

XXV International Conference Living and Walking in Cities - New scenarios for safe mobility in urban areas (LWC 2021), 9-10 September 2021, Brescia, Italy

Young drivers' pedestrian anti-collision braking operation data modelling for ADAS development

Matteo Miani^{a,1}, Matteo Dunnhofer^{b,1}, Christian Micheloni^b,

Andrea Marini^{c,d}, Nicola Baldo^{a,*}

^a*Polytechnic Department of Engineering and Architecture (DPIA), University of Udine, Via del Cotonificio 114, 33100 Udine, Italy*

^b*Department of Mathematics, Computer Science and Physics (DMIF), University of Udine, Via delle Scienze 206, 33100 Udine, Italy*

^c*Department of Languages, Literatures, Communication, Education and Society (DILL), University of Udine, Via Margreth 3, 33100 Udine, Italy*

^d*Claudiana—Landesfachhochschule für Gesundheitsberufe, I-39100 Bolzano, Italy*

Abstract

Smart cities and smart mobility come from intelligent systems designed by humans. Artificial Intelligence (AI) is contributing significantly to the development of these systems, and the automotive industry is the most prominent example of "smart" technology entering the market: there are Advanced Driver Assistance System (ADAS), Radar/LIDAR detection units and camera-based Computer Vision systems that can assess driving conditions. Actually, these technologies have become consumer goods and services in mass-produced vehicles to provide human drivers with tools for a more comfortable and safer driving. Nevertheless, they need to be further improved for progress in the transition to fully automated driving or simply to increase vehicle automation levels. To this end, it becomes imperative to accurately predict driver's decisions, model human driving behaviors, and introduce more accurate risk assessment metrics. This paper presents a system that can learn to predict the future braking behavior of a driver in a typically urban vehicle-pedestrian conflict, i.e., when a pedestrian enters a zebra crossing from the curb and a vehicle is approaching. The algorithm proposes a sequential prediction of relevant operational indicators that continuously describe the encounter process. A car driving simulator was used to collect reliable data on braking behaviours of a cohort of 68 licensed university students, who faced the same urban scenario. The vehicle speed, steering wheel angle, and pedal activity were recorded as the participants approached the crosswalk, along with the azimuth angle of the pedestrian and the relative longitudinal distance between the vehicle and the pedestrian: the proposed system employs the vehicle information as human driving decisions and the pedestrian information as explanatory variables of the environmental state. In fact, the pedestrian's polar coordinates are usually calculated by an on-board millimeter-wave radar which is typically used to perceive the environment around a vehicle. All mentioned information is represented in the form of time series data and is used to train a recurrent neural network in a supervised machine learning process. The main purpose of this research is to define a system of behavioral profiles in non-collision conditions

* Corresponding author. Tel.: +39-0432-558-745.

E-mail address: nicola.baldo@uniud.it

¹ M.M. and M.D. share equal contribution.

that could be used for enhancing the existing intelligent driving systems, e.g., to reduce the number of warnings when the driver is not on a collision course with a pedestrian. Preliminary experiments reveal the feasibility of the proposed system.

© 2022 The Authors. Published by ELSEVIER B.V.

This is an open access article under the CC BY-NC-ND license (<https://creativecommons.org/licenses/by-nc-nd/4.0>)

Peer-review under responsibility of the scientific committee of the Living and Walking in Cities

Keywords: ADAS; Traffic safety; Driver behaviour; Gated Recurrent Units; Driving simulator.

1. Introduction

Technological development of car-sensing devices, such as radars, cameras, and sensors, has allowed scientists to build increasingly reliable architectures, called Advanced Driving Assistant Systems (ADAS), that can improve driving safety and efficiency. In fact, the modern cars' capabilities to detect specific objects in the surrounding road environment and to represent a high number of traffic situations have been enhanced to the point that SAE Level 2 autonomous driving (i.e., the driving automation level provided by steering and brake/acceleration support to the driver, as defined by the Society of Automotive Engineers) is an established reality. However, the transition to fully autonomous driving requires an additional innovation: the capability to understand and predict the intentions of other traffic participants up to a few seconds in the future. Relating this prediction to current driving parameters, the autonomous or semi-autonomous vehicle can make appropriate driving decisions in advance, such as adjusting its trajectory or forgoing a specific action (e.g., overtaking), to avoid a predicted future collision with another moving road user. Nevertheless, the prediction of future driver behavior (or trajectory) based on the current ego-vehicle dynamic state and scene representation is also a concept that draws the attention of scientists. Such a prediction can be exploited to adapt assistance systems to the driver's intentions, i.e., to determine on the current situation representation whether (and when) to initiate or abort an intervention (avoiding unnecessary driving interference). In fact, most current ADAS systems do not have a medium-term predictive capability of the driver's intentions (who is still in charge of decision-making), but act reactively in high-risk situations, such as the emergency braking system. However, to reach this goal, there is a need to define a system of expected behavioral profiles in ordinary traffic situations to be compared with the system-generated predictions.

This study represents a contribution to the above-mentioned topic and presents a system, based on machine learning techniques, capable of predicting ahead in time (on different time scales) the severity level (intended as the potential of an elementary traffic event to become an accident) associated with a typical inner-city traffic scenario: the encounter between the ego-vehicle and a pedestrian who enters a crosswalk from the curb. This prediction depends on the current vehicle dynamic state, accelerator and brake pedal information, as well as external scene properties. In particular, it is the authors' intention to provide a proof-of-concept of both the learning method and the surrogate safety metrics for studying the considered scenario. Actually, although this study focuses on developing a system of expected behavioral profiles in a specific traffic scene, the exposed methodologies can be easily scaled to a higher number of inner-city traffic situations.

1.1. Related works

Machine learning techniques are emerging in the area of ADAS as the main approach to motion prediction, and a review of the relevant literature can be found in Meiring and Myburgh (2015). Specifically, existing techniques for learning driver or other road user behavior sequences from the set of features acquired by the vehicle sensor system can be divided between 2 methods: classification and regression.

Classification problems concern the identification of future movement intentions, also called "behavior primitives" to segment complex driving behavior into a sequence of basic elements, such as lane keeping, left/right lane change or left/right turn, go straight (when the behavior of vehicles surrounding the ego-vehicle is of interest), or speed maintenance, braking and stopping (when the object is the driver behavior). In the latter context, the approach of Ortiz et al. (2011) showed that multi-layer perceptron models (MLPs), can produce excellent predictions of future behavior primitives in a simple inner-city scenario (i.e., approaching a traffic light), over very long-time horizons (up to 6s), using a current "low-dimensional" representation of the scene. In fact, these authors derived object-level features,

focusing on the current ego-vehicle state (its speed) and the scene properties (distance and status of the nearby traffic light). Recently, Khairdoost et al. (2020) implemented a deep learning system that can anticipate (by 3.6 seconds on average) driver maneuvers (left/right lane change, left/right turn and go straight), exploiting the driver's gaze and head position as well as vehicle dynamics data.

Differently, regression problems are concerned with predicting the future positions of cars (Alché and de La Fortelle, 2017), cyclists (Zernetsch et al., 2016) and pedestrians (Alahi et al., 2016) surrounding the ego-vehicle, by a general understanding of their movement dynamics. In this context, learning models must learn from inputs that are time series, i.e., the current and past positions of the traffic participants, and produce outputs that are future sequences (Lipton et al., 2015). In recent years, Long Short-Term Memory (LSTM) and Gated Recurrent Units (GRU) architectures, variants of the more general Recurrent Neural Networks (RNNs), have shown excellent efficiency in time series prediction tasks: these models can capture the sequences' dynamics by extracting relevant information from an arbitrarily long context window and retaining a state of that information (Hochreiter and Schmidhuber, 1997). Among recent applications in problems relevant to autonomous vehicle movement within urban setting, Huang et al. (2019) encoded temporal and spatial interactions between pedestrians in a crowded space, by combining an LSTM and a Graph Attention Network (GAT), to obtain "socially" plausible trajectories.

As regression methods have gained prominence in the study of the interaction between the ego-vehicle trajectory and that of other traffic participants (whereas classification approaches merely provide rough information about the intentions of the scene participants), this paper focuses on predicting future severity levels associated with driver-pedestrian encounters by exploiting GRU architectures and surrogate safety assessment metrics. To the best of our knowledge, the existing literature has attempted to evaluate trajectories of individual road users and the presence of a collision course between them, neglecting the operational severity measures of their interaction. In fact, any traffic encounter can be analyzed as an interaction process with the potential to develop into an accident (Laureshyn et al., 2010). So, the presented approach differs from previous studies in its capability to describe continuously (and ahead in time) the safety relationship between two traffic participants, whether they are on a collision course or leaving it. In particular, our contribution is aimed at studying situations in which the driver has avoided collision with the pedestrian. The proposed method has been developed on data collected by means of a car driving simulator.

2. Risk Assessment Model

The driver-pedestrian encounter, i.e., the simultaneous arrival of a driver and a pedestrian at the crosswalk, is a traffic event characterized by a continuous interaction over time and space between the two road users (Várhelyi, 1998). In fact, the pedestrian's decision to enter the zebra crossing depends on the perceived speed and distance of the approaching vehicle; concurrently, the driver evaluates whether to grant or deny the priority to the pedestrian, based on the estimated arrival time at the crosswalk. Since the two traffic participants may enter a collision course during the encounter process, such a conflict has the potential to end in a collision. The latter would occur if, for example, the driver's attention levels, his/her ability to control the vehicle, or the vehicle's dynamic state were not adequate for a safe stopping behaviour (Baldo et al., 2020). Therefore, since the goal of current research studies (Sec. 1.1) is the development of an efficient Autonomous Emergency Braking Pedestrian (AEB-P) system for assisted or autonomous driving, there is a need to define a Risk Assessment Model (RAM) that can objectively and quickly capture the severity of the encounter process. In fact, a RAM that can fully understand the relationships between behavior and risk is essential to adequately judge the system's intervention timing (Yang et al., 2019).

In recent years, the application of traffic safety indicators, also referred to as Traffic Conflict Techniques (TCTs), has been very successful as a proactive surrogate approach (as it is complementary to accident statistical analysis) in the area of traffic event safety assessment due to its efficiency and short analysis time (Zheng et al., 2014). Specifically, Laureshyn et al. (2010) identified and developed a set of safety indicators to continuously describe the severity level of the encounter process between two road users and, thus, to relate "individual interactions to the general safety situation" of the event. In the following, we briefly present the two fundamental indicators for evaluating the encounter between a vehicle and a pedestrian on a crosswalk, the Time-To-Collision (TTC) and Time Advantage (T_{ADV}), whereas for the calculation procedure please refer to Appendix A of the study by Laureshyn et al. (2010).

The TTC, at an instant t , is a continuous parameter that represents the time it would take for two users on a collision course to collide if they maintained their current trajectory and relative speed. Specifically, among all the pedestrian-

vehicle front end contact points in a collision, the one leading to the lowest TTC value (TTC_{MIN}) should be selected. In principle, the lower the TTC_{MIN} , the higher the risk of a collision and vice versa. On the other hand, the T_{ADV} is the temporal extension of the Post-Encroachment Time (PET) concept, i.e., it represents at each instant t the time that would elapse between the passage of the first and the second user through the same conflict zone if both maintained the current trajectory and relative speed. In other words, this indicator expresses the potential to avoid an accident, although the two road users are on a collision course: values above 2-3 seconds indicate that a user has a temporal advantage over his opponent in a competition over the same spatial zone and is likely to pass first.

3. Driving information acquisition

Since this study focuses on the development of predictive models of the expected severity level in a common vehicle-pedestrian encounter process, there is a need to acquire vehicle and pedestrian information during the stopping maneuver of several drivers in such an inner-city scenario and in sufficient numbers to successfully train GRU networks. In addition, these data must allow for the proper evaluation of interaction safety indicators (especially during the online system application) and be easily acquired by on-board sensors (i.e., cameras, pedal potentiometers, IMU, GNSS and millimeter-wave radars). In particular, the analysis of the relevant literature (Sec. 1.1 and 2), allowed us to identify three main groups of features for studying the problem in question: the actuators and steering wheel states, the information about the car dynamics, the pedestrian speed and direction vector. In order to collect this information while keeping drivers safe, a cohort of 68 licensed university students was recruited to participate in a driving simulation experiment at the Roads Laboratory of the Polytechnic Department of Engineering and Architecture (DPIA) of the University of Udine. In the Lab's AutoSim 1000-M driving simulator (please see Baldo et al., 2020), it was possible to record a data set for GRU models development by observing test participants' behavior toward a planned traffic encounter: a child entering a crosswalk from the curb.

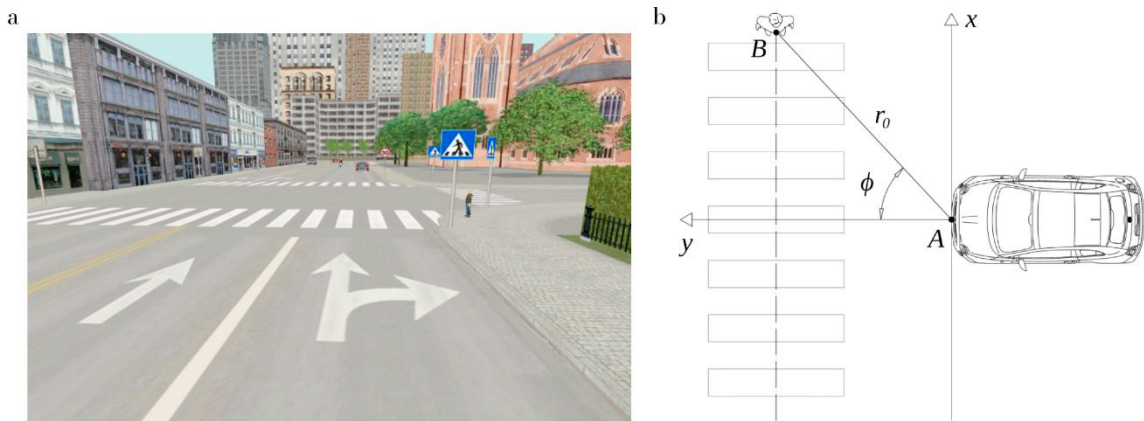


Fig. 1. (a) Frontal view of the simulated scenario and (b) ego-vehicle coordinate system.

The simulated scene (Fig. 1a) was set in a typical urban environment on a course that could take about 15 minutes to be completed, considering the 50 km/h speed limit. On the urban course, participants experienced many traffic light intersections, tight turns (90°), short straight streets, and occasional crossings of pedestrians, some of which occurred outside the crosswalks. The encounter that was used to record data (and compute surrogate safety indicators for each participant in an offline fashion, following the procedure described in Sec. 2) was set on a four-lane road (two lanes in each direction), placing traffic signs and markings that met European standards (Fig. 1a). The scenario was designed as follows: 1) the participant arrives at a red traffic light approximately 200 m from the crosswalk, on which he/she has a clear view, 2) as the participant starts driving, the pedestrian (initially hidden) walks at a 90° angle toward the road and then stops at the edge of the curb, 3) at the moment the vehicle, based on its current speed, is about 3s from the crosswalk, the pedestrian enters the zebra crossing and maintains a speed of 1.4 m/s. These conditions, consistent

with relevant literature (Bella et al., 2018), require the participant to stop, giving priority to the pedestrian. In this way, the data collected refers to stopping maneuvers that avoided a collision with the pedestrian.

The students recruited for the experiment were between the ages of 20 and 30, had a valid driver's license, and had driven at least 5,000 km behind the wheel in the past year. All participants were properly trained in the use of the simulator actuators before the experimental driving and were able to test their abilities on a simulated suburban course that took approximately 5 minutes. However, 5 of 68 participants were excluded from the analyses, as they could not complete the experiment due to excessive discomfort. The remaining 63 drivers were screened to identify a sample of males and females roughly balanced on age and non-verbal intellectual quotient (IQ). The cohort that was selected for GRU models development included 17 females and 31 males, with a mean age of 23.2 (SD 2.88) and 23.9 (SD 1.93) years and a mean IQ level of 33.1 (SD 2.56) and 34.0 (SD 1.58), respectively. Finally, it is worth pointing out that participation in the experiment was voluntary, there was no monetary reward, and all participants gave informed consent after being instructed on the simulated driving test procedure. Conversely, they were not informed about the objectives of the research.

4. Features, outputs and prediction strategy

Formally, the severity prediction system must perform a regression between the situation representation at the current time t , described by a set of sensor-observed features, and the expected severity level of the vehicle-pedestrian encounter, defined by surrogate safety indicators, for times $t+1s$, $t+2s$, and $t+3s$. In practice, a learning algorithm is trained to accurately predict the TTC and T_{ADV} at instants $t-2s$, $t-1s$, and t using the vehicle sensing system observations available at instant $t-3s$. Once the expected convergence between predicted outputs and ground-truth targets is reached, the trained model represents the desired medium-term prediction system and can be used to make forward-in-time predictions based on the current scene representation. In particular, it is assumed that the input features can all be acquired simultaneously and at a regular time interval given by the lowest sampling frequency among those of the involved sensors, since the goal is to perform real-time prediction on the running vehicle. However, considering that the current study is based on a driving simulator experiment, the calculation of surrogate safety indicators for each participant and the training of the learning algorithm were both performed in an offline fashion, at the end of the driving experiments.

The coordinate system established for the vehicle sensing system is shown in Fig. 1b. The following features form the 11 components of the learning model input vector and describe the stopping maneuvers for each participant in the experimental data set (Sec. 3): travel time, accelerator- and brake-pedal positions, steering angle, ego-vehicle's longitudinal and lateral velocity (v_{AX} ; v_{AY}), yaw rotation, pedestrian's velocity (v_{BX} ; v_{BY}) and direction vectors (r_0 ; ϕ), i.e., the azimuth angle ϕ of the pedestrian with respect to the travel direction and the distance r_0 between the ego-vehicle (point A or C) and the pedestrian (point B, please see Fig. 1b). Specifically, v_{BX} , v_{BY} , r_0 and ϕ are the pedestrian state data acquired by presumed millimeter-wave radars mounted on the vehicle front end (points A and C, Fig. 1b). In fact, while the recognition of the pedestrian presence in the scene is usually done using robust frame classification algorithms of vehicle camera videos, the status of the detected pedestrian is generally acquired by a radar: in this study, we assumed a data acquisition equipment consisting of a long-distance millimeter-wave radar installed at the center of the vehicle front end and a mid-range radar on either side of the vehicle front, with maximum detection distances of 100 m and 50 m and azimuth angles of $\pm 10^\circ$ and $\pm 45^\circ$, respectively, to ensure the optimal scene coverage (Yang et al. 2019). Therefore, the time sequences of the participants' stopping maneuvers begin with the detection of the pedestrian's status from the long-distance radar, approximately 100 m from the crosswalk, assuming the simultaneous detection of the pedestrian's presence by the ego-vehicle classification system. Furthermore, although the sampling rate of these systems (e.g., the current 77 GHz band millimeter-wave radar) is 20 Hz, we assumed that the encounter process was recorded at 10 Hz to account for limitations of other on-board equipment and raw data processing times.

Regarding the model output vector, we decided to split the severity level prediction problem into single-output learning tasks: TTC and T_{ADV} indicators, associated with each participant, were the learning target of two separate GRU models. Finally, before being inputted to the GRU, each feature was standardized, i.e., was subtracted by its minimum value and divided by the distance between its minimum and maximum value computed on the training set, to remain in an acceptable range with respect to the activation functions.

5. GRU networks

Due to their extended success in many time-series applications, we employed RNNs to implement our learning system. In particular, we selected the GRU architecture because of its increased robustness with respect to vanilla RNNs, and of its expressive power that matches the more sophisticated Long Short-Term Memory (LSTM) network but achieved with less training effort. Based on the input feature vector input x_t and previous output h_{t-1} , each GRU learning neuron performs different operations using so-called “gate” operations. The “update” gate decides the amount of past information to be forwarded to the future, while the “forget” gate focuses on which part of the past information to forget. These particular features allow a network to produce more meaningful gradients during the learning phase, ultimately making them acquire enhanced long-term relations between features. To additionally improve the quality and abstractness of such feature relations, the learning model can be organized as stacked layers of GRUs.

In our setting, we used 3 layers of GRUs each with 64 neurons. After those, a dense output layer with as many neurons as the outputs is applied to predict the target values. Overall, the role of the GRU layers is to abstract a meaningful representation that summarizes the driver’s maneuver up to time step t . Such higher-level features are in turn exploited by the output layer that produces the predicted future TTC and T_{ADV} states.

We used the PyTorch framework (Paszke et al., 2019) to implement the described architecture and its learning. The full dataset was split subject-wise in training and test sets using a 0.7 ratio, resulting in 33 subjects for training and 15 for testing. The model has been trained to minimize the RMSE between its predictions and the ground-truth TTC or T_{ADV} targets. The learning procedure has been conducted for 10’000 epochs using the Adam optimizer (Kingma and Ba, 2015) with a learning rate of 10^{-3} . After each epoch, the model was executed on the test set to assess its generalization capabilities to new maneuvers. The model obtaining the lowest RMSE score on such tests, hence the most general one, was retained as the final learned model.

6. Results

The aim of this paper was to train GRU neural networks to learn medium-term (up to 3s) relations between surrogate safety indicators and some relevant features of an encounter situation, in order to make predictions. In this section, the generalization capabilities of the trained models are evaluated for each time horizon (i.e., 1s, 2s, 3s ahead) by first computing both the Mean Absolute Error (MAE) and the Root Mean Square Error (RMSE) between the network prediction and the actual value of the safety indicators that characterized the 15 test set drivers stopping maneuvers, and then averaging the obtained results among these drivers. In what follows, we refer by the acronyms GRU_{TTC} and GRU_{TADV} to the GRU models predicting TTC and T_{ADV} , respectively. Table 1 presents the average MAE and RMSE values divided by model type and prediction time horizon, along with the 90th percentile of MAE. Differently, Fig. 2 presents the empirical cumulative distribution probability (ECDP) of the absolute prediction error (AE) over the whole test set for both GRU_{TTC} (Fig. 2a) and GRU_{TADV} (Fig. 2b): the 90th percentile of the AE is shown in Table 1.

Table 1. Model performance evaluation.

Prediction Horizon	GRU_{TTC} Model Evaluation				GRU_{TADV} Model Evaluation			
	90%th AE	MAE	90%th MAE	RMSE	90%th AE	MAE	90%th MAE	RMSE
1s	0.4712s	0.1352s	0.2446s	0.2477s	0.7251s	0.3410s	0.5440s	0.5107s
2s	0.6166s	0.2558s	0.4006s	0.4473s	1.0533s	0.4522s	0.8979s	0.7408s
3s	0.8688s	0.3805s	0.4478s	0.5844s	1.0367s	0.4587s	0.8884s	0.8266s

As expected, the quality of TTC and T_{ADV} predictions is higher, both in terms of MAE and RMSE (Table 1), when the time horizon is the shortest, but the worsening of both metrics as the time step grows increases is characterized by a deviation that is always less than a quarter of a second: the maximum deviation, equal to 0.2301s, was measured on the RMSE of the GRU_{TADV} model when moving from the 1s to the 2s prediction time-frame (Table 1, last column), whereas the prediction accuracy did not seem to change significantly moving from 2s to 3s. These results confirm that there is a strong correlation between selected input features and the shortest time-frame prediction, but such correlations become less evident for discriminating the expected safety indicators values at longer horizons. This is

particularly the case for the GRU_{TADV} model which provided less accurate (although very positive) predictions than the GRU_{TTC} model, as can be easily inferred by comparing the 90th percentile values of the MAE (columns 4 and 8 of Table 1) between the two models: it means that the T_{ADV} is more difficult to learn than the TTC, probably because the T_{ADV} varies more rapidly than the TTC when car and pedestrian enter a collision course (Laureshyn et al., 2010). However, overall, the obtained results are satisfactory, and the mean absolute error was on each time scale and for each safety indicator less than half a second (i.e., the magnitude order is that of 500 milliseconds), confirming the success of the proposed approach.

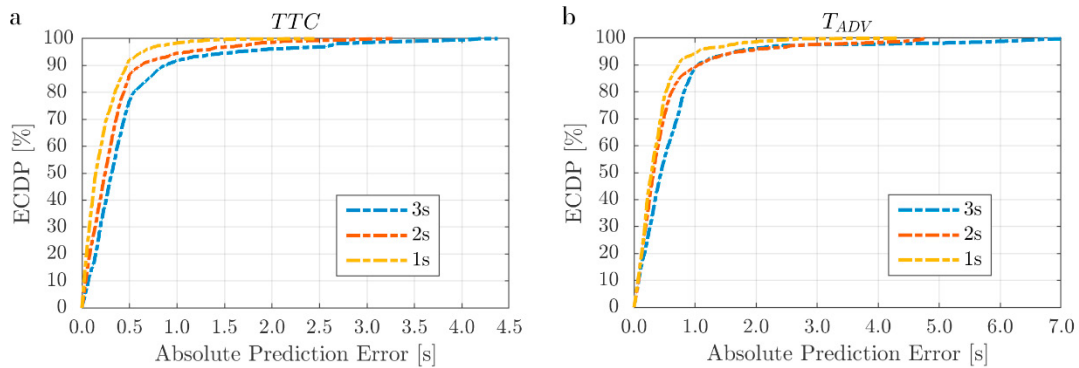


Fig. 2. Cumulative distribution of absolute prediction errors on the test set's TTC (a) and T_{ADV} (b) indicators.

Fig. 3 illustrates a feasible application of the proposed system: employing the three interaction severity classes proposed by Laureshyn et al. (2010) for the encounter between a car and a pedestrian, the TTC and T_{ADV} predictions were exploited to evaluate (for a participant randomly sampled from the test set) the accuracy of a severity level classification system.

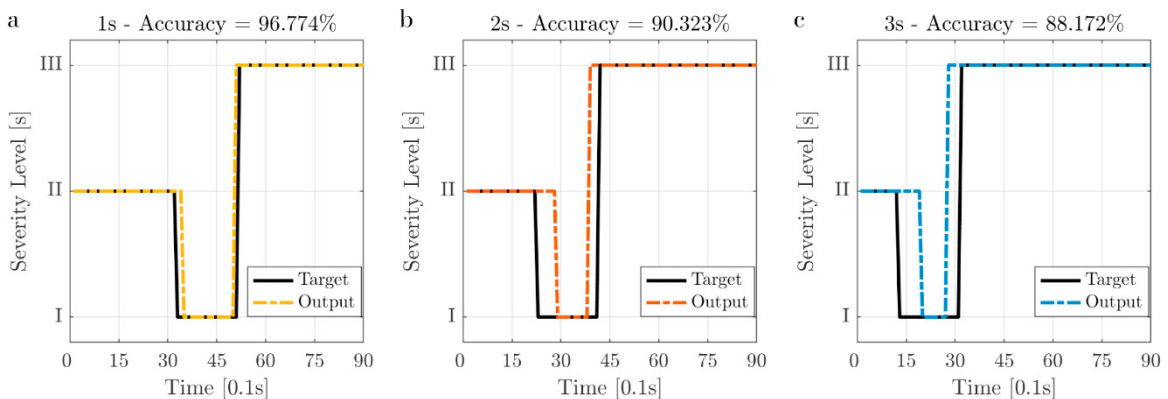


Fig. 3. Example of the encounter's severity level classification for a participant in the test set, for each time horizon: (a) 1s (b) 2s (c) 3s.

The observed accuracy exceeds 88% on each prediction timescale but, as such timescale increases, GRU predictions delay (anticipate) the expected severity level in the first (in the last) phase of the encounter process. Alché and de La Fortelle (2017) had already observed this phenomenon in predicting surrounding vehicle trajectories and blamed it on the lack of learning metrics that could capture the stochasticity and confidence levels of the prediction problem.

7. Conclusions

In this study, it was shown that it is possible for a Deep Learning system, based on GRU architectures, to predict ahead in time changes in the severity level of the encounter process between the ego-vehicle and a pedestrian on the crosswalk. The proposed method, which relies on data acquired with an advanced driving simulator, exploits a gradient-descent optimizer to learn how to describe the interaction process from individual driving characteristics. Although the prediction quality decreases as the time horizon increases, the trained models provide satisfactory and promising results for both managing the vehicle-pedestrian encounter in autonomous driving systems and improving current ADAS systems. In fact, the proposed predicting system could be applied to recognize an anomalous or hazardous interaction, to anticipate a braking maneuver, as well as to enhance vehicle deceleration, with the aim of pursuing an improvement in vehicle and pedestrian safety in the inner-city traffic.

Future research perspectives opened by the current study include: (1) generalizing the proposed GRU models to different encounter scenarios; (2) studying selected features to identify a low dimensional scene representation suitable for prediction purposes; (3) addressing the encounter severity level prediction as a behavior classification problem to compare the prediction quality against that of the proposed system; (4) modifying the system for online learning and operation; (5) validating the prediction system with data acquired during real vehicle-pedestrian encounters on inner-city roads; (6) designing metrics better suited to the stochastic problem at hand that consider probability distributions and confidence levels of prediction errors. Anyway, this research represents a valuable contribution to the study, understanding and knowledge of the interaction processes between road users from the point of view of traffic safety, especially in urban environment.

References

- Alahi, A., Goel, K., Ramanathan, V., Robicquet, A., Fei-Fei, L., Savarese, S., 2016. Social LSTM: Human trajectory prediction in crowded spaces, 2016 IEEE Conference on Computer Vision and Pattern Recognition. Las Vegas, NV, USA, pp. 961–971.
- Alché, F., de La Fortelle, A., 2017. An LSTM network for highway trajectory prediction, 2017 IEEE 20th International Conference on Intelligent Transportation Systems. Yokohama, Japan, pp. 353-359.
- Baldo, N., Marini, A., Miani, M., 2020. Drivers' braking behavior affected by cognitive distractions: an experimental investigation with a virtual car simulator. *Behavioral Sciences* 10.10, 150.
- Bella, F.; Borrelli, V.; Silvestri, M.; Nobili, F., 2018. Effects on Driver's Behavior of Illegal Pedestrian Crossings. *Advances in Intelligent Systems and Computing* 786, 802-812.
- Hochreiter, S., Schmidhuber, J., 1997. Long short-term memory. *Neural computation* 9.8, 1735–1780.
- Huang, Y., Bi, H., Li, Z., Mao, T., Wang, Z., 2019. STGAT: Modeling spatial-temporal interactions for human trajectory prediction, 2019 IEEE/CVF International Conference on Computer Vision. Seoul, Korea (South), pp. 6271-6280.
- Khairdoost, N., Shirpour, M., Bauer, M.A., Beauchemin, S.S., 2020. Real-time driver maneuver prediction using LSTM. *IEEE Transactions on Intelligent Vehicles* 5.4, 714-724.
- Kingma, D.P., Ba, J.L., 2015. Adam: A method for stochastic optimization. arXiv preprint: 1412.6980v9.
- Laureshyn, A., Svensson, Å., Hydén, C., 2010. Evaluation of traffic safety, based on micro-level behavioural data: Theoretical framework and first implementation. *Accident Analysis and Prevention* 42.6, 1637-1646.
- Lipton, Z.C., Berkowitz, J., Elkan, C., 2015. A Critical Review of Recurrent Neural Networks for Sequence Learning. arXiv preprint: 1506.00019v4.
- Meiring, G.M.M., Myburgh, H.C., 2015. A review of intelligent driving style analysis systems and related artificial intelligence algorithms. *Sensors* 15.12, 30653-30682.
- Ortiz, M.G., Fritsch, J., Kummert, F., Geppert, A., 2011. Behavior prediction at multiple time-scales in inner-city scenarios, 2011 IEEE Intelligent Vehicles Symposium. Baden-Baden, Germany, pp. 1068-1073.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S., 2019. PyTorch: An imperative style, high-performance deep learning library. arXiv preprint: 1912.01703v1.
- Várhelyi, A., 1998. Drivers' speed behaviour at a zebra crossing: a case study. *Accident Analysis and Prevention* 30.6, 731-743.
- Yang, W., Zhang, X., Lei, Q., Cheng, X., 2019. Research on longitudinal active collision avoidance of autonomous emergency braking pedestrian system (AEB-P). *Sensors* 19.21, 4671.
- Zernetsch, S., Kohonen, S., Goldhammer, M., Doll, K., Sick, B., 2016. Trajectory prediction of cyclists using a physical model and an artificial neural network, 2016 IEEE Intelligent Vehicles Symposium. Gothenburg, Sweden, pp. 833–838.
- Zheng, L., Ismail, K., Meng, X.H., 2014. Traffic conflict techniques for road safety analysis: open questions and some insights. *Canadian Journal of Civil Engineering* 41.7, 633–641.