



UNIVERSITY OF LEEDS

This is a repository copy of *Filtration analysis of pedestrian-vehicle interactions for autonomous vehicle control*.

White Rose Research Online URL for this paper:
<http://eprints.whiterose.ac.uk/133025/>

Version: Accepted Version

Proceedings Paper:

Camara, F, Giles, O, Madigan, R et al. (8 more authors) (2018) Filtration analysis of pedestrian-vehicle interactions for autonomous vehicle control. In: Proceedings of IAS-15. 15th International Conference on Autonomous Systems, 11-15 Jun 2018, Baden-Baden, Germany. interACT .

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.



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

Filtration analysis of pedestrian-vehicle interactions for autonomous vehicle control

Fanta Camara^{1,2}, Oscar Giles¹, Ruth Madigan¹, Markus Rothmüller³,
Pernille Holm Rasmussen³, Signe Alexandra Vendelbo-Larsen³,
Gustav Markkula¹, Yee Mun Lee¹, Laura Garach-Morcillo^{1,4},
Natasha Merat¹ and Charles Fox^{1,2,5}

¹ Institute for Transport Studies, University of Leeds, UK

² School of Computer Science, University of Lincoln, UK

³ School of Architecture, Design and Planning, Aalborg University, DNK

⁴ Department of Civil Engineering, University of Granada, ESP

⁵ Ibox Automation Ltd., UK

Abstract. Interacting with humans remains a challenge for autonomous vehicles (AVs). When a pedestrian wishes to cross the road in front of the vehicle at an unmarked crossing, the pedestrian and AV must compete for the space, which may be considered as a game-theoretic interaction in which one agent must yield to the other. To inform development of new real-time AV controllers in this setting, this study collects and analyses detailed, manually-annotated, temporal data from real-world human road crossings as they interact with manual drive vehicles. It studies the temporal orderings (filtrations) in which features are revealed to the vehicle and their informativeness over time. It presents a new framework suggesting how optimal stopping controllers may then use such data to enable an AV to decide when to act (by speeding up, slowing down, or otherwise signalling intent to the pedestrian) or alternatively, to continue at its current speed in order to gather additional information from new features, including signals from that pedestrian, before acting itself.⁶

1 Introduction

While localisation, mapping, route planning, and control are now largely solved problems for autonomous vehicles in static and ballistic environments [8], the major outstanding challenge for real-world autonomous vehicles is operation in environments containing people. Unlike static and ballistic environments, people are complex interactive agents having their own goals, utilities, and decision making systems, and interactions with them must take these into account in order to predict their actions and plan accordingly. Interaction is recursive and complex: an AV's own actions will affect the person's actions and vice versa. This is critical in environments where traffic rules do not clearly define priority, such as at unmarked intersections, where AVs and pedestrians have to negotiate over who will pass the other.

⁶ 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.

A game-theoretic model of such interactions was recently presented [4] proving that (under several assumptions including discretisable space and time, no lateral motion, and communication only via agent positioning on the road) the optimal strategy for both agents is probabilistic and recursive. As the two agents get closer over time, both should gradually increase the probability that they will yield at each time, then draw their yield or non-yield action from this probability. The probabilities both tend to unity as the agents get closer to a collision. The model proves that there must remain some small but strictly non-zero probability of the crash actually occurring, in both agents strategies, in order for the interaction to proceed optimally. (This formalises the intuition that if an AV is known to be perfectly safe, then it will make no progress as all other road users may push in front of it at every interaction.) A second study [1], then empirically measured human behaviours in a laboratory version of a road crossing scenario, and suggested it is possible to assign a single parameter to each agent which summarizes their entire behavioral preferences during such interactions. This parameter measures ‘assertiveness’ as $\theta = U_{time}/U_{crash}$, the ratio of the agents value of time (i.e. dollar value of losing 1 second of arriving at their destination, for example by yielding to the other agent for road priority), and the agents (negative) value of the collision actually occurring (which will be worse for an unarmoured pedestrian than for the driver of a heavy protective car, especially of a larger car such as an SUV).



Fig. 1: Intersection where pedestrian-vehicle road-crossing interactions were observed, by observers at locations X and Y. (WGS84: 53.8073, -1.5518)

Real-world AV controllers based on this game theoretic model would thus benefit from any additional information about θ for pedestrians which they encounter in such situations. θ is not directly observable but the previous study recommended further work to discover observable features which help to infer it.

To this end, the present study proposes a new temporal filtration-based framework for analyzing pedestrian-vehicle interactions during road crossings. It uses data collected manually from real-world road crossings interactions between pedestrians and human drivers in an urban environment, to study how information about the winner is revealed over time via a set of manually defined and collected features. For example, an AV encountering a pedestrian trying to cross may initially see information about the road geometry, then the demographic of the pedestrian, then movements made by the pedestrian. During this period, the AV may choose to act (e.g. changing its speed

or otherwise signalling to the pedestrian) or to not act and continue with its current ballistics for some period in order to collect more information. As a (literal!) Optimal Stopping problem, the framework shows the trade-off between time and information: if the AV waits too long then it will pass or hit the pedestrian before any decision is made; but if it acts too soon then it risks missing valuable information about the pedestrian which would improve the action selection.

1.1 Related work

This is a work-in progress report which presents early results from the data analysis and the initial conceptual framework (filtrations and optimal stopping applications to AVs), which together represent first steps towards building AV controllers based upon them. To our knowledge, there is no previous work related to filtering pedestrian-vehicle action sequences. A review on different approaches for pedestrian crossing behavior modelling and analysis is provided in [10]. Methods of analysis are often performed via video recording, semi-structured interviews and VR recording. Previous work on pedestrian crossing behavior analysis can be found in [14] [5] [9] [12] [19]. Rasouli et al. introduce [14] [15] a novel dataset composed of 650 video-clips for driver-pedestrian interactions in several locations and different weather conditions. The analysis of their data show that attention plays an important role, as in 90% of the time, pedestrians reveal their intention of crossing by looking at the approaching vehicles. Rasouli et al. also present some behavioral patterns that have been observed in their data, that show some frequent sequences of actions that are used by pedestrians in their crossing behavior. Similar to our approach [13] uses task analysis to divide pedestrian-vehicle interaction as a sequence of actions giving two outcomes, either the vehicle passes first or the pedestrian crosses and perform some experiments with participants on their crossing behavior using virtual reality. In [5], Gorrini et al. analyzed video data of interaction between pedestrians and vehicles at an unsignalized intersection using semi-automatic tracking. Their study shows that pedestrian crossing behaviour can be divided into 3 phases: approaching (stable speed), appraising (deceleration due to evaluation of speed and distance of oncoming vehicles) and crossing (acceleration). Papadimitriou et al. [12] made a comparison of 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 behavior is in accordance with their declared behaviors from a questionnaire survey and they report that female and male participants have similar crossing behavior. In [9] drivers' crossing behaviour model in China at unsignaled intersections is presented using game theory and their risk perception is inferred via an adaptative neuro-fuzzy inference system. Previous works [18] [9] [11] have 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 but not really into how the interaction can be modelled as a sequence of the actions, more meaningful for autonomous systems. Surprising results have suggested that for autonomous vehicles, some apparently intuitive human communication styles might not be necessary for interactions with pedestrians. [3] showed that facial communication cues such as eye contact do not play a major role in pedestrian crossing behaviors, and that the motion pattern and behavior of vehicles are more important. Human drivers and pedestrians check for eye contact when the vehicle moves in an unexpected manner [3] [17]. However [6] showed that pedestrians can use eye-contact to influence drivers behavior and make them stop more often at crossings, which has the advantage of increasing the pedestrians' confidence while crossing. Similar results

from [16] show that vehicle movement is sufficient for indicating the intention of drivers and present some motion patterns of road users such as advancing, slowing early and stopping short. Statistical Filtration is a concept to incorporate events over time which is widely used in Optimal Stopping problems [2], such as the classic marriage problem which asks how many Tinder dates ones should attend and discard in order to infer the statistics of the population before marrying one of them. It is also used in finance, for example in pricing options to trade a stock at the days volume weighted average price (VWAP), which requires trades to be made before this price is fully known [7].

2 Methods

The study consists of data collection from real-world urban pedestrian-vehicle crossing interactions; definitions of features including descriptor, event, and outcome features; and analysis of outcome probability and value of information given the filtration of the features at each point in time.

2.1 Data collection

An ethnographic observation study on pedestrian-vehicle interactions was conducted at an unsignalized intersection near the University of Leeds, UK, following an exploration phase including observation of 70 pedestrian-vehicle interactions. An observation protocol was designed in which 204 road-crossing interactions were observed for the presence or absence of 62 temporal event features and 12 static descriptor features within each interaction, as listed in 1 and 2. Observers were positioned near the intersection as shown in figure 1, and worked together to identify and agree on when a vehicle-pedestrian pair took place in an ‘interaction’, one observing the vehicle and driver behaviour and one observing the pedestrian behaviour. From the start of the observation, each observer talked out loud about how the observed subject moved, communicated and reacted to the other observers subject, which allowed collaborative explanation of timely correct behaviour sequences. After the interaction, both observers filled in a observation protocol form from start till end, one typing, one controlling. Observation and data collection was conducted in accordance with University of Leeds Ethics and Data Protection regulations.

2.2 Data preparation

The winner (i.e. the pedestrian or vehicle which takes priority in the conflicted space and passes by the other) for each interaction was determined, and one of two new events were inserted into the filtrations to annotate this at the time it becomes actualised: (*Vehicle passed the Pedestrian* or *Vehicle stopped for observed Pedestrian*). We say ‘the game is over’ at the time within the interaction when the outcome becomes actualised. In most cases, features continued to be collected after the end of the game, and these remain preserved in the sequences. We then re-estimated the frequencies $freq(W, f_i)$ and $freq(\neg W, f_i)$ using Good-Turing estimation (add-one to each observed frequency) to deal with unobserved events in our data, and computed the normalized likelihoods,


$$\lambda(W|f_i) = \frac{freq(W|f_i)}{freq(W|f_i) + freq(\neg W|f_i)},$$

Descriptor Features d_i	$\lambda(W d_i)$
Pedestrian: older person (60+ years)	0.87542
Pedestrian: teenager (13-18y)	0.7008
Weather: Rainy	0.568513
Pedestrian's Distraction	0.5506
Pedestrian: midage adult (30-60y)	0.54586
Pedestrian: Gender (Female)	0.54111
Weather: Sunny	0.5033
Weather: Overcast	0.4875
Pedestrian: young adult (18-30y)	0.46762
Group of Pedestrians	0.42586
Driver/Vehicle Interacting Vehicle Coming From right	0.37752
Driver/Vehicle Interacting Vehicle is Single	0.37137

Table 1: All 12 descriptor features used for the observation of Pedestrian-Vehicle Interaction listed by descending order of likelihood λ_{d_i}



(a) Vehicle passes first, Pedestrian stepped back on the pavement (b) Group of pedestrians crossing while the vehicle stops for them



(c) Pedestrian initiating crossing (d) Vehicle stopped due to traffic but finally decides to walk further while the Pedestrian is approaching on the left after the vehicle passes but the vehicle passed first

Fig. 2: Examples of observed pedestrian-vehicle interactions at the unsignalized intersection

Event Features e_i	$\lambda(W e_i)$
'Crossing Phase: Pedestrian Speeded up'	0.95471
'Crossing Phase: Driver/Vehicle Decelerated for observed pedestrian'	0.84051
'Crossing Phase: Driver/Vehicle Used signals Turn Indicator'	0.77844
'Approaching Phase: Driver/Vehicle Waved hand'	0.77844
'Approaching Phase: Driver/Vehicle Head Movements Other (elaborate in notes)'	0.7784
'Crossing Phase: Driver/Vehicle Movement Other (elaborate in notes)'	0.77844
'Crossing Phase: Pedestrian Raised hand in front'	0.77844
'Crossing Phase: Driver/Vehicle Raised hand in front'	0.77784
'Crossing Phase: Driver/Vehicle Head Turned in the direction of pedestrian'	0.77784
'Crossing Phase: Driver/Vehicle Stopped for observed pedestrian'	0.77784
'Crossing Phase: Pedestrian Looked at driver'	0.77784
'Approaching Phase: Driver/Vehicle Stopped due to other pedestrian'	0.77784
'Crossing Phase: Pedestrian Movements Other (elaborate in notes)'	0.77844
'Crossing Phase: Pedestrian Initiated crossing movement'	0.7712
'Approaching Phase: Driver/Vehicle Head Turned in the direction of pedestrian'	0.74541
'Crossing Phase: Pedestrian Head Movements Turned left'	0.7454
'Approaching Phase: Driver/Vehicle Interacting vehicle Bus / Truck'	0.72490
'Approaching Phase: Vehicle Stopped for observed pedestrian'	0.7008
'Crossing Phase: Pedestrian Looking at other pedestrians entering the road'	0.6372
'Crossing Phase: Pedestrian Waved Hand'	0.63725
'Approaching Phase: Driver/Vehicle Head Turned left'	0.6372
'Approaching Phase: Driver/Vehicle Movement Other (elaborate in notes)'	0.6372
'Approaching Phase: Pedestria Hand Movements Other (elaborate in notes)'	0.6372
'Crossing Phase: Driver/Vehicle Turned left'	0.6372
'Crossing Phase: Vehicle Waved hand'	0.63725
'Crossing Phase: Driver/Vehicle Accelerated'	0.63725
'Crossing Phase: Driver/Vehicle Turned right'	0.6372
'Approaching Phase: Pedestrian Looking at other pedestrians entering the road'	0.6372
'Approaching Phase: Pedestrian Looking at other RUs Others (elaborate in notes)'	0.6372
'Approaching Phase: Driver/Vehicle Used signals Flashed Lights'	0.6372
'Approaching Phase: Pedestrian Movements Kept pace'	0.6231
'Approaching Phase: Vehicle Used signals Turn Indicator'	0.559
'Crossing Phase: Driver/Vehicle Passed the pedestrian'	0.5394
'Approaching Phase: Pedestrian Movements Did not Stop'	0.5365
'Approaching Phase: Pedestrian Head Movements Turned right'	0.53485
'Approaching Phase:Driver/Vehicle approached From left'	0.5292
'Approaching Phase: Driver/Vehicle Decelerated due to other pedestrians'	0.5131
'Approaching Phase: Driver/Vehicle Stopped due to traffic'	0.51315
'Approaching Phase: Driver/Vehicle approached from Multiple'	0.5009
'Approaching Phase: Driver/Vehicle Decelerated for observed pedestrian'	0.4875
'Approaching Phase: Pedestrian Speeded up'	0.46762
'Crossing Phase: Pedestrian Raised hand sideways'	0.4676
'Approaching Phase: Driver/Vehicle Interacting vehicle Other (elaborate in Notes)'	0.4676
'Crossing Phase: Pedestrian Stepped back on pavement'	0.4676
'Approaching Phase: Driver/Vehicle Turned left'	0.45419
'Approaching Phase: Pedestrian Stopped at the edge of the pavement'	0.43844
'Approaching Phase: Pedestrian Stepped on road and stopped'	0.42951
'Approaching Phase: Pedestrian Head Movements Turned left'	0.42951
'Approaching Phase: Pedestrian Movements Slowed down'	0.4260
'Crossing Phase: Pedestrian Looking at Looked at vehicle'	0.41269
'Approaching Phase: Driver/Vehicle Decelerated due to traffic'	0.3874
'Crossing Phase: Pedestrian Hand Movements Other (elaborate in notes)'	0.36931
'Approaching Phase: Driver Head Turned right'	0.36931
'Approaching Phase: Driver/Vehicle Interacting vehicle Van'	0.3693
'Approaching Phase: Driver/Vehicle Kept pace'	0.36931
'Approaching Phase: Driver/Vehicle Turned right'	0.3598
'Crossing Phase: Pedestrian Head Movements Turned right'	0.3341
'Approaching Phase: Pedestrian Looked at approaching vehicle'	0.3129
'Crossing Phase: Pedestrian Looking at other RUs (elaborate in comments)'	0.26
'Crossing Phase: Pedestrian Slowed down / stopped	0.26
'Approaching Phase: Driver/Vehicle Accelerated'	0.163316
'Approaching Phase: Driver/Vehicle Passed the pedestrian'	0.11514

Table 2: All 62 event features used for the observation of Pedestrian-Vehicle Interaction listed by descending order of likelihood λ_{e_i}

for each feature (where W means pedestrian wins and $\neg W$ means that the pedestrian does not win, i.e. the vehicle wins) given that we observed individual features f_i anywhere during the interaction.

2.3 Filtration

The normalized likelihood representation allows for simple filtration-based fusion of the features. A filtration, $\mathcal{F}(t)$ is a monotonically growing set of observations (feature) available over time t , such that $\mathcal{F}(t + \Delta t)$ comprises $\mathcal{F}(t)$ and all new features observed in the interval $(t + \Delta t)$. Normalized likelihoods are single values in range $[0,1]$ which yield posteriors when Bayes-fused with priors, as,

$$P(W|f_1, f_2, f_n) = P(W|0) \otimes \lambda(W|f_1) \otimes \lambda(W|f_2) \otimes \dots \otimes \lambda(W|f_n),$$

where the Bayesian fusion operator is,

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

and $P(W|0)$ is the prior. Using this form of filtering, a real time system can begin with a prior estimate about the winner, and iteratively fuse in each normalized likelihood as it becomes available.

Two types of feature appear in the filtration: static *descriptor* features which describe a property of the entire interaction and are observable from start, such as age and gender of the pedestrian, weather, or time of day; and temporal *event* features which occur and are observable at a particular instant during the interaction, such as the presence of eye-contact, placing a foot in the road, or making a hand gesture. We assume that the prior $P(W|0)$ is available at $t = 0$; that all present descriptor features d_i become available *together* at the start of the interaction at time $t = 1$; and that the temporal event features e_i are revealed to us *one at a time* at times $t(e_i)$. In the present study we assume these are integer valued times which correspond directly to the events index in the observed sequence, rather than exact real-valued times of the events (this is for simplicity as this is a just proof of concept study.) Thus, $\mathcal{F}(t) = \{d_i\}_{\forall i} \cup \{e_j\}_{j=2:t}$, and for each interaction, we infer the sequence of probabilities $P(W|\mathcal{F}(t))$.

2.4 Residual filtration posterior volatility

From the shapes of the filtrations, we wish gain insights useful for the design of a real-time AV controller. In particular: what is the balance of information contained in the initial descriptors which are available right away, vs information arriving later in the filtration? This should give insight into how much future AV controllers need to care about optimal stopping issues: whether they should typically act immediately on detection of a new interaction and/or of its descriptors, or whether they should bide their time collecting more information before acting.

As a measure of this value of information over time, we thus define a series of *residual filtration posterior volatility* statistics,

$$s_t = \langle std\{P(W|\mathcal{F}(\tau))\}_{\tau=t:T} \rangle,$$

where the expectation is over interactions, and *std* is the standard deviation operator. These statistics are *second order* statistics: they measure the volatility of our own belief over time during interactions. The statistic s_t for time t measures how much fluctuation is expected to occur in our own beliefs over the remaining time within the interaction.

3 Results

Fig. 3(a) shows the distribution of durations of games (as number of features in the sequence before the game is over) and Fig. 3(b) shows the distribution of the longer interaction durations, which also include observations of features after the game is over. In the absence of any other information, the prior on the outcome is computed as $P(W|0) = 36\%$ (74 of 204) of the interactions which resulted in the pedestrian winning (defined as by their passing through the conflict area before the vehicle). There were no observed collisions between vehicles and pedestrians.




Fig. 3: Lengths of games and interactions

3.1 Filtering: effect of observations over time on outcome

Figure 4 and 5 show examples of the posterior $P(W|\mathcal{F}(t))$ for particular interactions. We inserted two new types of event, *Pedestrian wins* and *Vehicle wins*, when we observe a feature that determines the end of the interaction such as *Vehicle passed the pedestrian* or *Vehicle stopped for observed pedestrian*, the rest of the sequences then becomes no more interesting as we already know the winner of the interaction.

3.2 Residual filtration posterior volatility profiles

Fig. 7 shows the series of s_t statistics computed over all interactions and Fig. 6 for ten particular interactions. This shows that the expected volatility of belief decreases over time up to around time 5-10, then levels off. The decrease period includes both the initial observation of the descriptor features, plus the first few event features, but not later event features.

4 Discussion

The filtration results in Fig. 7 show that there is high volatility in belief about the interaction outcome at the start of an interaction, decreasing rapidly as first descriptors

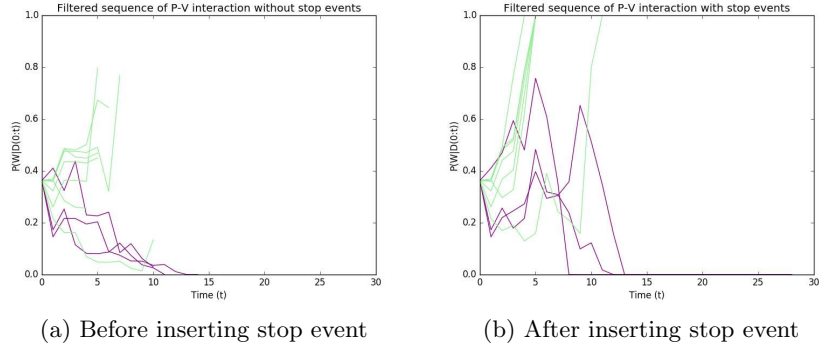


Fig. 4: Belief filtrations for 10 particular sequences

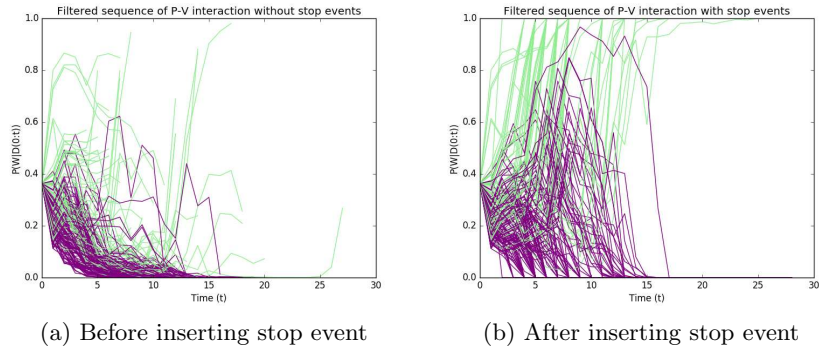


Fig. 5: Belief filtrations for all interaction sequences

are available then as the first few (1-10) event features are seen, then leveling off. From fig. 3 it is seen that these all occur before the end of most games. This means that waiting to observe these first event features is a useful strategy for an AV, but it then gains little additional information from waiting to observe any more, and may risk the game ending before making a decision if it waits much longer. Thus, the results suggest that AVs should not act right away on detecting a road crossing interaction, but rather wait to observe just the first few informative features before acting (by speeding up or slowing down). In turn, this suggests that Optimal Stopping based models would be a fruitful research area for AV controllers for pedestrian interactions. The game theory model of [4] in the absence of such filtration information has an optimal strategy which gradually increases the probability of yielding actions over time. This has the same general form as found here – to wait a while before acting – but future work should now determine how to fuse the value of information found in this present study with the values of arrival time and collision avoidance of the previous study. It seems likely that a combined model will continue to be probabilistic in optimal strategy.




Fig. 6: Residual filtration posterior volatility profile for each of the 10 interactions

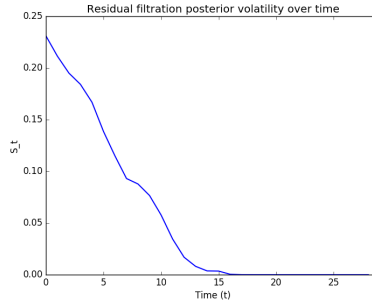


Fig. 7: Residual filtration posterior volatility profile for all interactions

The present study is intended as an early proof of concept only and as such makes several simplifying assumptions. The simple integer sequence indexes used to measure time here should be replaced in future work with the actual real times of events. The naive Bayes assumption that all features are independent given the class should be questioned in future, both in terms of features normalized likelihood strengths and also their occurrence in time. Some kind of Poisson process may be required to model their joint distribution over time, as Poisson processes model distributions of discrete events occurring randomly in continuous time.

We considered the prediction only of the interaction outcome – who wins – rather than the inference of the underlying latent assertiveness variables θ involved in causing the outcome. This is an important first step towards inferring θ , and the winners of historical interaction data might in some data-mining style cases be used as a hard (integer 0 or 1) approximation for their soft (real 0-1 ranged) θ . Future work should consider how to make the inference more precisely and also how to infer and separate the effect of the drivers own assertiveness θ_{driver} from the pedestrians, as ultimate it is historical and real-time pedestrian θ values which are required to inform the real-time game theoretic AV controller of [4].

We currently assume that all information is contained by the *presence* of features in the filtration, and that their time and sequence of their occurrence within the filtration is unimportant. More advanced models could consider the non-presence of a feature

within each filtration time interval as a potentially informative alternative feature; and also consider the presence and absence of motifs of sequences or noisy sequences of features as additional features.

The features used here were observed and logged by trained human ethologists. It seems likely that machine vision systems could match their performance for some of the features, such as macroscopic positions and motions of the agent, but less likely for others such as eye contact and gestures. The data collection operated in a single intersection in the UK and it is possible that features may have different likelihoods in other intersections and countries.

Some non-Bayesians such as Dempster-Shafer theorists still suggest that Bayesian theory is incapable of reasoning about the distinction between uncertainty -in-the-world and uncertainty-in-our-beliefs. This is not true, and the use of second-order Bayesian probability (or ‘meta probability’) is a counterexample to the claim. Whilst this is not the first application of such a method, it is an especially simple and concrete one which may be useful in rebutting such claims in future. Our use of a hand-crafted statistic s_t remains somewhat frequentist and should be replaced by a full Bayesian model at the second-order level in future work, which might for example employ mixtures of Beta distributions to model $P(P(W|\mathcal{F}(t + \tau))|\mathcal{F}(t))$, the expected future beliefs at time $t + \tau$ at wall-clock times t . Optimal stopping algorithms could then be constructed from expected entropy reductions in these second-order belief distributions given the extended observation times.

References

1. CAMARA, F., ROMANO, R., MARKKULA, G., MADIGAN, R., MERAT, N., AND FOX, C. Empirical game theory of pedestrian interaction for autonomous vehicles. In *Proceedings of Measuring Behavior* (2018), Manchester Metropolitan University, UK.
2. DEGROOT, M. H. *Optimal statistical decisions*, vol. 82. John Wiley & Sons, 2005.
3. DEY, D., AND TERKEN, J. 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* (New York, NY, USA, 2017), AutomotiveUI ’17, ACM, pp. 109–113.
4. FOX, C. W., CAMARA, F., MARKKULA, G., ROMANO, R., MADIGAN, R., AND MERAT, N. When should the chicken cross the road?: Game theory for autonomous vehicle - human interactions. In *Proc. 5th International Conference on Vehicle Technology and Intelligent Transport Systems (VEHITS)* (2018).
5. GORRINI, A., VIZZARI, G., AND BANDINI, S. Towards modelling pedestrian-vehicle interactions: Empirical study on urban unsignalized intersection. *CoRR abs/1610.07892* (2016).
6. GUGUEN, N., MEINER, S., AND EYSSARTIER, C. A pedestrian’s stare and drivers’ stopping behavior: A field experiment at the pedestrian crossing. *Safety Science* 75 (2015), 87 – 89.
7. HUMPHERY-JENNER, M. L. Optimal VWAP trading under noisy conditions. *Journal of Banking & Finance* 35, 9 (2011), 2319–2329.
8. KATO, S., TAKEUCHI, E., ISHIGURO, Y., NINOMIYA, Y., TAKEDA, K., AND HAMADA, T. An open approach to autonomous vehicles. *IEEE Micro* 35, 6 (Nov 2015), 60–68.

9. LIU, M., CHEN, Y., LU, G., AND WANG, Y. Modeling crossing behavior of drivers at unsignalized intersections with consideration of risk perception. *Transportation Research Part F: Traffic Psychology and Behaviour* 45 (2017), 14 – 26.
10. MAMIDIPALLI, S., SISOPIKU, V. P., SCHROEDER, B., AND ELEFTERIADOU, L. A review of analysis techniques and data collection methods for modeling pedestrian crossing behaviors. *Journal of Multidisciplinary Engineering Science and Technology Vol. 2, Issue 2* (February 2015).
11. MITROVIC, J. S., BASARIC, V., AND BOGDANOVIC, V. Pedestrian crossing behaviour at unsignalized crossings. *Horizons* (2016).
12. PAPADIMITRIOU, E., LASSARRE, S., AND YANNIS, G. Pedestrian risk taking while road crossing: A comparison of observed and declared behaviour. *Transportation Research Procedia* 14 (2016), 4354 – 4363. Transport Research Arena TRA2016.
13. PILLAI, A. Virtual reality based study to analyse pedestrian attitude towards autonomous vehicles. G2 pro gradu, diplomity, 10 2017.
14. RASOULI, A., KOTSERUBA, I., AND TSOTSOS, J. K. Agreeing to cross: How drivers and pedestrians communicate. *2017 IEEE Intelligent Vehicles Symposium (IV)* (2017), 264–269.
15. RASOULI, A., KOTSERUBA, I., AND TSOTSOS, J. K. Understanding pedestrian behavior in complex traffic scenes. *IEEE Transactions on Intelligent Vehicles* 3, 1 (March 2018), 61–70.
16. RISTO, M., EMMENEGGER, C., VINKHUYZEN, E., CEFKIN, M., AND HOLLAN, J. Human-vehicle interfaces: The power of vehicle movement gestures in human road user coordination. pp. 186–192.
17. ROTHENBUECHER, D., LI, J., SIRKIN, D., MOK, B., AND JU, W. Ghost driver: A field study investigating the interaction between pedestrians and driverless vehicles. In *2016 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)* (Aug 2016), pp. 795–802.
18. SCHNEEMANN, F., AND GOHL, I. Analyzing driver-pedestrian interaction at crosswalks: A contribution to autonomous driving in urban environments. In *2016 IEEE Intelligent Vehicles Symposium (IV)* (June 2016), pp. 38–43.
19. SUN, D., UKKUSURI, S., BENEKOHAL, R., AND TRAVIS WALLER, S. Modeling of motorist-pedestrian interaction at uncontrolled mid-block crosswalks.