Game Theory, Proxemics and Trust for Self-Driving Car Social Navigation

Fanta Camara and Charles Fox

***Abstract*— To navigate in human social spaces, self-driving cars and other robots must show social intelligence. This involves predicting and planning around pedestrians, understanding their personal space, and establishing trust with them. The present paper gives an overview of our ongoing work on modelling and controlling human–self-driving car interactions using game theory, proxemics and trust, and unifying these fields via quantitative models and robot controllers.**

I. INTRODUCTION

Recent years have witnessed the rapid deployment of robotic systems in many places such as roads, pavements, workplaces and care homes [27], [35]. Robot navigation in environments with static objects is largely solved, but navigating around humans in dynamic environments remains an active research question. To operate in human social spaces, robots must show social intelligence, i.e. the ability to understand human behaviour via explicit and implicit communication cues for better human-robot interactions (HRI) [33]. Autonomous vehicles (AVs), also known as "self-driving cars" are appearing on the roads but need better understandings of pedestrians' social behaviour, especially in urban areas [31]. In particular, previous work showed that pedestrians may take advantage over autonomous vehicles [17] by intentionally and constantly stepping in front of AVs, hence preventing them from making progress on the roads, this is known as the 'freezing robot problem' [34]. This inability of current AVs to read the intention of other road users, predict their future behaviour and interact with them is described as 'the big problem with self-driving cars' [2]. Thus, AVs need better decision-making models and must find a good balance between stopping for pedestrians when required and driving to reach their final destination as quickly as possible for their on-board passengers.

We recently performed two comprehensive reviews of existing pedestrian models for AVs, ranging from low-level sensing, detection and tracking models [3] to high-level interaction and game theoretic models of pedestrian behaviour [4]. These reviews found that existing lower-level models are accurate and mature enough to be deployed on AVs but more research is needed in the higher-level models. Hence, in our work, we focus on modelling, learning and operating pedestrian high-level social behaviour on self-driving cars using game theory, proxemics and trust.

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

Fanta Camara and Charles Fox are with the Institute for Transport Studies (ITS), University of Leeds, UK, and with the Lincoln Centre for Autonomous Systems (L-CAS), University of Lincoln, UK.

Correspondence: Fanta Camara, tsfc@leeds.ac.uk/fcamara@lincoln.ac.uk

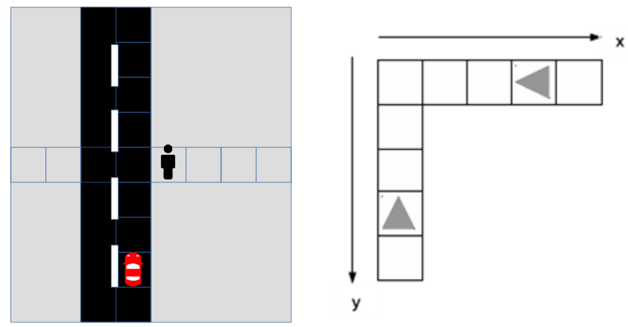

(a) Road-Crossing Scenario (b) Game Theory Model [19]

Fig. 1: Two agents try to cross over an intersection as quickly as possible while avoiding a collision. The first agent to pass wins the game (reward), the second looses (small penalty) and they are both bigger losers if there is a collision (large penalty).

II. SEQUENTIAL CHICKEN GAME THEORY MODEL

Game theory is widely used to model decision-making between rational agents, in economics [26] and in multi-agent systems coordination [25]. Its core concept is Nash equilibrium, which is a set of optimal actions for the agents in the absence of information about the other's choices.

We have created a game theoretic model [19] of a pedestrian and an AV negotiating for shared space as a pedestrian considers whether to cross in front of the AV's path as in Fig. 1a. This is modelled as a discrete sequential game theory model based on the 'game of chicken' as shown in Fig. 1b. The model has two utility parameters (U_{time} , U_{crash}), which refer to the value of time and the utility of avoiding a collision. The pedestrian X and autonomous driver Y are moving towards each other at an unmarked intersection. This occurs over a discrete space (the path is formed of squares) as in Fig. 1b and discrete times ('turns') during which the agents can adjust their discrete speeds. Here a turn corresponds to one discrete time step, i.e. the time offered to the agents to make a new decision. They simultaneously select their speed of either 1 square per turn (SLOW) or 2 squares per turn (FAST), at each turn. Space and time are discrete to keep the model simple and computationally tractable. 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 penalty, $-TU_{time}$, where T is the time from the start of the game and U_{time} is the value of one turn of travel time.

The model assumes the two players choose their actions (speeds) $a_Y, a_X \in \{1, 2\}$ simultaneously then implement them simultaneously, in 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 optimal strategies are derivable from game theory together with a novel meta-strategy convergence solution concept, via recursion. Sequential chicken can be viewed as a sequence of one-shot sub-games, whose payoffs are the expected values of new games resulting from the actions, and are solvable by standard game theory.

Results from the mathematics and simulations agreed on a solution to the freezing robot problem. If the vehicle is programmed to be perfectly safe and always yield to pedestrians, it will freeze and never make any progress in a series of pedestrian interactions. Instead, in this model, it *must* be programmed to *deliberately collide* with the pedestrian with a small but non-zero probability. This provides a *credible threat* sufficient for the vast majority of interactions to proceed without collision, but with some of the pedestrians yielding to allow the vehicle to make progress.

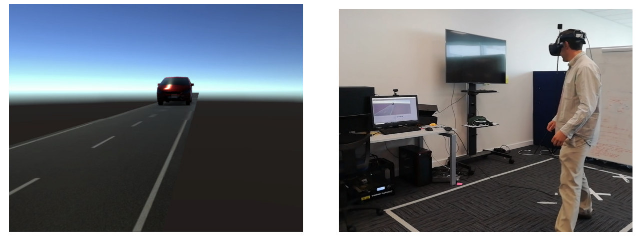

Fig. 2: Virtual AV and participant in the VR Experiment.

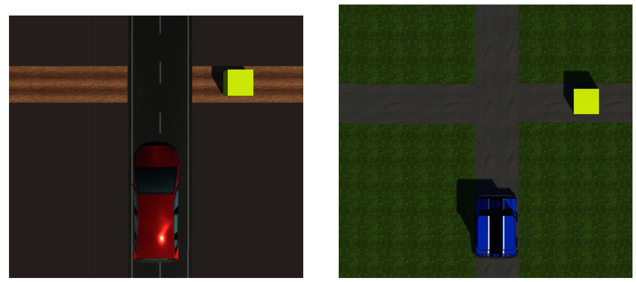

Fig. 3: Top view of the scenes used for the VR experiment.

III. Learning parameters from real-world & virtual reality experiments

To validate this game theory model [19], we ran several experiments with human participants to infer the utility parameters. In a first empirical study [14], we measured participants' behaviour whilst playing the game theory model as a board game. The parameters of the model could be inferred via a Gaussian Process (GP) regression [30] and the results showed that participants had a preference for saving time, U_{time}, rather than avoiding a collision, U_{crash}. This study provided a first empirical understanding of the human factors required by future game theoretic autonomous vehicles. In a second study [5], we developed a novel empirical method based on tracking real humans in a semi-structured environment and used their discrete positions to model and predict their behaviour with game theory. We made use of dynamic programming to compute the optimal game theoretic solution form, then found the behavioural parameters via empirical observation and a GP regression analysis. This method formed a step towards game-theoretic controllers for autonomous vehicles in similar real-world situations such as negotiations over priority at un-signalled road-crossings. This second study showed that participants were globally playing rationally, 11% of them deviated from their optimal behaviour. It also confirmed participants' preference for time saving rather than collision avoidance, this unusual result was due to the high safety conditions of the experiment. In a third study [6], we extended our method to work with continuous trajectory data and found similar results as in [5]. The biggest drawback with these three empirical experiments is that participants were behaving as if they were in a competition and so their preferences for saving time was rather unrealistic and discarding the high negative utility associated with a real-world crash, which could potentially lead to death. As a consequence, we moved to virtual reality (VR) to further investigate human interaction preferences in two virtual environments, as shown in Figs. 2 and 3, because VR offers the opportunity to experiment on human behaviour in simulated real-world environments that can be dangerous or difficult to study. We developed a virtual game theoretic autonomous vehicle that interacted with human participants [7], [8], [18]. The results showed a more realistic crossing behaviour from participants, preferring avoiding collisions with the virtual AV rather than saving time. When presented with different AV behaviours via a gradient descent approach, participants preferred an AV that makes its decisions quickly. Finally, we found similar crossing behaviours in both virtual environments, as previously shown in [28].

IV. Learning Interaction Sequence patterns from Pedestrian–Vehicle data

To learn social behavioural patterns from current pedestrian–vehicle interactions, we collected a large-scale dataset from real-world human road crossings [12]. Pedestrian-vehicle interactions were decomposed into sequences of independent discrete events. We looked for common patterns of behaviour that can predict the winner of an interaction. We used logistic regression, decision tree regression, and motif analysis to find sub-sequences of actions used by both pedestrians and human drivers while crossing. We found predictive features that could inform the AV about the eventual winner of an interaction. We then used the same dataset to study the temporal orderings (filtration) in which features (including signals from the pedestrian) can be revealed to an autonomous vehicle and their informativeness over time during pedestrian-vehicle interactions [11]. This framework suggested 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

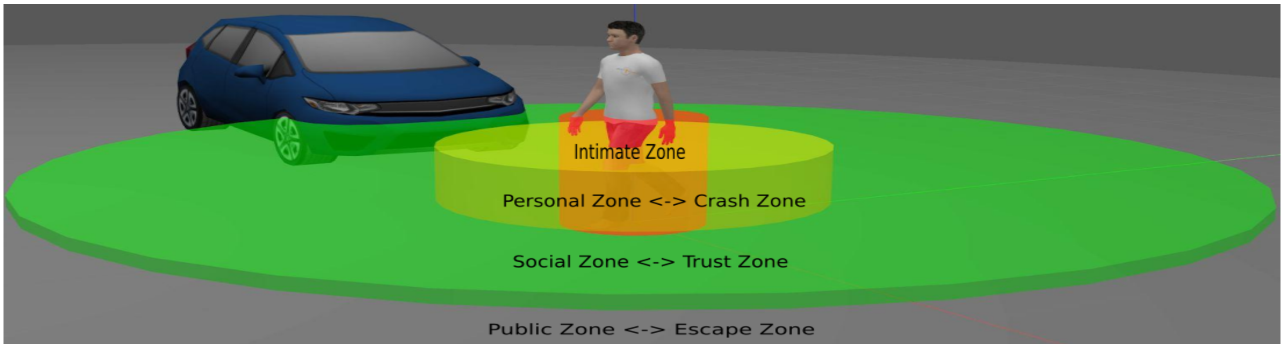

Fig. 4: Autonomous vehicle entering pedestrian's social zone, which can also be viewed and quantified as a trust region.

features, including signals from that pedestrian, before acting itself. Here we found that the AV should wait and observe about 7 to 10 features before acting/making its decision. Using the public Daimler pedestrian dataset [21] we also developed simple heuristic features which can be fused to predict road crossing intent to some extent [13]. These predictions could be integrated into Sequential Chicken based AV controllers as priors to improve their predictions and interactions.

V. Unifying & Quantifying Proxemics and Trust

The Sequential Chicken model showed that if the vehicle's only way to inflict negative utility onto pedestrians is to actually hit them, then it must be programmed to deliberately provoke a crash occasionally in order to make progress. This is clearly an undesirable and unethical solution to the real-world freezing robot problem. However, the equations of the model still work if other, less violent, forms of negative utility are made available for the vehicle to inflict upon members of the public, with higher frequency traded for lower damage. Possible solutions that are currently under debate might include spraying jets of water at anti-social pedestrians intentionally blocking the AV's path, or humiliating them in public using horns as is often done my human drivers to penalise other road users' anti-social behaviour [36].

An intriguing additional option which we have chosen to explore is to make over-assertive pedestrians feel uncomfortable by invading their proxemic space. The theory of proxemics was introduced by Hall [20] to describe humans' psychological sense of comfort or discomfort during physical interactions. Hall proposed four distinct zones: the intimate up to 0.45m, the personal ranges from 0.45m to 1.2m, the social between 1.2m to 3.6m, and the public beyond 3.6m [23]. Social robotics experiments have shown that these proxemic zones change in size when humans interact with robots of different heights, appearances, speeds, voices, and also for different HRI activities [32].

We recently developed the first mathematical model of proxemics and trust concepts for self-driving cars and pedestrians interactions [10]. It defined the trust zone as the area of the proxemics zones where trust is required i.e., one agent has to rely on the other during the interaction. In [10], we define a trust zone as a set of physical states in which a first agent, called *Agent1* is in a position of vulnerability and has to rely on the actions of a second agent, called *Agent2*. The trust zone comprises those states in which Agent2 can choose whether to stop for the Agent1 but Agent1 cannot avoid a collision by themselves, given their physical speeds and stopping distances. Formally, the work in [10] defines *physical trust requirement* (PTR) as a Boolean property of the physical state of the world (not of the psychology of the agents) with respect to Agent1 during an interaction, true if and only if Agent1's future utility is affected by an immediate decision made by Agent2. This model assumes that the two agents are approaching each other at a right angle, as is the case where one crosses the other's path, as shown in Fig. 4. It then defines the following three zones based on the PTR:

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

$$d_{crash} = v_2 t_2 + \frac{v_2^2}{2\mu_2 g}, \quad (1)$$

in which a crash is guaranteed and neither party can prevent it [24], v_2 is Agent2's speed. The first term depends on Agent2's thinking reaction time, t_2, and the second term represents the physical braking distance if Agent2 is a wheeled agent, μ_2 is the coefficient of friction between Agent2's tyres and tarmac, and g is gravity. If Agent2 is a walking agent, we will here assume this second term is omitted as walkers are always in static equilibrium and can stop instantly once a decision is made. Running agents [22] or finer detailed models of walkers [29] could use different braking models.

Escape zone is the area where Agent1 is able to choose their own action to avoid collision, without needing to trust Agent2 to behave in any particular way. If w_2 is the width of Agent2, which Agent1 must cross at speed v_1 to pass first, the escape zone is then the set $\{d : d_{escape} < d\}$ with

$$d_{escape} = v_2 t_1 + w_2 \frac{v_2}{v_1}. \quad (2)$$

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

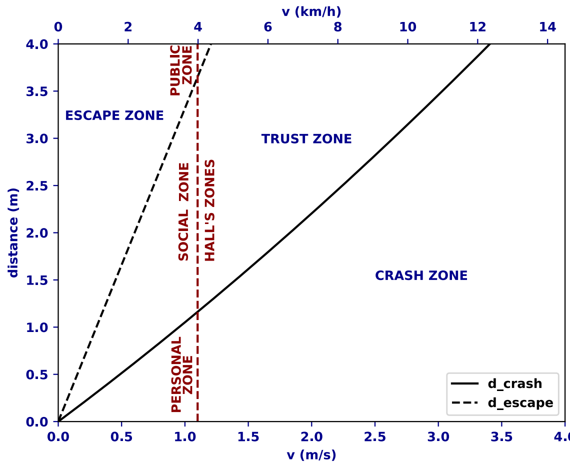

Fig. 5: Distances and zones predicted by the PTR model for different car speeds v at lower speeds.

collision, but Agent2 is able to slow to prevent the collision if it chooses to yield.

Note that these zones are not symmetric between Agent1 and Agent2. They describe when Agent1 must trust Agent2. Their roles must be swapped and the zones recomputed to see when Agent2 must trust Agent1. The crash, escape, and trust zones were mapped to Hall's personal, public, and social zones respectively, for Agent1 [10]. The trust/social zone is the region in which physical trust is required. This may be a prerequisite for some types of interactions, with physical trust being useful to enable the content of the interaction. The evidence for this mapping came from the observation that if an autonomous vehicle Agent2 is set to drive at the same speed as a pedestrian Agent1, the model generates Hall's proxemic social zone to within 4% quantitative accuracy of Hall's original empirical sizes. This unexpected result, found by studying how an AV should interact with pedestrians, may now explain a larger question about how humans interact with each other and with other types of robots. Hall's zone sizes have previous been only empirical observations but the PTR model now explains them generatively and to 4% accuracy for the first time. We then extended this model for more general human-human interactions and HRI by taking different interaction headings into account and found an error down to 1% [9].

VI. OpenPodcar: open source hardware platform for social navigation research

Laboratory and VR experiments on pedestrians are limited in realism, so to scale our models towards the real world, we need a real autonomous vehicle. Commercial AVs are very expensive, beyond reach of our most other labs who may wish to replicate and extend our work. So we have developed a new low-cost, autonomous vehicle research platform, OpenPodcar shown in Fig. 6 and based on an off-the-shelf, hard-canopy, mobility scooter donor vehicle. We are

Fig. 6: OpenPodcar: open source hardware AV [15].

releasing OpenPodcar [15] as open source hardware (OSH, [1]) together with a full automation open source software (OSS) stack. This will enable other groups to replicate our complete system and experiments, and to use their own research to extend and contribute to a single shared system, which can evolve over time towards real-world use.

Hardware and software designs are provided to convert the donor vehicle into a low-cost and fully autonomous platform. The open source platform consists of (a) hardware components: CAD designs, bill of materials, and detailed build instructions; (b) Arduino, ROS [16] and Gazebo software files which provide standard interfaces and physical simulation for the vehicle; (c) higher-level ROS implementations of standard robot control, including the move_base interface with Timed-Elastic-Band planner which enacts command to get the vehicle move from one pose to another, (d) lidar based pedestrian detection and tracking. The platform is large enough to transport one person at speeds up to 15km/h, for example for use as a last-mile autonomous taxi service or to transport delivery containers similarly around a city center. It is small and safe enough to be parked in a standard research lab and be used for realistic human-vehicle interaction studies. System build cost from new components is around 7,000USD in total in 2022, with our own build from second-hand components costing around 2,000USD.

VII. Summary

The sequential chicken model showed that if an AV's only available actions are to yield to pedestrians or drive forward, then it must be programmed with a small collision probability in order to create a credible threat to solve the freezing robot problem. However, we found that the rare large negative utility of a crash can be replaced by more frequent but smaller penalties using human proxemic preferences. Our proxemics study has provided a generative explanation for the numerical sizes of Hall's empirical zones for the first time, and can now be used to control and adapt the AV's behaviour more safely. OpenPodcar platform will enable other researchers to run and extend this interaction model and we invite the community to join in.

REFERENCES

[1] J. Bonvoisin, J. Molloy, M. Häuer, and T. Wenzel. Standardisation of practices in open source hardware. *arXiv preprint arXiv:2004.07143*, 2020.

[2] R. Brooks. The big problem with self-driving cars is people. *IEEE Spectrum: Technology, Engineering, and Science News*, 2017.

[3] 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.

[4] 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 behavior. *IEEE Transactions on Intelligent Transportation Systems*, 2020.

[5] 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.

[6] F. Camara, S. Cosar, N. Bellotto, N. Merat, and C. W. Fox. *Continuous Game Theory Pedestrian Modelling Method for Autonomous Vehicles*. River Publishers, 2020.

[7] F. Camara, P. Dickinson, and C. Fox. Evaluating pedestrian interaction preferences with a game theoretic autonomous vehicle in virtual reality. *Transportation Research Part F: Traffic Psychology and Behaviour*, 78:410–423, 2021.

[8] F. Camara, P. Dickinson, N. Merat, and C. Fox. Towards game theoretic AV controllers: measuring pedestrian behaviour in virtual reality. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2019) Workshops*, 2019.

[9] F. Camara and C. Fox. Extending quantitative proxemics and trust to HRI. *(under review)*.

[10] F. Camara and C. Fox. Space invaders: Pedestrian proxemic utility functions and trust zones for autonomous vehicle interactions. *International Journal of Social Robotics*, 2020. https://doi.org/10.1007/s12369-020-00717-x.

[11] 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.

[12] F. Camara, O. Giles, R. Madigan, M. Rothmüller, P. H. Rasmussen, S. A. Vendelbo-Larsen, G. Markkula, Y. M. Lee, L. Garach, N. Merat, and C. W. Fox. Predicting pedestrian road-crossing assertiveness for autonomous vehicle control. In *Proceedings of IEEE 21st International Conference on Intelligent Transportation Systems*, 2018.

[13] F. Camara, N. Merat, and C. Fox. A heuristic model for pedestrian intention estimation. In *IEEE Intelligent Transport Systems Conference 2019*, 2019.

[14] F. Camara, R. Romano, G. Markkula, R. Madigan, N. Merat, and C. 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.

[15] F. Camara, C. Waltham, D. Churchill, and C. Fox. OpenPodcar: an open source vehicle for self-driving car research. *(under review)*.

[16] S. Cousins, B. Gerkey, K. Conley, and W. Garage. Sharing software with ROS. *IEEE Robotics & Automation Magazine*, 17(2):12–14, 2010.

[17] S. Deb, L. Strawderman, D. W. Carruth, J. DuBien, B. Smith, and T. M. Garrison. Development and validation of a questionnaire to assess pedestrian receptivity toward fully autonomous vehicles. *Transportation research part C: emerging technologies*, 84:178–195, 2017.

[18] F., Camara, P. Dickenson, N. Merat, and C. Fox. Examining pedestrian-autonomous vehicle interactions in virtual reality. In *Proc. Transport Research Arena 2020 (TRA2020)*, 2020.

[19] 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.

[20] E. T. Hall. *The hidden dimension*, volume 609. Garden City, NY: Doubleday, 1966.

[21] J. F. P. Kooij, N. Schneider, F. Flohr, and D. M. Gavrila. Context-based pedestrian path prediction. In *Proc. of ECCV*, pages 618–633, 2014.

[22] T. Kwon and J. K. Hodgins. Control systems for human running using an inverted pendulum model and a reference motion capture sequence. In *Symposium on Computer Animation*, pages 129–138, 2010.

[23] D. Lambert. *Body language*. HarperCollins, 2004.

[24] D. Lyubenov. Research of the stopping distance for different road conditions. *Transport Problems*, 6:119–126, 2011.

[25] Y. Meng. Multi-robot searching using game-theory based approach. *International Journal of Advanced Robotic Systems*, 5(4):44, 2008.

[26] O. Morgenstern and J. Von Neumann. *Theory of games and economic behavior*. Princeton university press, 1953.

[27] R. R. Murphy, V. B. M. Gandudi, and J. Adams. Robots are playing many roles in the coronavirus crisis – and offering lessons for future disasters, 2020. The Conversation.com.

[28] J. P. Nuñez Velasco, H. Farah, B. van Arem, and M. P. Hagenzieker. Studying pedestrians' crossing behavior when interacting with automated vehicles using virtual reality. *Transportation research part F: traffic psychology and behaviour*, 66:1–14, 2019.

[29] L. Patnaik and L. Umanand. Physical constraints, fundamental limits, and optimal locus of operating points for an inverted pendulum based actuated dynamic walker. *Bioinspiration & Biomimetics*, 10(6):064001, 2015.

[30] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press, 2005.

[31] A. Rasouli and J. K. Tsotsos. Autonomous vehicles that interact with pedestrians: A survey of theory and practice. *IEEE Transactions on Intelligent Transportation Systems*, 21(3):900–918, 2020.

[32] J. Rios-Martinez, A. Spalanzani, and C. Laugier. From proxemics theory to socially-aware navigation: A survey. *International Journal of Social Robotics*, 7(2):137–153, 2015.

[33] A. Thomaz, G. Hoffman, and M. Cakmak. Computational human-robot interaction. *Found. Trends Robot*, 4(2–3):105–223, dec 2016.

[34] 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.

[35] N. Tsoi, J. Connolly, E. Adéníran, A. Hansen, K. T. Pineda, T. Adamson, S. Thompson, R. Ramnauth, M. Vázquez, and B. Scassellati. Challenges deploying robots during a pandemic: An effort to fight social isolation among children. In *Proceedings of the 2021 ACM/IEEE International Conference on Human-Robot Interaction*, HRI '21, pages 234–242, New York, NY, USA, 2021. Association for Computing Machinery.

[36] UK Law Commission. Automated Vehicles: a joint preliminary consultation paper, 2019. `https://s3-eu-west-2.amazonaws.com/lawcom-prod-storage-11jsxou24uy7q/uploads/2018/11/6.5066_LC_AV-Consultation-Paper-5-November_061118_WEB-1.pdf`.