



Deposited via The University of Leeds.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/150884/>

Version: Accepted Version

Proceedings Paper:

Camara, F, Merat, N and Fox, CW (2019) A heuristic model for pedestrian intention estimation. In: Proceedings of 2019 IEEE Intelligent Transportation Systems Conference (ITSC). 2019 IEEE Intelligent Transportation Systems Conference (ITSC), 27-30 Oct 2019, Auckland, New Zealand. IEEE, pp. 3708-3713. ISBN: 978-1-5386-7024-8.

<https://doi.org/10.1109/ITSC.2019.8917195>

© 2019 IEEE. This is an author produced version of a paper published in Proceedings of 2019 IEEE Intelligent Transportation Systems Conference (ITSC). Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works. Uploaded in accordance with the publisher's self-archiving policy.

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.

A heuristic model for pedestrian intention estimation

Fanta Camara^{1,2}, Natasha Merat¹ and Charles W. Fox^{1,2,3}

Abstract—Understanding pedestrian behaviour and controlling interactions with pedestrians is of critical importance for autonomous vehicles, but remains a complex and challenging problem. This study infers pedestrian intent during possible road-crossing interactions, to assist autonomous vehicle decisions to yield or not yield when approaching them, and tests a simple heuristic model of intent on pedestrian-vehicle trajectory data for the first time. It relies on a heuristic approach based on the observed positions of the agents over time. The method can predict pedestrian crossing intent, crossing or stopping, with 96% accuracy by the time the pedestrian reaches the curbside, on the standard Daimler pedestrian dataset. This result is important in demarcating scenarios which have a clear winner and can be predicted easily with the simple heuristic, from those which may require more complex game-theoretic models to predict and control.

Index Terms—Pedestrian Intention Crossing Estimation; Agent-Human Interactions; Autonomous Vehicles

I. INTRODUCTION

Fully autonomous vehicles (AVs) promise a future with better mobility systems, fewer on-road accidents, and reduced congestion in cities [15]. While their localisation, mapping and route-planning are well understood problems [22], [14], there remain major concerns about their *interaction* with other road users, especially more vulnerable ones such as pedestrians and cyclists. Other road users are less predictable than static environments, being active agents making their own complex intelligent decisions, and are sometimes in direct competition rather than co-operation with the vehicle for priority on the road. Brooks has identified other road users as ‘the big problem with self-driving cars’ [3].



Fig. 1. Pedestrian intention with a vehicle, from its dashcam, in the Daimler dataset [17]. The present study uses a heuristic model to predict the outcomes of interactions from this data – whether the pedestrian will cross in front of the vehicle or yield to it.

This project 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.

¹ Institute for Transport Studies (ITS), University of Leeds, UK

² School of Computer Science, University of Lincoln, UK

³ Ibex Automation Ltd., UK

Pedestrians are the most challenging road users for AVs. Unlike cyclists – who must follow clearly defined traffic rules – pedestrian behaviour is less predictable and more complex to model. Predicting pedestrians’ behaviour is hard due to the multiple uncertainties about their posture, gestures, and demographics, and the predictive links between these visible features and their underlying intentions and emotions. A recent UK government report has emphasised the need for AV-pedestrian interaction research and its incompleteness as a research area, having possible implications that could require changes in traffic legislation [7]. These arise from the problem of automating the concept of assertiveness in competition for road space. Usually engineering systems are designed to be as safe as possible, but vehicles driven completely safety must always yield to other road users in competition with them, and will hence make zero progress in busy areas, known as the ‘freezing robot problem’ [23]. Like human-driven vehicles, AVs may thus sometimes need to maintain a credible threat of colliding with or otherwise causing some form of lesser but real harm to pedestrians in order to sometimes obtain priority over them and make progress. The law will require this notion of credible threats of deliberate harm to be rationalized and quantified so that AVs using it can demonstrate – in court in the event of accidents – that they have acted optimally and legally.

A game theoretic model, ‘Sequential Chicken’, was recently proposed to provide optimal strategies for such scenarios where an autonomous vehicle is encountering a pedestrian intending to cross the road at a crossing where priority is not clearly defined [10]. In this model, two agents X and Y are moving towards an intersection. They communicate their crossing intent through their positions, i.e. their distances x and y from the intersection point. If neither player yields, they collide and receive a negative utility U_{crash} . Otherwise they get a utility depending on the time taken to reach their destination, U_{time} . [6] and [4] then presented a set of laboratory experiments with human participants playing artificial version of the Sequential Chicken game, and used them to find a unique behavioural parameter $\theta = U_{crash}/U_{time}$ for the players, summarising their balance of preferences for avoiding collision versus saving travel time.

In cases where there is a clear winner visible at the start of an interaction – such as when one player is clearly much closer to reaching the intersection point than the other – this game theory model reduces to a simple deterministic prediction that this player will win. But in other cases, where both players have a similar predicted time to arrival at the intersections, the model does not reduce to a simple prediction and instead models complex game theoretic interactions as the two players signal to each other and negotiate for priority with their physical positions.

The previous laboratory studies were in artificial environments so it is not clear from them alone how often the full game theory model vs the reduced predictions are needed in real world interactions. We would like to know how often a simple heuristic matched to the reduced model will predict accurately. This will tell us whether and when the full game theory model is needed and appropriate in the real world. To this end, a standard dataset from pedestrian-vehicle interactions [17] is analysed with such a simple heuristic to test for its accuracy.

A. Other related work

Intention estimation plays an important role in interactions scenarios, and predicting the behaviour of road users and controlling robotic interactions with them now forms an active research area. [25] proposed a model that learns a set of features from a database of LIDAR pedestrian trajectories using support vector machines (SVM) and predicts whether or not a pedestrian will cross the street. In [24] they compared this with LSTM neural networks to predict pedestrian crossing intent and in [26], an extended version of random forests – quantile regression forests – is used to predict pedestrian road-crossing behaviour. [8] used the Theory of Planned Behaviour (TPB) to predict adolescents road crossing intentions. [1] developed an intention-aware motion planning system based on Mixed Observable Markov Decision Processes (MOMDPs) where pedestrian intention is incorporated to the planning model as the unobserved variables, implemented on an autonomous robot navigating in a university campus and dealing with jaywalkers. [11] developed a Gaussian Process model that incorporates contextual features such as the distance to the curbside and to the traffic light. Pedestrian trajectories are predicted using a Transferable ANSC algorithm. [2] defined a set of features that models an inner-city and proposed a generic context-based model to predict pedestrians behaviour at zebra crossings. [17] developed a dynamic Bayesian network including pedestrian head orientation to the trajectory and crossing intent prediction. [21] used a latent-dynamic conditional random field model including pedestrian dynamics and awareness of the environment such as their head orientation for intention recognition. [16] proposed a motion contour histogram of oriented gradients descriptor method to predict pedestrian crossing intent with a high accuracy, tested on data collected from laboratory conditions. [18] used Fictitious Play to compute equilibria over similar 2D trajectories, for pedestrian-pedestrian interactions in pedestrian-only environments. [27] proposed a probabilistic method based on Markov chains to predict pedestrian motion in urban environments. [13] proposed a three-layer trajectory prediction model composed of a trajectory planner, a force-based (social force) model and a game theoretic decision model. A deep learning model was proposed in [12] to predict pedestrian crossing intention, which a good accuracy on another Daimler dataset. [20] presented a pedestrian crossing intention detection using contextual features and a SVM classifier. [9] developed a CNN model based for 2D human pose estimation combined with random forests and a radial basis function support vector machine for pedestrian intention estimation.

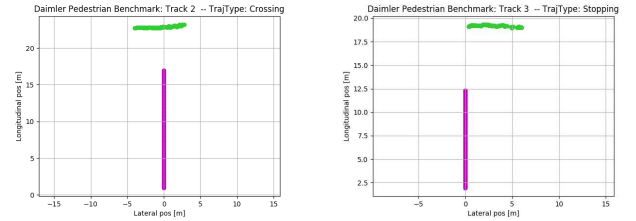
II. METHODS

The objective of the present study is to show where game theory is not required for pedestrian-vehicle interactions. We show how accurately and with how much available time before a potential collision, it is possible for the proposed heuristic model, that we name *heuristic ratio model*, to predict pedestrian intent to cross or not cross in front of a vehicle, in trajectory data of vehicle-pedestrian interactions.

A. Data

We used the well-known Daimler pedestrian trajectory prediction dataset [17] composed of a set of data from 58 pedestrian-vehicle interactions. The dataset includes vehicle telematics together with dashcam video of vehicle-pedestrians interactions between a vehicle driving around and pedestrians that it encountered¹. In this dataset, all the pedestrians enter the road from the right hand side of the upcoming vehicle. Pedestrians are then detected and tracked in the images to produce trajectory information for pedestrians [17]. Available data thus includes: ground truth and measured positions of the pedestrian and the vehicle during an interaction; outcome (whether the pedestrian ultimately crosses or yields); whether the

¹Video sequences last several seconds (min /max / mean: 2.53 s / 13.27 s / 7.15 s) [17]



(a) Pedestrian crossing, vehicle yielding (b) Pedestrian yielding, vehicle passing

Fig. 2. Examples of interactions from the Daimler dataset. (green=pedestrian trajectory; magenta=car trajectory).

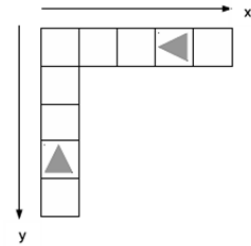


Fig. 3. Sequential Chicken game model

pedestrian sees the vehicle; and pedestrian head orientation during the interaction. Two interactions are shown in Fig. 2 with the trajectories of the pedestrian in *green* and the vehicle in *magenta*. These trajectories² were used to test our method and for each interaction, we derived from the dataset the following inputs:

- *pedestrian trajectory*, as the pedestrian distance to the curbside is used as giving the pedestrian position over time
- *vehicle trajectory*, as the vehicle speed is used to derived the vehicle positions over time
- *distance between vehicle and pedestrian*, as the pedestrian longitudinal distance to the ego vehicle is used to measure the distance between the pedestrian and the upcoming vehicle.

B. Heuristic Ratio Model

1) *Parametrization*: Each interaction between the vehicle and the pedestrian is described as a sequence of discrete games where the goal is to cross first the intersection while avoiding a collision. From the first detection of the pedestrian, time is discretized into integer ticks, t . This heuristic model is inspired from the Sequential Chicken model [10], thus space is discretized into *cells* as shown in Fig. 3, with the pedestrian being player X and the vehicle being player Y, and their locations discretized as integer cells x and y as they approach each other at right angles.

In the original Sequential Chicken model [10], at each time t , the vehicle and pedestrian players both choose simultaneously between two discrete actions, to move either SLOW (one square forwards) or FAST (two squares forwards), and must negotiate via these moves to avoid collision where their paths overlap. This original model cannot fit data where a player stops completely, so here we extend it to allow pedestrians an extra third action, STOP.

At each t , we wish to know the probability $P(X) = 1 - P(Y)$ that the pedestrian will be the eventual winner of the game by crossing the road before the vehicle passes.

²Three interactions in the dataset contain trajectories where the pedestrian is already crossing the road or at the edge of the curbside, we thus removed these three interactions and tested our model on the remaining 55 interactions.

TABLE I
MAPPING BETWEEN DISCRETE AND REAL SPEEDS FOR PEDESTRIAN
AND VEHICLE

Player	SLOW action	FAST action
Pedestrian	1. $\frac{X_{cellsize}}{\Delta t} \approx 0.75m/s$	2. $\frac{X_{cellsize}}{\Delta t} \approx 1.5m/s$
Vehicle	1. $\frac{Y_{cellsize}}{\Delta t} \approx 13.4km/h$	2. $\frac{Y_{cellsize}}{\Delta t} \approx 26.9km/h$

The discrete cell scales are chosen to fit the scales of typical interactions in the Daimler data: pedestrian cells are chosen to be of size $X_{cellsize} = 0.045m$ square and vehicle cells $Y_{cellsize} = 0.225m$ square, so that the pedestrian and vehicle speeds are,

$$speedX = N \frac{X_{cellsize}}{\Delta t}, \quad (1)$$

where N is the number of discrete cells they move in the game grid as in [6] and $\Delta t = 0.0603s$ is defined as one video frame duration in the Daimler data. This gives the speeds in Table I: for the pedestrian, x_{speed} is about 0.75m/s in a slow move and higher than 1.5m/s in normal/fast move while for the vehicle y_{speed} is about 13.4km/h and higher than 26.9km/h in normal/fast motion. These cell sizes and speeds are coherent with the normal walking speed of pedestrians, about 1.5m/s [19], and the driving speed of the vehicle in the data recording environment [17]. When pedestrian walking speed is lower than their cell size $X_{cellsize}$ be consider it to be a STOP action. Using these cell sizes, we then create the discrete game grid as in Fig. 3, with the two players initial locations x and y set to cells corresponding to their continuous locations in the video at the start of the interaction.

2) *Algorithm*: Given the proposed discretization of the space, we obtained X_{cell_number} which represents the number of discrete cells for player X and Y_{cell_number} the number of cells for player Y. The game starts with an observation of the position X_t of the pedestrian and the speed V_{y_t} of the vehicle. If the pedestrian's speed is higher than the normal walking speed of 1.5m/s then the selected action is to *go fast*, a decrease of the speed but kept within the interval $[X_{cellsize}, 2.X_{cellsize}]$ is regarded as *yield* and if the speed is even smaller than $X_{cellsize}$ then it is a *stop* action. For the vehicle, only the speed variation informs about the action selection between *go fast* and *yield*. Once each player's action is selected, they move on the grid. For each new observation, the likelihoods of winning $\lambda_{win}X$ and $\lambda_{win}Y$ of the agents are computed given their current position. The likelihood of winning is calculated based on the remaining cells that separate the players from the intersection/collision, which is defined as a *ratio* of their distances to the intersection. These likelihoods are normalized as:

$$normalize(a) = a/Z \quad (2)$$

where Z is a normalizer. We then *fuse* the likelihood of winning with the prior π to form the probabilities of winning $P(X | X_t)$ and $P(Y | Y_t)$ where the Bayesian fusion operator is:

$$p \otimes q = \frac{pq}{pq + (1-p) + (1-q)} \quad (3)$$

In the present study, the prior π is chosen to be flat 0.5, i.e both pedestrian and vehicle have equal probability to cross the road first. These probabilities of winning are finally stored into vectors S_X and S_Y that inform about the winner of the interaction over time with:

$$S_X = \{P(X | x_0), \dots, P(X | x_n)\}, \quad S_Y = \{P(Y | y_0), \dots, P(Y | y_n)\} \quad (4)$$

Algorithm 1 Crossing Intent Probability Computation

```

1:  $P(X | X_0) \leftarrow \pi$ 
2:  $P(Y | Y_0) \leftarrow 1 - \pi$ 
    $X_{cell\_number} \leftarrow \frac{X_{distance2curbside}}{X_{cellsize}}$ 
3:  $Y_{cell\_number} \leftarrow \frac{Y_{distance2X}}{Y_{cellsize}}$ 
    $x_{index} \leftarrow X_{cell\_number}$ 
4:  $y_{index} \leftarrow Y_{cell\_number}$ 
5: for each new observation  $X_t$  and  $V_{y_t}$ : do
6:   if  $X_t > 0$  then
7:      $x_{speed} \leftarrow X_t - X_{t-1}$ 
8:     if  $x_{speed} \geq 2.X_{cellsize}$  then
9:        $x_{action} \leftarrow FAST$ 
10:    else
11:      if  $x_{speed} > X_{cellsize}$  and  $x_{speed} < 2.X_{cellsize}$  then
12:         $x_{action} \leftarrow SLOW$ 
13:      else
14:         $x_{action} \leftarrow STOP$ 
15:      end if
16:    end if
17:   if  $V_{y_t} - V_{y_{t-1}} > 0$  then
18:      $y_{action} \leftarrow FAST$ 
19:   else
20:      $y_{action} \leftarrow SLOW$ 
21:   end if
22:    $x_{index} \leftarrow x_{index} - x_{action}$ 
23:    $y_{index} \leftarrow y_{index} - y_{action}$ 
24:    $\lambda_{win}X = 1 - \frac{x_{index}}{\max(X_{cell\_number}, Y_{cell\_number})}$ 
25:    $\lambda_{win}Y = 1 - \frac{y_{index}}{\max(X_{cell\_number}, Y_{cell\_number})}$ 
26:    $normalize(\lambda_{win}X, \lambda_{win}Y)$ 
27:    $P(X_{win} | X_t) \leftarrow \pi \otimes \lambda_{win}X$ 
28:    $P(Y_{win} | Y_t) \leftarrow \pi \otimes \lambda_{win}Y$ 
29:    $S_X \leftarrow store(P(X_{win} | X_t))$ 
30:    $S_Y \leftarrow store(P(Y_{win} | Y_t))$ 
31:    $X_{t-1} \leftarrow X_t$ 
32:    $V_{y_{t-1}} \leftarrow V_{y_t}$ 
33:   end if
34: end for
35: Return  $S_X, S_Y$ 

```

C. Variants of the heuristic ratio model

1) *Ratio model (Model A)*: We define Model A as being the heuristic ratio model described above and detailed in Algorithm 1. The pedestrian is the winner if their probability of winning is greater than the vehicle probability of winning at the time pedestrian reaches the curbside as:

$$P(X_{win} | T_{Xcurbside}) > P(Y_{win} | T_{Xcurbside}) \quad (5)$$

2) *Ratio model with scaling factor (Model B)*: This model relies on the heuristic ratio model (Model A) and we define a scaling factor $\alpha = 2.15$ and multiply it with the vehicle likelihood of winning $\lambda_{win}Y$ before its normalization in Algorithm 1, such that:

$$\begin{cases} \lambda_{win}Y = \alpha * \lambda_{win}Y \\ P(X_{win} | T_{Xcurbside}) > P(Y_{win} | T_{Xcurbside}) \end{cases} \quad (6)$$

3) *Evaluation*: We used the mean absolute error (MAE) to gain insight about the accuracy of the predictions given by,

$$MAE = \langle |p - g| \rangle \quad (7)$$

where p is the predicted outcome and g is the ground truth value. MAE informs about how confident we can be with the predictions.

III. RESULTS

The results of the heuristic ratio model are shown in Table II. The ratio model (Model A) has a prediction accuracy of 76.4% while the ratio model with a scaling factor (Model B) can predict pedestrian crossing with 96.3% accuracy by the time the pedestrian reaches the curbside. In Fig. 4 and 5, we show the evolution of Algorithm 1 for two interaction scenarios using the best model (Model B). The estimated trajectories in meters and discrete cells show how the pedestrian and vehicle move over time in the discrete game grid according to the heuristic ratio model.

TABLE II
RESULTS FOR THE RATIO MODEL VARIANTS

Type of Model	Starting Frame	Prediction	Accuracy
Model A: Ratio Model	Frame 0	44/55	80.0%
Model A: Ratio Model	Mixed starting frames	42/55	76.4%
Model B: Ratio Model with Scaling factor	Frame 0	40/55	72.7%
Model B: Ratio Model with Scaling factor	Mixed starting frames	53/55	96.3%

		Prediction outcome		Total
		Crossing	Stopping	
Actual value	Crossing	40	1	41
	Stopping	1	13	14
Total		41	14	

TABLE III
CONFUSION MATRIX

In total, there were 41 crossing scenarios and 14 stopping scenarios. As shown in Table III, our model correctly predicts 40 crossing scenarios out of 41 and 13 stopping scenarios out of 14, hence reaching an accuracy of 96.3%. The two probabilities for the interactions that were not correctly classified are shown in Fig. 6, it is to be noted that these interactions were categorized as either ‘critical’ or ‘anomalous’ situations in the original dataset [17]. Given the good results found with Model B, we can interpret that this is due to the scaling factor α , which is compensating the lack of good prior used in our model. Due to the small amount data available, we have used the Daimler dataset as the test set only but in the future, the prior could be learnt from training data.

Fig. 7(a) shows the mean average error of the pedestrian’s probability of winning from the start of the interaction using Model B predictions. Pedestrian’s behaviour is uncertain at the beginning of the interaction and becomes more and more certain. This tells us that autonomous vehicles should wait and observe a certain number of features before acting. This is coherent with the previous work in [5]. Fig. 7(b) represents the uncertainty over pedestrian’s

probability of winning from the end of the interactions using Model B predictions. It shows that Model B can predict pedestrian crossing intention with about 90% confidence around 1s before the end of the interaction and it reaches 96% confidence in the intention estimation at the time the pedestrian reaches the curbside.

IV. CONCLUSIONS

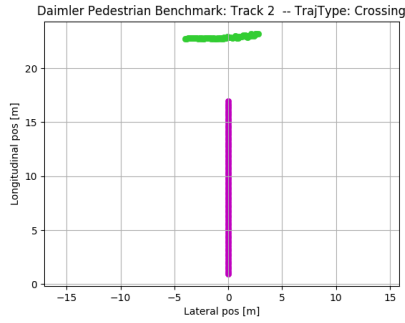
A simple heuristic model can accurately estimate pedestrian crossing intent on the standard Daimler dataset with about 96% accuracy and about one second before the crossing action occurs. The remaining one in every twenty interactions, which are those in which both players have similar initial times to arrival at the intersection, may thus require more complex models such as the Sequential Chicken game theory model. This shows that such complex models do appear to be necessary for AVs in general use and should continue to be developed and refined to bring accuracy closer to 100%.

V. ACKNOWLEDGEMENTS

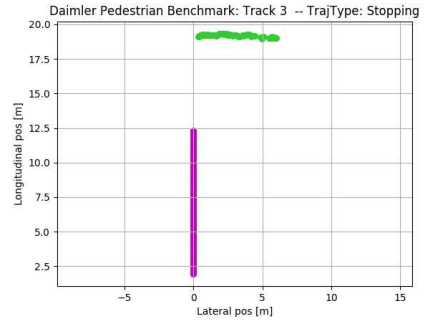
We are grateful to Jingyuan Wu and Johannes Ruenz at Robert Bosch GmbH for useful discussions and suggestions.

REFERENCES

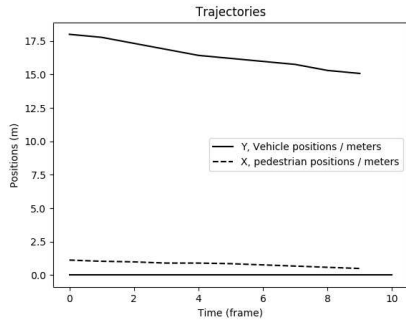
- [1] T. Bandyopadhyay, C. Z. Jie, D. Hsu, M. H. Ang, D. Rus, and E. Frazzoli. Intention-aware pedestrian avoidance. In *Experimental Robotics*, pages 963–977. Springer, 2013.
- [2] S. Bonnin, T. H. Weisswange, F. Kummert, and J. Schmuuederich. Pedestrian crossing prediction using multiple context-based models. In *17th International IEEE Conference on Intelligent Transportation Systems (ITSC)*, pages 378–385, Oct 2014.
- [3] R. Brooks. The big problem with self-driving cars is people and we’ll go out of our way to make the problem worse, 2017.
- [4] F. Camara, S. Cosar, N. Bellotto, N. Merat, and C. W. Fox. Towards pedestrian-av interaction: method for elucidating pedestrian preferences. In *IEEE/RSJ Intelligent Robots and Systems (IROS) Workshops*, 2018.
- [5] F. Camara, O. Giles, R. Madigan, M. Rothmüller, P. Holm Rasmussen, S. A. Vendelbo-Larsen, G. Markkula, Y. M. Lee, L. Garach, N. Merat, and C. W. Fox. Filtration analysis of pedestrian-vehicle interactions for autonomous vehicles control. In *Proceedings of the 15th International Conference on Intelligent Autonomous Systems workshops*, 2018.
- [6] F. Camara, R. Romano, G. Markkula, R. Madigan, N. Merat, and C. W. Fox. Empirical game theory of pedestrian interaction for autonomous vehicles. In *Measuring Behavior 2018: 11th International Conference on Methods and Techniques in Behavioral Research*. Manchester Metropolitan University, March 2018.
- [7] UK Law Commission Consultation. Automated vehicles a joint preliminary consultation paper, 2018.
- [8] D. Evans and P. Norman. Predicting adolescent pedestrians’ road-crossing intentions: an application and extension of the theory of planned behaviour. *Health Education Research*, 18(3):267–277, 2003.
- [9] Z. Fang, D. Vázquez, and A. López. On-board detection of pedestrian intentions. *Sensors*, 17(10):2193, 2017.
- [10] 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 *VEHITS 2018: 4th International Conference on Vehicle Technology and Intelligent Transport Systems*, January 2018.
- [11] G. Habibi, N. Jaipuria, and J. P. How. Context-aware pedestrian motion prediction in urban intersections. In *IEEE 21st International Conference on Intelligent Transportation Systems*, 2018.
- [12] M. Hoy, Z. Tu, K. Dang, and J. Dauwels. Learning to predict pedestrian intention via variational tracking networks. In *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pages 3132–3137, 2018.
- [13] F. T. Johora and J. Müller. Modeling interactions of multimodal road users in shared spaces. In *21st International Conference on Intelligent Transportation and Systems (ITSC)*, 12 2018.
- [14] 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, Nov 2015.



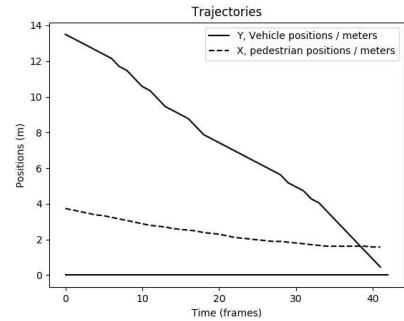
(a) Interaction 2



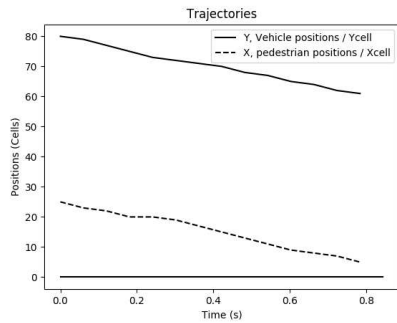
(a) Interaction 3



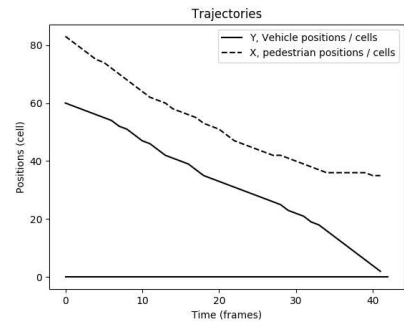
(b) Estimated trajectories in meters (from 16th frame)



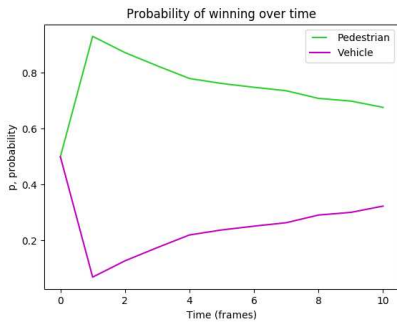
(b) Estimated trajectories in meters (from 16th frame)



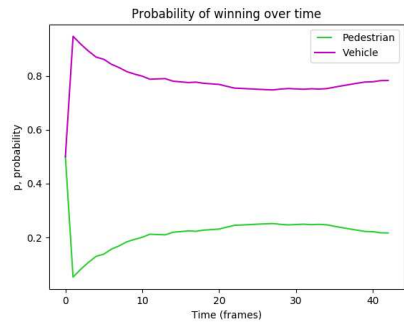
(c) Estimated trajectories in discrete cells (from 16th frame)



(c) Estimated trajectories in discrete cells (from 16th frame)



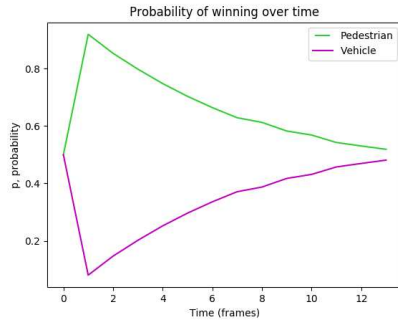
(d) $P(\text{player} \mid \tau_i)_{\tau_i=t:T}$



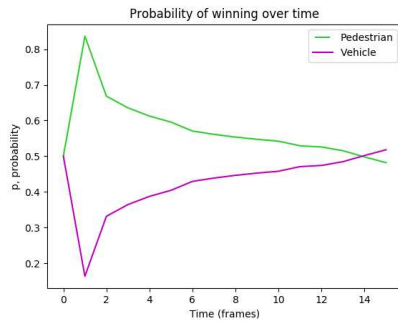
(d) $P(\text{player} \mid \tau_i)_{\tau_i=t:T}$

Fig. 4. Results obtained from Model B for pedestrian crossing scenario

Fig. 5. Results obtained from Model B for pedestrian stopping scenario

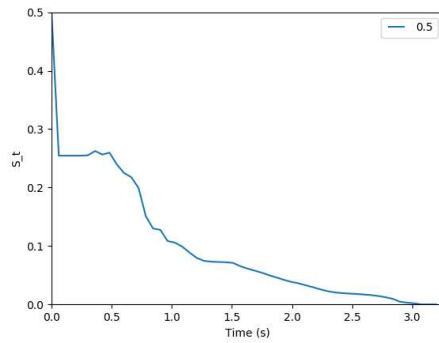


(a) $P(\text{player} \mid \tau_i)_{\tau_i=t:T}$ for a crossing scenario predicted as pedestrian stopping

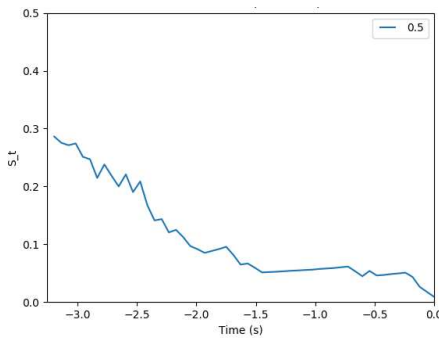


(b) $P(\text{player} \mid \tau_i)_{\tau_i=t:T}$ for a stopping scenario predicted as pedestrian crossing

Fig. 6. Incorrect predictions



(a) Time from the start (s)



(b) Time from the end (s)

Fig. 7. Mean average error for Model B

- [15] T. J. Kim. Automated autonomous vehicles: Prospects and impacts on society. *Journal of Transportation Technologies*, 2018.
- [16] S. Köhler, M. Goldhammer, S. Bauer, K. Doll, U. Brunsmann, and K. Dietmayer. Early detection of the pedestrian's intention to cross the street. In *2012 15th International IEEE Conference on Intelligent Transportation Systems*, pages 1759–1764. IEEE, 2012.
- [17] J. F. P. Kooij, N. Schneider, F. Flohr, and D. M. Gavrila. Context-based pedestrian path prediction. In *Computer Vision – ECCV 2014*, pages 618–633, 2014.
- [18] W. Ma, D. Huang, N. Lee, and K. M. Kitani. Forecasting interactive dynamics of pedestrians with fictitious play. *CoRR*, abs/1604.01431, 2016.
- [19] H. C. Manual. Highway capacity manual. *Washington, DC*, 2, 2000.
- [20] F. Schneemann and P. Heinemann. Context-based detection of pedestrian crossing intention for autonomous driving in urban environments. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2243–2248, Oct 2016.
- [21] A. T. Schulz and R. Stiefelhagen. Pedestrian intention recognition using latent-dynamic conditional random fields. In *2015 IEEE Intelligent Vehicles Symposium (IV)*, pages 622–627, June 2015.
- [22] S. Thrun, W. Burgard, and D. Fox. *Probabilistic Robotics*. Intelligent robotics and autonomous agents. MIT Press, 2005.
- [23] P. Trautman and A. Krause. Unfreezing the robot: Navigation in dense, interacting crowds. In *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 797–803, Oct 2010.
- [24] B. Völz, K. Behrendt, H. Mielenz, I. Gilitschenski, R. Siegwart, and J. Nieto. A data-driven approach for pedestrian intention estimation. In *Intelligent Transportation Systems (ITSC), 2016 IEEE 19th International Conference on*, pages 2607–2612. IEEE, 2016.
- [25] B. Völz, H. Mielenz, G. Agamennoni, and R. Siegwart. Feature relevance estimation for learning pedestrian behavior at crosswalks. In *2015 IEEE 18th International Conference on Intelligent Transportation Systems (ITSC)*, pages 854 – 860, 2015.
- [26] B. Völz, H. Mielenz, R. Siegwart, and J. Nieto. Predicting pedestrian crossing using quantile regression forests. In *2016 IEEE Intelligent Vehicles Symposium (IV)*, pages 426–432, June 2016.
- [27] J. Wu, J. Ruenz, and M. Althoff. Probabilistic map-based pedestrian motion prediction taking traffic participants into consideration. In *IEEE Intelligent Vehicles Symposium (IV) June 26-30, 2018, Changshu, Suzhou, China*, 2018.