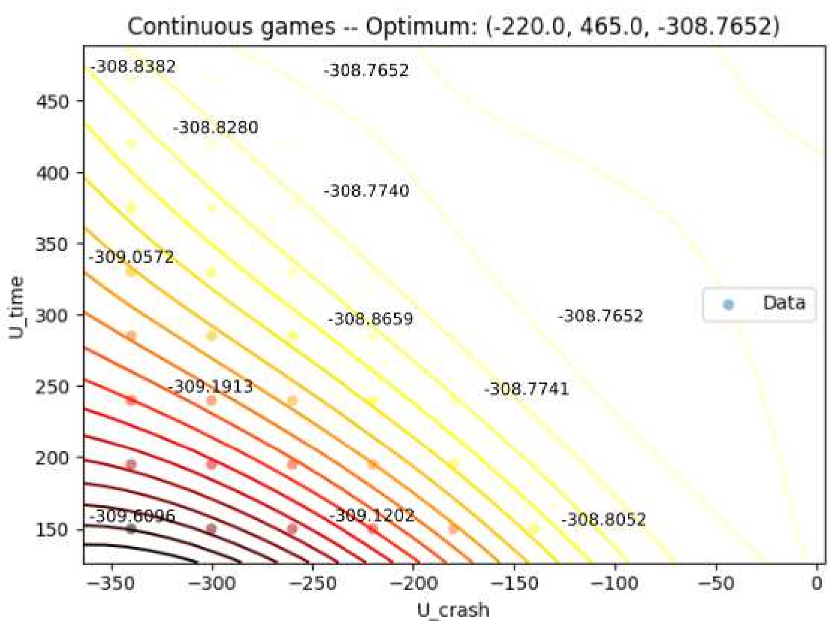

Fig. 8. Gaussian process log-posterior over behavioural parameters.

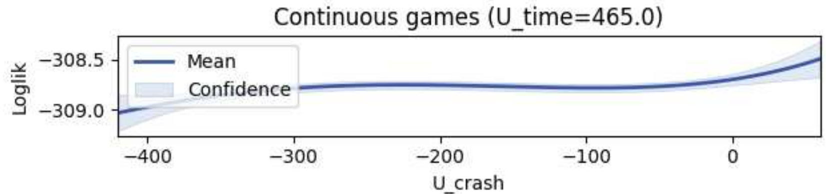

Fig. 9. Slice through the Gaussian process showing standard deviation log-posterior confidence.

The behavioural parameter ($\theta \simeq -\frac{1}{2}$) shows that participants were having higher preferences for time saving rather than for collision avoidance, which is similar to the findings in [9] [11]. As in these studies, the high ratio may be explained by the artificial laboratory nature of the environment: subjects want to win the game and know there is no significant negative utility for collisions as the laboratory environment is designed to be safe. The method is now well developed enough to move to the real world for future studies, and we expect to see lower ratios there, where the cost of collisions with vehicles and other pedestrians is much higher.

5 Discussion

The results shown are from a small sample of data and are intended as a proof-of-concept of the proposed method. This shows how a full stack of real-time detection and tracking, and game theoretic modelling can work together to understand and predict continuous pedestrian interactions with another road user. The data used here is from pedestrian–pedestrian interactions and is only from a small sample of 14 interactions. Previous work performed this on highly artificial discrete time, turn taking human experiments. This is the first time that a method now exists for more natural continuous data as would be found in real-world AV interactions. The key concept in moving from discrete to continuous data is that we were able to discretise both players actions into just two discrete categories, SLOW and FAST, which enables the sequential chicken model to then operate with minimal changes.

Future work could now make use of this method, firstly to collect and analysis much larger experimental pedestrian–pedestrian data sets; and secondly to deploy a model trailed from this data as a controller in a real AV. It is possible that when trained on larger data sets, the model might show different preferences for different types of pedestrians. For example, real-time detectable features such as age [36], gender [48], body pose [12], activity recognition [14], gait [44], and style of dress might give information about pedestrian intention and behavioural preferences, which if found from training data could then be used to refine real-time AVs pedestrian predictions and active speed controls. This method could then possibly enable new AV online-learning algorithms that adapt to the environment or passenger's preferences.

Acknowlegdements

This project has received funding from the EU H2020 interACT project: Designing cooperative interaction of automated vehicles with other road users in mixed traffic environments under grant agreement No. 723395.

References

1. R. Adkins. Autonomous vehicle intersection management. Technical report, 2016.
2. N. Bellotto, S. Cosar, and Z. Yan. *Human Detection and Tracking*, pages 1–10. Springer Berlin Heidelberg, 2018.
3. N. Bellotto and H. Hu. Computationally efficient solutions for tracking people with a mobile robot: an experimental evaluation of Bayesian filters. *Autonomous Robots*, 28:425–438, 2010.
4. B. Bontrager. The race to fully autonomous cars, 2018. `https://medium.com/swlh/the-race-to-fully-autonomous-cars-8212ff73aad`.
5. R. Brooks. The big problem with self-driving cars is people. In *IEEE Sprectrum*, 2017.
6. C. Cadena, L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I. Reid, and J. Leonard. Past, present, and future of simultaneous localization and mapping: Towards the robust-perception age. *IEEE Transactions on Robotics* , 2016.
7. F. Camara, N. Bellotto, S. Cosar, D. Nathanael, M. Althoff, J. Wu, J. Ruenz, A. Dietrich, and C. W. Fox. Pedestrian models for autonomous driving Part I: low-level models, from sensing to tracking. *IEEE Transactions on Intelligent Transportation Systems*, 2020.
8. F. Camara, N. Bellotto, S. Cosar, F. Weber, D. Nathanael, M. Althoff, J. Wu, J. Ruenz, A. Dietrich, A. Schieben, G. Markkula, F. Tango, N. Merat, and C. W. Fox. Pedestrian models for autonomous driving Part II: high-level models of human behaviour. *IEEE Transactions on Intelligent Transportation Systems*, 2020.
9. F. Camara, S. Cosar, N. Bellotto, N. Merat, and C. W. Fox. Towards pedestrian-AV interaction: method for elucidating pedestrian preferences. In *Proceedings of the IEEE/RSJ Intelligent Robots and Systems (IROS) Workshops*, 2018.
10. F. Camara, N. Merat, and C. W. Fox. A heuristic model for pedestrian intention estimation. In *Proceeedgins of the 22nd IEEE Intelligent Transportation Systems Conference (ITSC)*, 2019.
11. F. Camara, R. Romano, G. Markkula, R. Madigan, N. Merat, and C. W. Fox. Empirical game theory of pedestrian interaction for autonomous vehicles. In *Proceedings of Measuring Behavior 2018: 11th International Conference on Methods and Techniques in Behavioral Research*, 2018.
12. Z. Cao, G. Hidalgo Martinez, T. Simon, S. Wei, and Y. A. Sheikh. Openpose: Realtime multi-person 2d pose estimation using part affinity fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–1, 2019.
13. P. Chen, C. Wu, and S. Zhu. Interaction between vehicles and pedestrians at uncontrolled mid-block crosswalks. *Safety Science*, 82:68 – 76, 2016.
14. C. Coppola, D. R. Faria, U. Nunes, and N. Bellotto. Social activity recognition based on probabilistic merging of skeleton features with proximity priors from RGB-D data. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5055–5061, 2016.
15. D. Dey and J. Terken. Pedestrian interaction with vehicles: Roles of explicit and implicit communication. In *Proceedings of the 9th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*, 2017.

16. R. Elvik. A review of game-theoretic models of road user behaviour. *Accident Analysis & Prevention*, 62:388 – 396, 2014.
17. I. for Advanced Motorists. *How to Be A Better Driver: Advanced Driving the Essential Guide*. Motorbooks, 2007.
18. C. W. Fox, F. Camara, G. Markkula, R. Romano, R. Madigan, and N. Merat. When should the chicken cross the road?: Game theory for autonomous vehicle - human interactions. In *Proceedings of VEHITS: 4th International Conference on Vehicle Technology and Intelligent Transport Systems*, 2018.
19. A. Gorrini, G. Vizzari, and S. Bandini. Towards modelling pedestrian-vehicle interactions: Empirical study on urban unsignalized intersection. In *Proceedings of the 8th International Conference on Pedestrian and Evacuation Dynamics (PED)*, volume abs/1610.07892, 2016.
20. S. Kato, E. Takeuchi, Y. Ishiguro, Y. Ninomiya, K. Takeda, and T. Hamada. An open approach to autonomous vehicles. *IEEE Micro*, 35(6):60–68, 2015.
21. R. Kim, Changwon & Langari. Game theory based autonomous vehicles operation. *International Journal of Vehicle Design (IJVD), Vol. 65*, 2014.
22. T. J. Kim. Automated autonomous vehicles: Prospects and impacts on society. *Journal of Transportation Technologies*, 2018.
23. Y. Li, M. Li, L. Dou, Q. Zhao, Z. Wang, and J. Li. A socially multi-robot target tracking method based on extended cooperative game theory. In *Proceedings of the IEEE International Conference on Robotics and Biomimetics (ROBIO)*, pages 247–252, 2014.
24. W. Ma, D. Huang, N. Lee, and K. M. Kitani. Forecasting interactive dynamics of pedestrians with fictitious play. In *Proc. of IEEE CVPR*, 2017.
25. R. Madigan, T. Louw, M. Dziennus, T. Graindorge, E. Ortega, M. Graindorge, and N. Merat. Acceptance of automated road transport systems (ARTS): An adaptation of the utaut model. *Transportation Research Procedia*, 14:2217 – 2226, 2016.
26. F. Meng, J. Su, C. Liu, and W. H. Chen. Dynamic decision making in lane change: Game theory with receding horizon. In *Proceedings of UKACC 11th International Conference on Control (CONTROL)*, pages 1–6, 2016.
27. A. Millard-Ball. Pedestrians, autonomous vehicles, and cities. *Journal of Planning Education and Research*, 38(1):6–12, 2018.
28. O. Morgenstern and J. Von Neumann. *Theory of games and economic behavior*. Princeton University Press, 1953.
29. E. Papadimitriou, S. Lassarre, and G. Yannis. Pedestrian risk taking while road crossing: A comparison of observed and declared behaviour. *Transportation Research Procedia*, 14:4354 – 4363, 2016.
30. H. A. Rakha, I. Zohdy, and R. K. Kamalanathsharma. Agent-based game theory modeling for driverless vehicles at intersections. Technical report, 2013.
31. A. Rane, S. Krishnan, and S. Waman. Conflict resolution of autonomous cars using game theory and cellular automata. In *Proceedings of the International Conference on Reliability Optimization and Information Technology (ICROIT)*, pages 326–330, 2014.
32. A. Rasouli, I. Kotseruba, and J. K. Tsotsos. Agreeing to cross: How drivers and pedestrians communicate. *Proceedings of the IEEE Intelligent Vehicles Symposium (IV)*, pages 264–269, 2017.
33. M. Risto, C. Emmenegger, E. Vinkhuyzen, M. Cefkin, and J. Hollan. Human-vehicle interfaces: The power of vehicle movement gestures in human road user coordination. In *Proceedings of the 9th International Driving Symposium on Human Factors in Driver Assessment*, pages 186–192, 2017.
34. D. Rothenbücher, J. Li, D. Sirkin, B. Mok, and W. Ju. Ghost driver: A field study investigating the interaction between pedestrians and driverless vehicles. In *Proceedings of the 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 795–802, 2016.

35. F. Schneemann and I. Gohl. Analyzing driver-pedestrian interaction at crosswalks: A contribution to autonomous driving in urban environments. In *Proceedings of the IEEE Intelligent Vehicles Symposium (IV)*, pages 38–43, 2016.
36. S. Sithungu and D. Van der Haar. Real-time age detection using a convolutional neural network. In *Business Information Systems*, pages 245–256, 2019.
37. K. Skrzypczyk. Game theory based target following by a team of robots. In *Proceedings of the 4th IEEE International Workshop on Robot Motion and Control*, pages 91–96, 2004.
38. S. Thrun, W. Burgard, and D. Fox. *Probabilistic Robotics*. Intelligent robotics and autonomous agents. MIT Press, 2005.
39. P. Trautman and A. Krause. Unfreezing the robot: Navigation in dense, interacting crowds. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 797–803, 2010.
40. A. Turnwald, D. Althoff, D. Wollherr, and M. Buss. Understanding human avoidance behavior: Interaction-aware decision making based on game theory. *International Journal of Social Robotics*, 8(2):331–351, 2016.
41. A. Turnwald, W. Olszowy, D. Wollherr, and M. Buss. Interactive navigation of humans from a game theoretic perspective. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 703–708, 2014.
42. Udacity. Open source self-driving car, 2017.
43. M. Wade. Silicon valley is winning the race to build the first driverless cars, 2018. `https://theconversation.com/silicon-valley-is-winning-the-race-to-build-the-first-driverless-cars-91949`.
44. L. Wang, T. Tan, H. Ning, and W. Hu. Silhouette analysis-based gait recognition for human identification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(12):1505–1518, 2003.
45. C. K. Williams and C. E. Rasmussen. Gaussian processes for regression. In *Advances in Neural Information Processing Systems*, pages 514–520, 1996.
46. M. Yan. Multi-robot searching using game-theory based approach. *International Journal of Advanced Robotic Systems*, 5(4):44, 2008.
47. Z. Yan, T. Duckett, and N. Bellotto. Online learning for human classification in 3d lidar-based tracking. In *Proceedings of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, pages 864–871, 2017.
48. W. Zhang, M. L. Smith, L. N. Smith, and A. Farooq. Gender recognition from facial images: two or three dimensions? *JOSA A*, 33(3):333–344, 2016.