



UNIVERSITY OF LEEDS

This is a repository copy of *Comparing merging behaviors observed in naturalistic data with behaviors generated by a machine learned model*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/175332/>

Version: Accepted Version

Proceedings Paper:

Srinivasan, AR orcid.org/0000-0001-9280-7837, Hasan, M, Lin, Y-S et al. (4 more authors) (2021) Comparing merging behaviors observed in naturalistic data with behaviors generated by a machine learned model. In: 2021 IEEE International Intelligent Transportation Systems Conference (ITSC). 2021 IEEE International Intelligent Transportation Systems Conference (ITSC), 19-22 Sep 2021, Indianapolis, IN, United States. IEEE , pp. 3787-3792. ISBN 978-1-7281-9143-0

<https://doi.org/10.1109/ITSC48978.2021.9564791>

© 2021 IEEE. 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.

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/>

Comparing merging behaviors observed in naturalistic data with behaviors generated by a machine learned model

Aravinda Ramakrishnan Srinivasan^{1*}, Mohamed Hasan², Yi-Shin Lin¹, Matteo Leonetti², Jac Billington³, Richard Romano¹, and Gustav Markkula^{1*}

Abstract—There is quickly growing literature on machine-learned models that predict human driving trajectories in road traffic. These models focus their learning on low-dimensional error metrics, for example average distance between model-generated and observed trajectories. Such metrics permit relative comparison of models, but do not provide clearly interpretable information on how close to human behavior the models actually come, for example in terms of higher-level behavior phenomena that are known to be present in human driving. We study highway driving as an example scenario, and introduce metrics to quantitatively demonstrate the presence, in a naturalistic dataset, of two familiar behavioral phenomena: (1) The kinematics-dependent contest, between on-highway and on-ramp vehicles, of who passes the merging point first. (2) Courtesy lane changes away from the outermost lane, to leave space for a merging vehicle. Applying the exact same metrics to the output of a state-of-the-art machine-learned model, we show that the model is capable of reproducing the former phenomenon, but not the latter. We argue that this type of behavioral analysis provides information that is not available from conventional model-fitting metrics, and that it may be useful to analyze (and possibly fit) models also based on these types of behavioral criteria.

I. INTRODUCTION

There is an increasing presence of vehicles with autonomous capabilities on the roadways [1], [2]. As road users share the road space, *interactions* occur, in the form of situations “where the behavior of at least two road users can be interpreted as being influenced by a space-sharing conflict between the road users” [3]. For these interactions between autonomous vehicles and other road users to be natural and safe, the autonomous vehicles need to understand other road users and anticipate their behaviors, and for this reason road user behaviour has been modeled at various levels, ranging from making a prediction about when a pedestrian will cross [4], [5] to deriving driving models for different human driven vehicles [6].

From a machine learned modeling perspective, predicting other road user behaviors can be formulated as the problem of predicting their trajectories. In recent years, recurrent neural networks (RNNs), a class of deep neural networks, have been used for these prediction tasks due to their ability

to model time series data, and not least Long Short-Term Memory (LSTM) based RNN models have been used to predict human trajectories while taking into account the neighbor’s trajectories [7]. Deo and Trivedi expanded on this idea for vehicle trajectory prediction, by utilizing a convolutional layer to preserve the spatial relationship between neighboring vehicles and predicted trajectory for the vehicle of interest [8]. This convolutional social pooling (CSP) LSTM algorithm has since been used as a benchmark to measure the accuracy of newer models [9], [10]. Mozaffari et al. [11] present a comprehensive literature review of deep-learning-based prediction algorithms.

The most common metric used in the machine learning literature to support the performance of a trajectory prediction algorithm is the root mean square error (RMSE) between the prediction and the actual trajectory; either the average displacement error over the entire prediction horizon or the the final displacement error. In models that give probability over the different maneuvers, negative log likelihood (NLL) is also reported. The drawback in simplifying the performance of the model to a single quantitative value is loss of higher-level, qualitative context about the different types of behaviors actually exhibited by the models. Conventional comparisons between different machine learned models can indicate which model reproduces the human trajectories more closely, but how low RMSE or NLL values are low enough? Do the machine learned models actually reproduce the higher level behaviors exhibited by humans?

Highway merging of vehicles can be considered as a microcosm of the complex interactions that happen during everyday driving and might thus serve as a valid example scenario to test the ability of the machine learned model to successfully navigate in traffic. During highway merging, space-sharing conflicts are common between the on-ramp vehicles and vehicles on the outermost lane of the highway, and the interaction between human drivers in this type of situation has been the topic of substantial white box (non machine-learned) modeling in the past [12], [13]. Implicitly, these models suggest that when one vehicle has a kinematic lead over the other, it will tend to pass the merging point first, but in situations where it is less clear that one agent is kinematically leading, who will go first is harder to predict, and will in practice depend on a competition (or cooperation) between the drivers. Another behavior that has been studied and modeled extensively in literature is the lane change behavior of the highway vehicle and the merging vehicle [14], [15]. The lane changes were attributed to different goals

This project has received funding from UK Engineering and Physical Sciences Research Council under fellowship named COMMOTIONS - Computational Models of Traffic Interactions for Testing of Automated Vehicles - EP/S005056/1

¹ Institute for Transport Studies, University of Leeds, UK

² School of Computing, University of Leeds, UK

³ School of Psychology, University of Leeds, UK

* Corresponding authors: A.R.Srinivasan@leeds.ac.uk, G.Markkula@leeds.ac.uk




Fig. 1. CSP-LSTM network architecture block diagram.

like, highway vehicles preferring to avoid the deceleration to accommodate the merging vehicles [16], the merging vehicle forcing the lane change due to end of on-ramp or driver’s aggressive preference, or the merging vehicle finding suitable gap in the traffic to merge safely [12]. Here we specifically study the “courtesy” lane change behavior of highway vehicles to accommodate the on-ramp vehicles. As far as we are aware, neither of these two phenomena (tendency of kinematically leading agents to pass first in merging situations, and tendency of on-highway vehicles to change lane to provide space for on-ramp vehicles) have been explicitly studied in naturalistic data, let alone in the behavior of machine-learned driver models.

Thus, the primary aim of this work is to outline an approach for obtaining richer insights—compared to RMSE or NLL—about the output of machine-learned models. Specifically, we propose analysis methods designed for identifying the two above-mentioned behavioral phenomena in naturalistic data, and then apply the same methods to the behavior predictions of one of the benchmark RNN models, the CSP-LSTM of Deo and Trivedi [8]. This model was chosen due to the availability of the code [17] and the dataset [18] used.

II. METHODS

A. Machine-learned model

Deo and Trivedi [8] introduced a CSP-LSTM architecture for vehicle trajectory prediction, as shown in Fig. 1. The model takes as input $t_{CSP_p} = 3$ seconds of trajectory data for both the vehicle of interest and the surrounding vehicles. It encodes each vehicle’s trajectory with a LSTM based encoder. A convolutional layer is utilized to preserve the spatial relationship between the surrounding vehicles’ encoded trajectories. The encoded trajectory of the vehicle of interest along with the convolutional layer output for the surrounding vehicles are passed through a LSTM-based decoder layer which produces a 5 seconds trajectory prediction for the vehicle of interest. This can then be repeated for each vehicle in a given driving scene.

B. Dataset

Deo and Trivedi [8] utilized the NGSIM dataset [18] for training their CSP-LSTM network. The NGSIM dataset consists of human driven vehicle trajectories from two different highways in USA, the US101 and I80 highways. We reproduced the same model fitting regime as Deo and

Trivedi, with exactly the same 70-10-20 split of the dataset into training, validation, and test sets, respectively, and we verified that our trained network achieved the same RMSE performance as reported in [8]. In this paper, all the presented and depicted behavioral comparisons were made between the naturalistic trajectories and the machine learned models output for only the validation and the test splits of the dataset. There was a total of 1,667 unique vehicles on the US101 and 1,268 unique vehicles on the I80 in the test and validation set. Out of these, 111 and 147 vehicles merged onto the highways from the on-ramp respectively.

C. Behavioral analysis

Humans are known to rely a lot on the first-order motion information when judging collision conflicts [19]. It is also known that humans have a fundamental ability to estimate the time-to-arrival (TTA) of approaching objects [20]. For these reasons we hypothesized that we would be able to observe salient behavioral patterns by analysing only the first-order kinematics. Let, t_m be the time the merge actually happened and τ be the look back window. From the naturalistic trajectories, the time, t_m and position of each merging scenario can be extracted. Since, we are interested in understanding what contributed to the merge happening the way it had happened, the look back window is essentially the length of history before the merge happened, which we utilize to help our behavioral understanding/analysis. Since, CSP-LSTM was capable of predicting up to 5 seconds of trajectory, the history or look back window, τ was chosen from 1 second up to 5 seconds at 1 second intervals before the actual merge (t_m) happened.

For both the vehicle on the highway and the on-ramp vehicle, the interacting pair, it is straightforward to find their distance to merging point at time, $t_m - \tau$, and also compute the instantaneous velocities of the vehicle from the naturalistic driving data. Let, v_t^h, d_t^h and v_t^m, d_t^m denote the instantaneous velocity and distance to the merging point for the highway vehicle and the on-ramp merging vehicle respectively at a given time, t . We can then compute the lead time for the highway vehicle for a given look back window τ , T_τ with the formula $T_\tau = d_{t_m-\tau}^h / v_{t_m-\tau}^h - d_{t_m-\tau}^m / v_{t_m-\tau}^m$.

The behaviors of interest were analyzed for the naturalistic data and the machine-learned model with the exact same procedure utilizing the lead time for the highway vehicle. A schematic of the behavioral comparison is shown in fig. 2. The trajectories has been represented in one (lower) dimension. In reality, both the lateral and longitudinal positions of the vehicles were utilized for training the machine-learned model and the behavioral analysis. The trained CSP-LSTM model was given trajectory history of 3 seconds as inputs from appropriate time stamp for the vehicle of interest (the highway vehicle in the interacting pair) and neighboring vehicles. The output of the trained model, the 5 seconds predicted trajectory of the vehicle of interest was analyzed for the behavior exhibited (fig. 2).

1) *Bias for kinematically leading agent to pass first:* Fig. 3(a) shows a schematic example for the bias for kine-

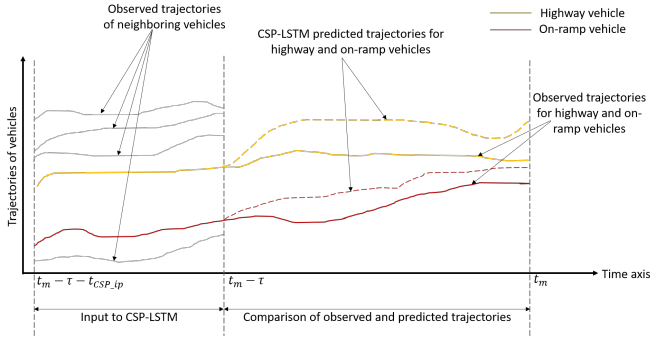


Fig. 2. Schematics of behavior comparison between naturalistic and CSP-LSTM generated trajectories.

matically leading agent to pass first behavior. Given a lead time for the highway vehicle, a positive value indicates that highway vehicle had an apparent lead over the merging on-ramp vehicle kinematically at that particular look back time and negative value indicates that the merging on-ramp vehicle had a lead. We would expect that when the absolute apparent lead $|T_\tau|$ is large enough, the kinematically leading agent would always pass the merging point first. Additionally, we would expect this pattern to break down as $|T_\tau|$ approaches zero.

2) *Changing lane to yield*: Fig. 3(b) illustrates a hypothetical lane change that happened between the merging time, t_m and the look back window, τ to accommodate the on-ramp vehicle merging into the highway. For the lane change behavior statistics, we count all lane changes which happened between the look-back window $t_m - \tau$ and the merging time t_m , since these can be considered as a potential courtesy lane changes. We expect the frequency of lane changes to increase when there is imminent space-sharing conflict, i.e., for small $|T_\tau|$.

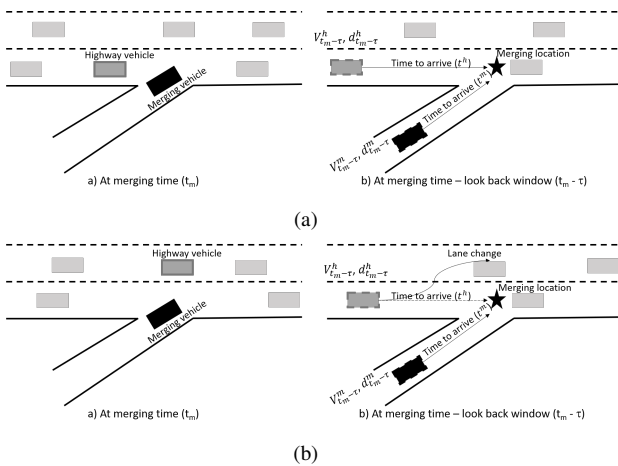


Fig. 3. A pictorial representation of behavioral phenomena a) Bias for kinematically leading agent to pass the merging point first b) Lane changing by highway vehicle to accommodate the on-ramp vehicle

III. RESULTS

A. Bias for kinematically leading agent to pass first

The frequencies of the initially kinematically leading agent passing the merging point first, as exhibited by the human driven vehicles for the US101 and I80 highways are shown in fig. 4(a) and fig. 4(c) respectively, as functions of the look-back time τ and the apparent kinematic lead T_τ for the highway vehicle at the look-back time. There are two obvious patterns: First, in line with our expectations, when either vehicle had a clear kinematic lead over the other, they ended up passing the merging point first; as soon as $|T_\tau|$ is above about 1 s for the (higher speed) US101, or above 2 s for the (lower speed) I80, the frequency of the leading agent passing first is at or close to 100%. Second, when there is imminent space-sharing conflict, that is when $|T_\tau|$ is closer to 0 seconds, there is a gradual shift to lower percentages, indicating a further level of interaction. This is especially so for larger τ , indicating—quite naturally—that the ultimate outcome of a close merging cooperation/competition may be harder to predict from snapshots further into the past.

Comparing the behaviors exhibited by the naturalistic data, fig. 4(a) and 4(c) with the CSP-LSTM model generated behaviors, fig. 4(b) and 4(d), it is clear that the CSP-LSTM model was able to produce trajectories which exhibited very similar levels of bias for the kinematically leading agent to pass the merging point first. Thus, at this level of analysis, the RMSE focused learning algorithm was enough to learn and reproduce this behavior pattern.

B. Changing lane to yield

Fig. 5 and 6 show the prevalence of lane change behavior in the naturalistic driving and the CSP-LSTM generated trajectories in US101 and I80 respectively. These observations are somewhat noisy due to the limited sample size, but is nevertheless clear that in the US101 highway data (fig. 5), when there is a space sharing conflict and there is sufficient time to successfully change lanes, $\tau \geq 3$ s there is a spike in lane change occurrences. However, the CSP-LSTM generated trajectories does not reproduce this increase in frequency of lane changes. The same data is shown in summary form in the bottom right panes of the fig. 5, indicating that human drivers in space-sharing conflict situations (lead time for highway vehicle $T_\tau \in [-1, 1]$ s), show an increased tendency to switch lanes, presumably out of courtesy to avoid the conflict, whereas the CSP-LSTM predicted trajectories do not reproduce this pattern. Also in non-space sharing conflict situations ($|T_\tau| > 1$ s), human and CSP-LSTM lane changing frequencies are not matching. With the I80 highway (fig. 6), the lane change peaks are not as clear as with the US101 highway in the naturalistic data (fig. 6). Nevertheless, in the space-sharing conflict zone, the bottom right corner in fig. 5 and fig. 6, the lane changing behaviors are similar to one another. The analyzes when done at the space-sharing conflict zone and non-space conflict zone, clearly show that the CSP-LSTM is not faithfully reproducing the lane changing behavior exhibited by the humans.

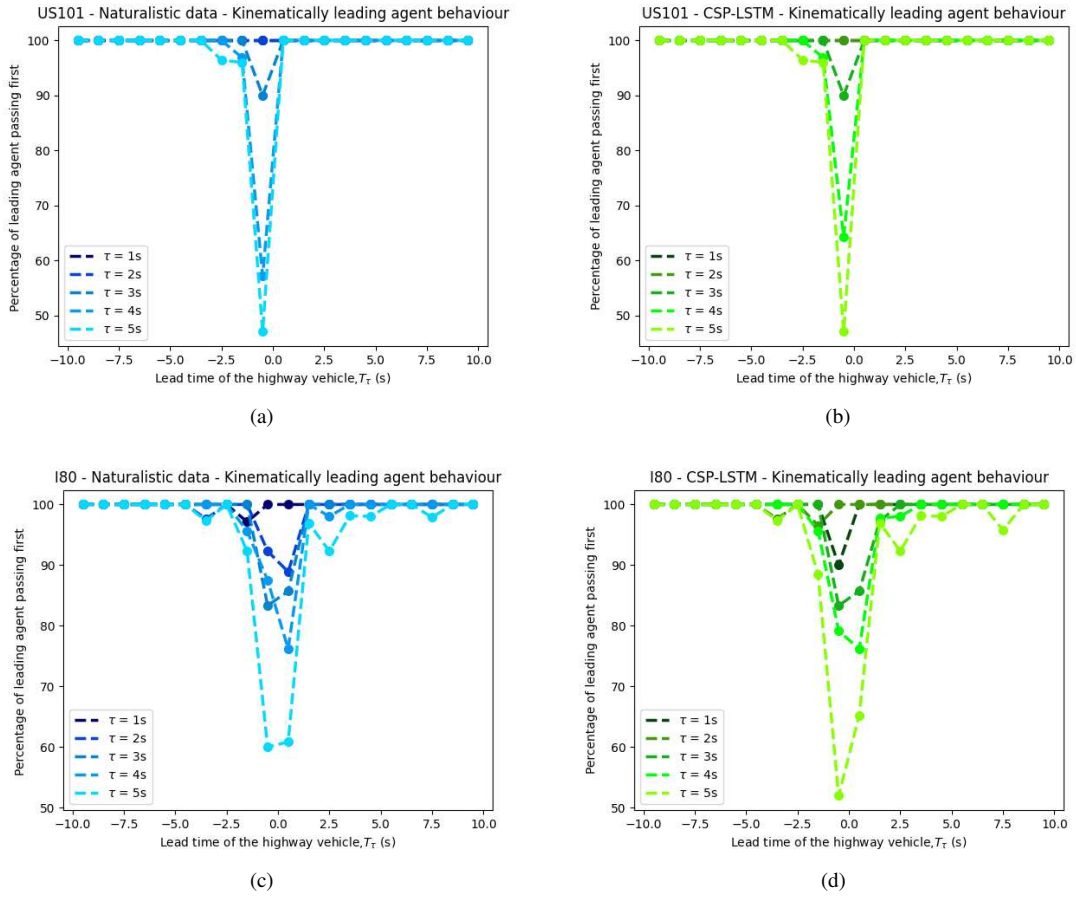


Fig. 4. Bias for kinematically leading agent to pass first behavior statistics




Fig. 5. US101 - Lane change analysis for the naturalistic driving and trajectories generated by the CSP-LSTM.

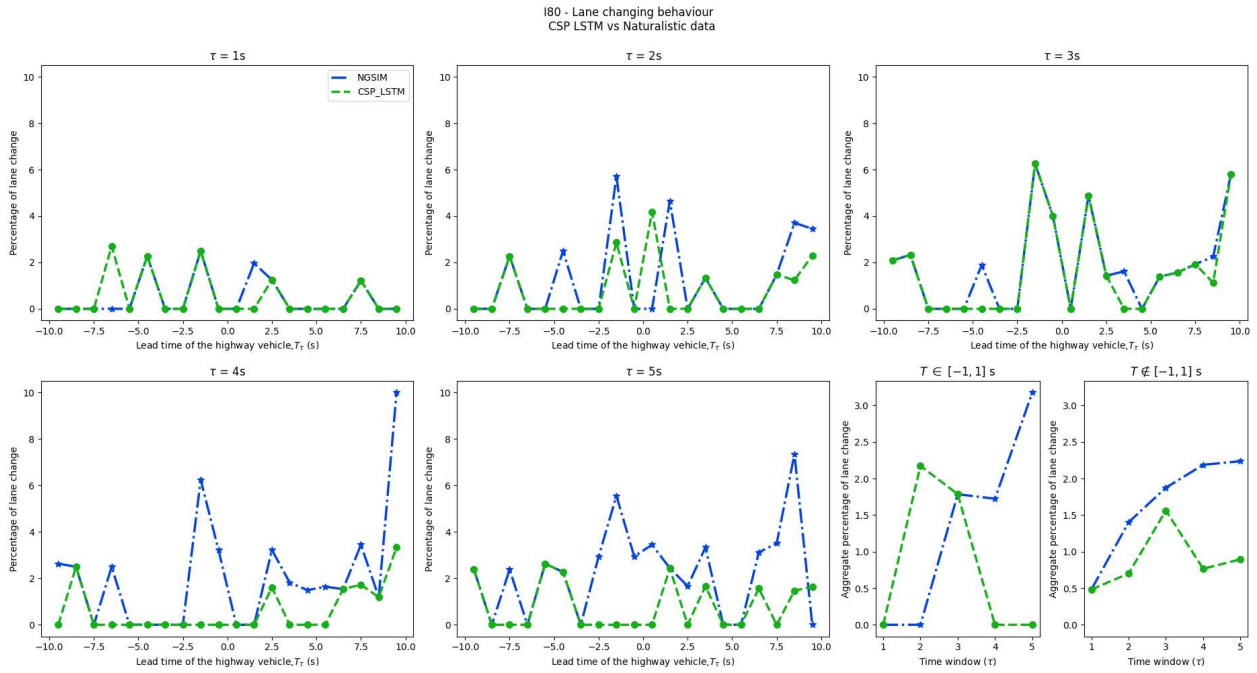


Fig. 6. I80 - Lane change analysis for the naturalistic driving and trajectories generated by the CSP-LSTM.




Fig. 7. Naturalistic and CSP-LSTM generated vehicle trajectories in potential lane changing scenarios for the vehicle on the outermost lane of the highway to accommodate the merging vehicle.

IV. DISCUSSION AND CONCLUSION

This paper presented a new, arguably more human-centric, way to analyze the capabilities of machine-learned models which are used in the domain of autonomous driving for predicting vehicle trajectories. To illustrate the proposed approach, two example behavioral phenomena in highway merging were targeted, analysis methods were developed to demonstrate their presence in naturalistic data, and subsequently also applied to the model's predictions. The results showed that the machine-learned model was able to reproduce well a phenomenon whereby the kinematically leading agent passes the merging point first when the apparent kinematic lead is large ($> 1 - 2$ s), up to 5 seconds in advance of the merge, and where this predictability also deteriorates for less clear kinematic leads. Thus, the simplification of the road user behavior learning to a pure trajectory learning task seems to be justified for replicating the kinematically leading agent bias to pass the merging point first as observed in the naturalistic driving data.

The other targeted phenomenon was courtesy lane changing behavior by the highway vehicles, to facilitate the merging of on-ramp vehicles. In short, the machine learned model was not able to reproduce this phenomenon well. To further understand the human and model lane changing behavior, the trajectories of the vehicles in both highways from the naturalistic and CSP-LSTM generated data are shown in fig. 7. It is clear from this figure that the limitation of the CSP-LSTM model does not lie in a general inability of predicting lane changing behavior; it clearly does predict some of the observed naturalistic lane changes. However, it does in general underpredict their frequency, and, crucially, it is clear from the bottom right panes of especially fig. 5 and 6, the model does not capture the context-sensitive increased tendency of human drivers to change lanes when there is an apparent space-sharing conflict with a merging vehicle.

Overall, the findings presented here demonstrate how a richer analysis of human and model-predicted behavior can provide a better understanding of the capabilities of machine-learned models. It is clear that the CSP-LSTM model is capable of capturing some advanced behavioral phenomena in impressive detail, yet is unable to capture other phenomena. Notably, neither of these insights are accessible from conventional performance metrics such as RMSE or NLL. Our results also open for many interesting future research opportunities: Our overall analysis approach can be generalized to a wider range of salient behavioral phenomena, across a wider range of interaction scenarios [3]. The added insights into the behavioral capabilities of the machine-learned models may be leveraged both in algorithms making use of the models, and to develop improved learning targets for the models, to help ensure that future models can be more behaviorally competent.

REFERENCES

[1] R. Ono, W. Ike, and Y. Fukaya, "Pre-collision system for toyota safety sense," SAE Technical Paper, Tech. Rep., 2016.

- [2] M. Dikmen and C. M. Burns, "Autonomous driving in the real world: Experiences with tesla autopilot and summon," in *Proceedings of the 8th international conference on automotive user interfaces and interactive vehicular applications*, 2016, pp. 225–228.
- [3] G. Markkula, R. Madigan, D. Nathanael, E. Portouli, Y. M. Lee, A. Dietrich, J. Billington, A. Schieben, and N. Merat, "Defining interactions: a conceptual framework for understanding interactive behaviour in human and automated road traffic," *Theoretical Issues in Ergonomics Science*, vol. 21, no. 6, pp. 728–752, 2020. [Online]. Available: <https://doi.org/10.1080/1463922X.2020.1736686>
- [4] F. Camara, N. Bellotto, S. Cosar, D. Nathanael, M. Althoff, J. Wu, J. Ruenz, A. Dietrich, and C. Fox, "Pedestrian models for autonomous driving part i: low-level models, from sensing to tracking," *IEEE Transactions on Intelligent Transportation Systems*, 2020.
- [5] F. Camara, N. Bellotto, S. Cosar, F. Weber, D. Nathanael, M. Althoff, J. Wu, J. Ruenz, A. Dietrich, G. Markkula *et al.*, "Pedestrian models for autonomous driving part ii: high-level models of human behavior," *IEEE Transactions on Intelligent Transportation Systems*, 2020.
- [6] L. Claussmann, M. Revilloud, D. Gruyer, and S. Glaser, "A review of motion planning for highway autonomous driving," *IEEE Transactions on Intelligent Transportation Systems*, vol. 21, no. 5, pp. 1826–1848, 2019.
- [7] A. Alahi, K. Goel, V. Ramanathan, A. Robicquet, L. Fei-Fei, and S. Savarese, "Social lstm: Human trajectory prediction in crowded spaces," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [8] N. Deo and M. M. Trivedi, "Convolutional social pooling for vehicle trajectory prediction," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2018, pp. 1468–1476.
- [9] Y. Ma, X. Zhu, S. Zhang, R. Yang, W. Wang, and D. Manocha, "Trafficpredict: Trajectory prediction for heterogeneous traffic-agents," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, no. 01, pp. 6120–6127, Jul. 2019. [Online]. Available: <https://ojs.aaai.org/index.php/AAAI/article/view/4569>
- [10] X. Li, X. Ying, and M. C. Chuah, "Grip: Graph-based interaction-aware trajectory prediction," in *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, 2019, pp. 3960–3966.
- [11] S. Mozaffari, O. Y. Al-Jarrah, M. Dianati, P. Jennings, and A. Mouzakis, "Deep learning-based vehicle behavior prediction for autonomous driving applications: A review," *IEEE Transactions on Intelligent Transportation Systems*, 2020.
- [12] C. F. Choudhury, V. Ramanujam, and M. E. Ben-Akiva, "Modeling Acceleration Decisions for Freeway Merges," *Transportation Research Record: Journal of the Transportation Research Board*, vol. 2124, no. 1, pp. 45–57, Jan. 2009. [Online]. Available: <http://journals.sagepub.com/doi/10.3141/2124-05>
- [13] K. Kang and H. A. Rakha, "Game theoretical approach to model decision making for merging maneuvers at freeway on-ramps," *Transportation Research Record*, vol. 2623, no. 1, pp. 19–28, 2017.
- [14] H. Yu, H. E. Tseng, and R. Langari, "A human-like game theory-based controller for automatic lane changing," *Transportation Research Part C: Emerging Technologies*, vol. 88, pp. 140–158, 2018.
- [15] Z. Zheng, "Recent developments and research needs in modeling lane changing," *Transportation research part B: methodological*, vol. 60, pp. 16–32, 2014.
- [16] A. Kondyli and L. Elefteriadou, "Driver behavior at freeway-ramp merging areas: focus group findings," *Transportation research record*, vol. 2124, no. 1, pp. 157–166, 2009.
- [17] N. Deo. Csp-lstm github code repository. Accessed: 2021-03-21. [Online]. Available: <https://github.com/nachiket92/conv-social-pooling>
- [18] U. J. Colyar. Next generation simulation (ngsim) vehicle trajectories and supporting data. Accessed: 2021-03-21. [Online]. Available: <https://data.transportation.gov/Automobiles/Next-Generation-Simulation-NGSIM-Vehicle-Trajectory/8ect-6jgj>
- [19] G. Markkula, J. Engström, J. Lodin, J. Bärgrman, and T. Victor, "A farewell to brake reaction times? kinematics-dependent brake response in naturalistic rear-end emergencies," *Accident Analysis & Prevention*, vol. 95, pp. 209–226, 2016.
- [20] D. N. Lee, "A theory of visual control of braking based on information about time-to-collision," *Perception*, vol. 5, no. 4, pp. 437–459, 1976.