Research article

Johannes Reschke*, Cornelius Neumann and Stephan Berlitz

# Personalised neural networks for a driver intention prediction: communication as enabler for automated driving

https://doi.org/10.1515/aot-2020-0035
Received June 29, 2020; accepted October 6, 2020;
published online October 28, 2020

**Abstract:** In everyday traffic, pedestrians rely on informal communication with other road users. In case of automated vehicles, this communication can be replaced by light signals, which need to be learned beforehand. Prior to an extensive introduction of automated vehicles, a learning phase for these light signals can be set up in manual driving with help of a driver intention prediction. Therefore, a three-staged algorithm consisting of a neural network, a random forest and a conditional stage, is implemented. Using this algorithm, a true-positive rate (TPR) of 94.0% for a 5.0% false-positive rate (FPR) can be achieved. To improve this process, a personalization procedure is implemented, using driver-specific behaviours, resulting in TPRs ranging from 91.5 to 96.6% for a FPR of 5.0%. Transfer learning of neural networks improves the prediction accuracy of almost all drivers. In order to introduce the implemented algorithm in today's traffic, especially the FPR has to be improved considerably.

**Keywords:** automotive lighting; learning signals; recurrent neural networks; time sequence processing; vehicle-pedestrian-communication.

## 1 Introduction

In today's traffic, communication between traffic participants is absolutely necessary in order to guarantee traffic flow, solve unclear situations and signalize yielding [1, 2]. Especially pedestrians rely and depend on interactions with drivers. When crossing the street, they seek eye contact with drivers of approaching vehicles to make sure that they have been seen or that the driver yields them the right of way [3]. Pedestrians' trust in other road users and their perceived safety decreases when there is no communication present [4, 5]. Thus and because pedestrians are the most vulnerable road users in traffic, authorities such as the NHTSA recommend to establish eye contact before crossing the street [6].

Recent development in the automation of vehicles continuously increases the number of automated and autonomous vehicles in our everyday traffic [7]. As the automation level of vehicles increases, the driver becomes more and more a passenger instead of actively taking part in the traffic [8]. This results in a low acceptance of automated vehicles, as there is no active communication partner for pedestrians. Thus, they lose trust when sharing space in traffic with the automated driving systems [9].

It is now the task of automated vehicles to communicate with other road users, especially with pedestrians, to give them certainty about their intention and restore trust (see Figure 1) [10]. Unfortunately, recent surveys regarding the vehicle-pedestrian-communication show that the intuitivity of signals used for this task is quite low and mostly not sufficient for this task. Participants in these surveys often cannot interpret signals correctly in first place [11–13]. Especially simple signals, such as when utilizing led-strips, are not intuitive [11]. In this context, a survey regarding the supporting role of colours for the intuitivity of symbols shows that it enhances the participants' confidence only for certain situations. Most often white and green show the best results, while blue-green is not favourable for the communication of automated vehicles [12]. Blue-green is currently under discussion in standardization organizations (SAE, GTB, ISO and GB standard) although this colour is not intuitive for traffic participants.

This shows clearly that we need a method to introduce a huge portion of the population to the new signals prior to the introduction of automated driving systems. A possibility requiring little effort to teach others is the

***Corresponding author: Johannes Reschke**, Development Functions Light, AUDI AG, Ingolstadt, Germany, E-mail: johannes.reschke@audi.de
**Cornelius Neumann,** Light Technology Institute (LTI), Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany
**Stephan Berlitz,** Lighting Innovations, AUDI AG, Ingolstadt, Germany

**Figure 1:** Symbol-based vehicle-pedestrian-communication [12].

introduction of these newly developed signals in manual driving mode of today's vehicles. Therefore the light signals are shown to other road users, while the driver is communicating as usual. Thus, for example, pedestrians can still rely on drivers' cues, while vehicles display new light signals in parallel. This way, other road users and especially pedestrians can learn to associate both modalities, build an inner model for the vehicle-pedestrian-communication and be prepared for a more frequent contact with automated driving systems.

In order to predict the driver's intention to show light signals automatically, algorithms utilizing artificial neural networks have proven suitable for this task. This is because they can interpret temporal dependencies especially well. Such algorithms were already developed in recent years, still inheriting the disadvantage of being trained on a general dataset and thus, an average driving profile [14, 15]. Artificial neural networks are especially well suited for predicting and processing time series. In literature, they clearly outperform other algorithms such as hidden Markov models [16], Gaussian mixture models [17], auto regressive (integrated) moving average or Box-Jenkins [18].

Optimizing the developed algorithms to predict a driver's intention to stop, they can be adapted to driver and vehicle specific behaviours. Therefore mostly transfer learning methods show excellent results [19]. Adapting the algorithm is especially important in order to decrease the classification's false-positive rate (FPR). False-positives result in potentially dangerous situations where the vehicle indicates an expected stop, while the driver wants to continue driving. Using transfer learning, algorithms can apply knowledge from a previously learned domain in a second novel task [19]. Utilizing driver and vehicle specific data, a general algorithm can be adapted to the driver after delivery of the vehicle to the customer. Here, already little training kilometres can be used to generate a suitable dataset to personalize neural networks for the driver intention prediction. The more data we collect, the better the neural network can adapt to the driver and the vehicle. Nevertheless, we expect the transfer learning's training effect to plateau as the amount of data increases.

Therefore, an algorithm predicting a driver's intention at pedestrian crossings is implemented. Using transfer learning, the algorithm is adapted to driver- and vehicle-specific behaviours to predict whether a driver wants to stop or continue driving more reliable.

## 2 Technical background

Neural feedforward networks process, as most deep learning algorithms, input variables at a single point in time without considering previous time steps and information. Recurrent neural networks (RNNs), especially LSTMs (long short-term memories) use feedback loops in their architecture to analyze sequences of data such as time variant signals [20]. LSTMs are gated RNNs that use multiple gates and an inner cell state to determine which information to store in the LSTM cell, which to output or forget. Because of these gates, LSTMs are capable of avoiding vanishing and exploding gradients, which typically occur in RNN architectures. This is essential to train deep neural networks and classify the driver intention reliably [20, 21].

Transfer learning is used to apply already learned knowledge inherited in a model to train a task specific, modified model. Thus, the classification accuracy can be increased using only a small amount of data [19]. Developers often use transfer learning when there is only little data available for the specific task and high accuracies are to be achieved. For neural networks, there are multiple ways of utilizing transfer learning, for example fine-tuning all layers of the architecture, retrain certain layers or adding new layers. Fine-tuning the network and adding new layers requires in comparison to only retraining certain layers more data and training effort. Retraining layers is often used for similar tasks, where little data are available [22, 23].

As in the previously developed algorithm, we semi-automatically label the data with 39 rules that we obtain from manually reviewing multiple time series. Utilizing these rules to train a random forest, we can predict whether the driver wants to stop or continue driving. For most stopping situations, vehicles do not come to a full standstill, which is why we need to use a stopping velocity below 7 km/h. Therefore, classifying whether the vehicle will come to a stop within the prediction horizon, we use a velocity threshold of 7 km/h [14]. This threshold is also used for detecting stops in automatic start–stop systems [24].

## 3 Driver intention prediction algorithm

Figure 2 depicts the driver intention prediction algorithm, which consists of three stages [14]. In stage one, a recurrent neural network with four LSTM layers and two fully connected layers predicts from all 22 input signals 5 output signals' time series over a prediction horizon of 2 s. These output signals are velocity, break pressure, steering wheel angle, longitudinal and lateral acceleration. The second stage implements a random forest classifier, which predicts from the 5 output signals (each spanning over 2 s) of stage one the driver's willingness to stop, i.e. driver's intention to stop in the next 2 s. Therefore, the random forest consists of 25 decision trees with a depth of 12 each. At the end of stage two, their prediction probabilities are averaged and processed with a threshold. The first two stages are trained independently from each other. The third stage compares the vehicle's speed at the end of the prediction horizon with the velocity threshold of 7 km/h, which is also used for automatic start–stop systems. Herewith, stage three predicts consecutive to the driver's willingness to stop, if the vehicle will come to a stop in the prediction horizon of 2 s.

We chose this three-staged approach in order to make the algorithm more interpretable. Deep neural networks are quite complex and very difficult to interpret [25], which is why we decided to use a more interpretable algorithm, a random forest, for stage two. In this stage, the actual prediction of the driver's willingness to stop is made. The decision whether the vehicle will comes to a stop or continues driving is based on a simple rule set, which is easy to interpret. Dividing the algorithm into several smaller subsections is based on the divide and conquer approach [26]. We totally excluded an end-to-end system, which is considered to be uninterpretable [27], as this algorithm applies to a safety related task.

The dataset for a general driver intention prediction algorithm at pedestrian crossing consists of 7114 time series from multiple test drivers [15]. These time series build a training dataset with a total length of approximately 120 hours and contain 1114 stops at zebra crossings. Test drivers had collected the data from April 2017 until May 2018. The dataset consisting of time series spanning over a full year guarantees to have diverse training samples. In total, we use 22 relevant signals with a temporal resolution of 10 ms, including vehicle speed, breaking pressure, steering wheel angle, camera objects, etc. We normalize all data with a min–max normalization in order to have similar orders of magnitude for all signals and therefore similar influence on the training process [28].

With this implemented algorithm and dataset, stage two achieves a false-positive rate (FPR) of 5.0% with a true-positive rate (TPR) of 93.6%. After processing it with the 7 km/h threshold, the algorithm results in stage three in a TPR of 94.0% at the same FPR (see Figure 3). While training, the algorithm reaches a minimum validation-MSE

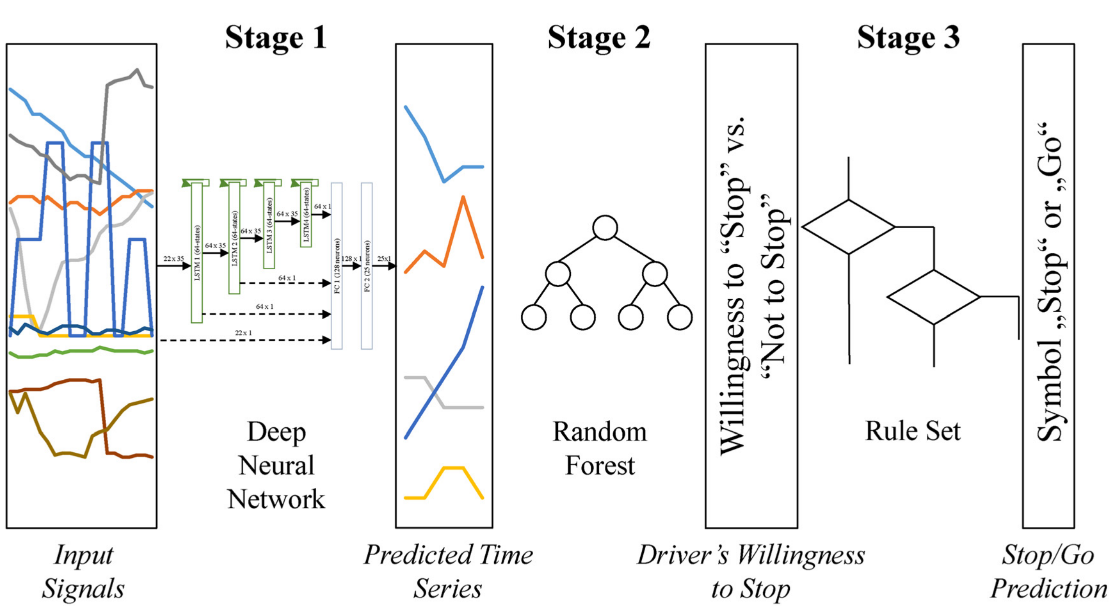


**Figure 2:** Architecture of the driver intention prediction algorithm [14].

(mean squared error) of $3.16 \times 10^{-4}$ and therefore stops training after 47 epochs. With the resulting TPRs and FPRs, visualized by the ROC curves (receiver operating characteristic) in Figure 3, the algorithm misses the requirements set by the authors with TPR≥95 and FPR≤5% slightly.

The results of the algorithm after stage three show higher accuracies than after stage two. This is because of the filtering and comparison with the 7 km/h threshold, with which some uncertain classification results are set to label "continue driving". Above 5.0% FPR both curves converge. None of both ROC curves has a smooth shape and especially the graph of stage three shows a saddle point around an FPR of 1.0%. This can be explained with the small test dataset with a low diversity of test samples. We use quite a small test set in order to have a higher number of training sequences to result in a well-trained algorithm.

## 4 Personalized neural networks using transfer learning

As described in chapter 2, transfer learning can help personalizing neural networks in order to get task-specific models. Therefore, a specialized dataset is necessary, which can be much smaller than the dataset for the pre-trained algorithm. Essential for this task and for getting a personalized driver intention prediction is the development of a representative dataset, which contains as many driver-specific characteristics as possible.

The first stage of the driver intention prediction algorithm needs to be adapted to driver and vehicle specific properties, which could be driving behaviour, breaking or steering characteristics. Preferably, an as small as possible dataset is to be used, in order to allow a relevant and most practical application in later series application. Therefore, we collect the data of four randomly selected drivers on a reference route in and around Ingolstadt, Germany. The recorded data have the same format and feature set as the above-described dataset for pre-training the algorithm. The reference route is 36.6 km long and consists of 56.2% city traffic, 31.0% country roads and 12.8% highways, as average routes of German drivers do as well [29, 30]. In total, the test drivers pass six zebra crossings, 47 traffic lights and nine unsignalized pedestrian crossings.

We obtain the optimal length of the reference route by evaluating various lengths regarding their test-MSE after applying transfer learning and selected data to the baseline network. Figure 4 confirms the expected results that larger datasets lead to lower MSEs, but also that there is a saturation effect at a certain point. Applying the elbow-method indicates that a dataset length of approximately 55 min is optimal for this task. This gives a compromise between high accuracies and low training efforts. Driving the reference route takes about 60 min depending on traffic and driving speed. In order to test the algorithm with an independent test set, a second reference route in Ingolstadt, Germany is set up. This reference route contains only city traffic and is approximately 3.3 km long.

A comparison of multiple transfers learning methods shows that personalization of the neural network's last two layers is favourable (see Table 1.). The improvement of the neural network's validation-MSE shows for adapting the last two and three layers the best results of the surveyed transfer learning methods. These two options show almost identical results for the validation-MSE's improvement in comparison to the pre-trained network. The less layers need to be adapted to driver-specific characteristics, the faster the training process is. Therefore, we choose to personalize the neural network's last two layers. Adding and training a new layer leads to a less accurate driver intention prediction than the baseline's network, most likely because there is not enough training data to adopt to

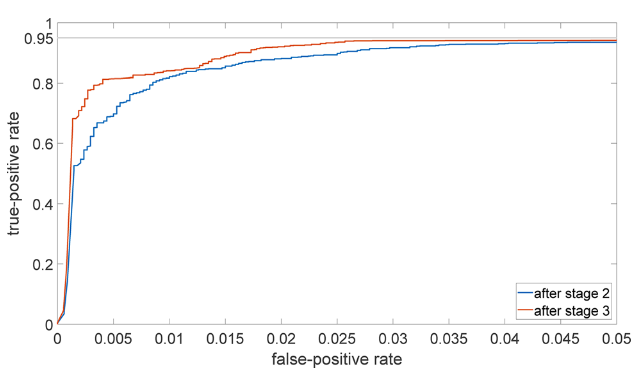


**Figure 3:** ROC curves of the pre-trained driver intention prediction algorithm after stage two and three.

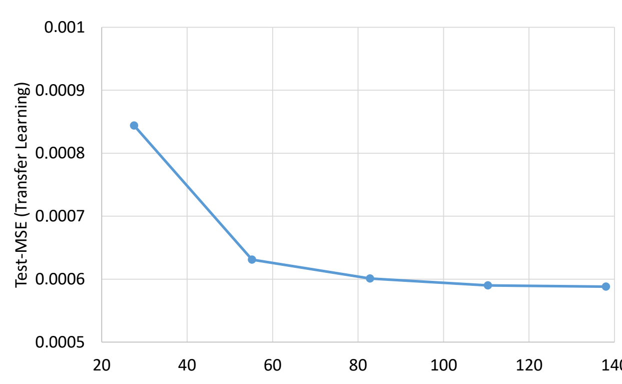


**Figure 4:** Evaluation of the optimal length of the training dataset.

**Table 1:** Improvement of validation-MSE for different transfer learning methods.

| | All layers | First layer | New layer | Last layer | Last 2 layers | Last 3 layers | Ø |
|---|---|---|---|---|---|---|---|
| Driver 1 | −15.10 | −12.05 | 8.92 | −14.58 | −14.98 | −15.00 | −10.46 |
| Driver 2 | −9.18 | −8.62 | 12.30 | −9.40 | −9.58 | −9.58 | −5.68 |
| Driver 3 | −18.04 | −15.30 | −6.72 | −19.86 | −20.57 | −20.47 | −16.83 |
| Driver 4 | −14.63 | −12.45 | 6.38 | −13.66 | −13.77 | −13.94 | −10.34 |
| Ø | −14.24 | −12.10 | 5.22 | −14.37 | −14.73 | −14.74 | −10.83 |

driver characteristics sufficiently. Retraining of the neural network's last layers outperforms the personalization of the first layer clearly. This is due to the network's architecture, where the first four LSTM layers learn temporal dependencies of certain features, while the last two fully connected layers combine them to output vectors. For all drivers, temporal dependencies stay constant or are very similar for all drivers, while output values of the signal vectors might differ quite a lot.

Figure 5 shows the ROC curves of the driver intention prediction algorithm's third stage with personalized neural networks. The network is adapted to one of the four drivers each. Applying the transfer learning for all drivers, the algorithm results in TPRs between 91.5 and 96.6% at an FPR of 5.0%. These values are realistic interpolations of the ROC curve's measured values [31]. The transfer learning for driver four's dataset shows best results underneath all test persons with a TPR of 95.2% and FPR of 1.8%. Furthermore, Figure 5 shows that above an FPR of 2.9%, the personalized ROC curves and therefore TPRs of almost all drivers (except driver two) stay clearly above the pre-trained, non-personalized neural network. For driver two, the personalization of the driver intention prediction decreases the classification accuracy in the entire range. This is mostly because of unsteady and hardly predictable driving behaviour in both reference routes. For all drivers, especially driver one and four, Figure 5 shows a steady increase of the TPR from 0 to the 1st measurement point. The reason therefore is a lack of data points due to the small test datasets. Adapting the algorithm to driver-specific behaviour of driver three, the classification accuracy outperforms the baseline above an FPR of 2.1%. Above this point, the TPR is higher than 95.5%. Thus, the classification accuracy is only superior to the baseline algorithm above a TPR of 92.5%, which is for this task sufficient, since the minimal required TPR is 95% and thus, lies above this point.

After stage two, the ROC curves of the personalized neural networks show in general similar results to stage three (see Figure 6). Again, results for driver one and four show best results underneath these test persons and result in higher TPRs as the baseline/pre-trained algorithm. TPRs of driver three and especially of driver two stay clearly below the baseline algorithm's ROC curve, which means that personalization after stage two only works for two drivers well, while the accuracy for two drivers decreases.

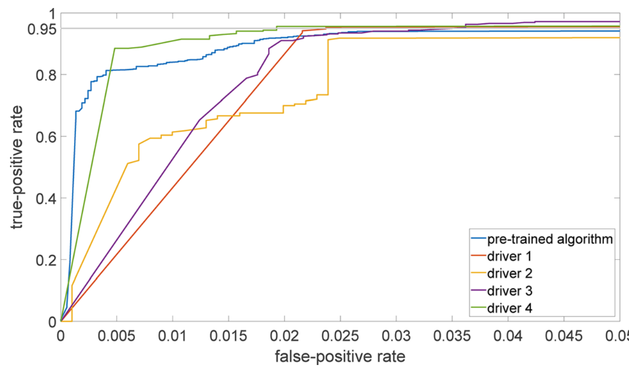


**Figure 5:** ROC curves of the personalized algorithm after stage three.

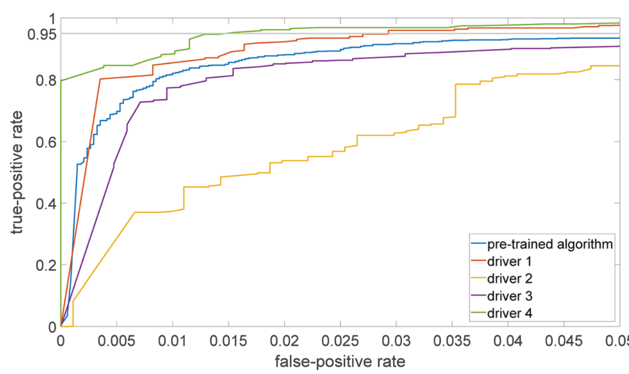


**Figure 6:** ROC curves of the personalized algorithm after stage two.

The comparison between stage two and three shows, that as for the pre-trained network, stage three can add certainty to the classification results. Nevertheless, we achieve the highest TPRs using only stage two, which indicates that distinguishable examples are filtered out in stage three.

Even though that there is an increase in prediction accuracies for driver one and four, this increase is not significant. Nevertheless, not only the networks are adjusted to a specific driver, but also the truths for data labelling are personalized. Thus, the driver-specific algorithms are representing their actual intention better.

## 5 Summary and outlook

Using a three-staged algorithm, we implemented a driver intention prediction as part of a vehicle-pedestrian-communication at pedestrian crossings. The first stage predicts time series of five different signals, which the algorithm interprets in stage two. In the final stage, the predictions from stage two are validated and result in a generalized classifier.

In order to meet the requirements for the driver intention prediction defined in the study by Reschke et al. [14], one possibility is to retrain certain parts of the neural network. For this, a driver- and vehicle-specific datasets are needed to personalize the neural network. Mostly, transfer learning is used for this task and in this contribution, adapting the last two fully connected layers of the neural network has shown best results. For driver four, whose personalized data worked best to adapt the network, a TPR of 95.2% and an FPR of 1.8% result from the evaluation? The personalization utilizing transfer learning worked especially well for driver one and four and led to clearly higher TPRs, but for driver two, the ROC curve shows much lower true-positive rates. Summarizing, the developed algorithm and the personalized neural networks are capable of predicting a driver's intention at pedestrian crossings to realize a vehicle-pedestrian-communication for manual driving. Thus, it serves as an enabler for the future communication of automated vehicles and therefore their acceptance and trust in these systems.

The introduced algorithm clearly outperforms reported results from literature of similar algorithms. The reported accuracies by Zhu et al. [32] and Tran et al. [33] reach 80.4 and 90% while our algorithm has an accuracy of 96.6%. The algorithm of Garcia et al. [34] achieves over a prediction horizon of 2 s a TPR of 65% and an FPR of 10%, which are clearly below the quality measures of our algorithm.

Even though that the requirements of the driver intention prediction are fulfilled for three drivers using personalized neural networks, at least 1.8% false positives occur. This means that for 1.8% of the positive predictions to stop, the driver actually wants to continue driving. This obviously is a serious risk for all traffic participants and especially for pedestrians.

Therefore, for further development of this and future algorithms, the FPR is required to stay significantly below 5% or needs to be secured with additional methods. Focus of future evaluations is the personalisation not only on both, driver- and vehicle-specific data, but also on each characteristics separately. This could be advantageous when multiple people use one vehicle or if one customer changes vehicles quite often. For this, the collected datasets need to be enlarged or enriched. Furthermore, a prediction of time series with help of time-convolutional neural networks (TCNs) is possible and potentially the classification accuracy can be increased [35]. Also, the personalization of neural networks using less data would be helpful for vehicles' series development as well as lifelong learning or continuous learning over a full vehicle life period. Finally, the driver could be integrated in the system's visualization loop and be informed what the systems predicts. This information can potentially avoid that a driver continues driving while the system indicates the pedestrian that the vehicle will stop.

**Acknowledgement:** We would like to thank especially Nawel Attia, Benedict Schleyer and Thomas Hoess, whose master theses layed the foundation for this contribution.
**Author contribution:** All the authors have accepted responsibility for the entire content of this submitted manuscript and approved submission.
**Research funding:** None declared.
**Competing interest:** The authors declare no conflicts of interest regarding this article.

## References

[1] K. Merten, *Informelle Zeichengebung im Straßenverkehr: Bericht zum Forschungsprojekt 7521 der Bundesanstalt für Straßenwesen Bereich Unfallforschung*, Köln: Bundesanstalt für Straßenwesen Bereich Unfallforschung, 1981.
[2] R. Risser, *Kommunikation und Kultur des Straßenverkehrs*, 1st ed. Wien, Österreich, Literas-Universitäts-Verlag, 1988.
[3] A. Rasouli, I. Kotseruba, and J. K. Tsotsos, "Understanding pedestrian behavior in complex traffic scenes," *IEEE Trans. Intell. Veh.*, vol. 3, pp. 61–70, 2018.
[4] D. Rothenbücher, J. Li, D. Sirkin, B. Mok, and W. Ju, "Ghost driver: A field study investigating the interaction between pedestrians and driverless vehicles," in *25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, IEEE, ed., New York, NY, USA: IEEE, 2016, pp. 795–802.
[5] V. M. Lundgren, A. Habibovic, J. Andersson, et al., "Will there be new communication needs when introducing automated vehicles to the urban context?," in *Advances in Intelligent Systems and Computing*, *Advances in Human Aspects of Transportation: Proc. of the AHFE 2017 International Conf. on Human Factors in Transportation*, vol. 597, Los Angeles, CA, USA, N. A. Stanton, Ed., Cham, Schweiz, Springer International Publishing, 2018, pp. 485–497.

[6] National Highway Traffic Safety Administration and US, *Department of Transportation*, Washington, DC, USA, Traffic Safety Facts: Pedestrians. 2016 Data, 2018.
[7] C. R. Hudson, S. Deb, D. W. Carruth, J. McGinley, and D. Frey, “Pedestrian perception of autonomous vehicles with external interacting features,” in *Advances in Human Factors and Systems Interaction: Proceedings of the AHFE*, Orlando, FL, USA: Springer, 2018, pp. 33–39.
[8] J. Domeyer, A. Dinparastdjadid, J. D. Lee, G. Douglas, A. Alsaid and M. Price., “Proxemics and Kinesics in Automated Vehicle–Pedestrian Communication: Representing Ethnographic Observations,” in vol. 2673, *Transportation Research Record: Journal of the Transportation Research Board*, National Academy of Sciences, EdSAGE journals., Washington, DC, 2019, pp. 70–81.
[9] A. Schaudt and S. Russell, “Judging a car by its cover: Human factors implications for automated vehicle external communication,” in *Lecture Notes in Mobility, Road Vehicle Automation 5*, G. Meyer and S. Beiker, Eds., Cham, Schweiz, Springer International Publishing, 2018, pp. 69–76.
[10] J.-H. Willrodt, H. Strothmann, and J. Wallaschek, “Optical car-to-human communication for automated vehicles,” in *Darmstädter Lichttechnik, Proc. of the 12th International Symposium on Automotive Lighting: ISAL 2017*, T. Q. Khanh, Ed., München: utzverlag, 2017, pp. 579–588.
[11] J. Reschke, P. Rabenau, M. Hamm, and C. Neumann, “Symbolische fahrzeug-fußgänger-interaktion,” in *Optische Technologien in der Fahrzeugtechnik: Proc. der 8. VDI-Fachtagung*, vol. 8, V. D. I Wissensforum GmbH, Ed., Düsseldorf, VDI Verlag GmbH, 2018, pp. 95–106.
[12] J. Reschke, M.-T. Auburger, R. Marichalar, and C. Neumann, “Kommunikation zwischen automatisierten Fahrzeugen und Fußgängern,” *ATZ - Automob. Z.*, vol. 121, no. 9, pp. 16–21, 2019.
[13] I. Othersen, S. Cramer, and C. Salomon, “HMI for external communication - kann die Fahrzeugbewegung als Kommunikationskanal zwischen einem Fahrzeug und einem Fußgänger dienen?,” in: VDI Wissensforum GmbH, editor. 10. *VDI-Tagung Mensch-Maschine-Mobilität*, Braunschweig: VDI Verlag GmbH; 2019, vol. 2360, pp. 145–154.
[14] J. Reschke, T. Höß, B. Schleyer, S. Berlitz, and C. Neumann, “How vehicles learn to display symbols to pedestrians,” in *Darmstädter Lichttechnik Proc. of the 13th International Symposium on Automotive Lighting: ISAL 2019*, vol. 18, T. Q. Khanh, Ed., München: utzverlag, 2019, pp. 590–599.
[15] J. Reschke, S. Prösl, M. Hamm, and C. Neumann, “Assistance system for vehicle-pedestrian-interaction: Deep learning and driver intention prediction,” in *SIA VISION* 2018: *Vehicle Infrastructure Safety Improvement in Adverse Conditions and Night Driving, SIA - French Society of Automotive Engineers*, Paris, Frankreich, 2018, pp. 19–26.
[16] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” in *Neural Computation, Massachusetts Institute of Technology*, vol. 9, Boston, MA, USA: MIT Press, 1997, pp. 1735–1780.
[17] H.Salehinejad, S.Sankar, J.Barfett, E.Colak, and S.Valaee, “Recent advances in recurrent neural networks,” [Online]. Available at: https://arxiv.org/abs/1801.01078 [accessed: Nov. 28, 2019].
[18] B.Oancea and S. C.Ciucu, “Time series forecasting using neural networks,” [Online]. Available at: https://arxiv.org/abs/1401.1333 [accessed: May. 21, 2019].
[19] S. J. Pan and Q. Yang, “A survey on transfer learning,” *IEEE Transactions on Knowledge and Data Engineering*, Piscataway, NJ, USA: IEEE, 2010, vol. 22, pp. 1345–1359.
[20] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*, Cambridge, Massachusetts, London, England, MIT Press, 2016.
[21] F. Gers, “Long short-term memory in recurrent neural networks,” Lausanne, Schweiz, Disseration, École polytechnique fédérale de Lausanne, 2001.
[22] Y.Guo, H.Shi, A.Kumar, et al, “SpotTune: transfer learning through adaptive fine-tuning,” [Online]. Available at: https://arxiv.org/abs/1811.08737 [accessed: Nov. 28, 2019].
[23] S. Yoon, H. Yun, Y. Kim, G.-T. Park, and K. Jung, “Efficient transfer learning schemes for personalized language modeling using recurrent neural network,” *31st Association for the Advancement of Artificial Intelligence: Workshop on Crowdsourcing, Deep Learning and Artificial Intelligence Agents*, San Francisco, CA, USA: AAAI Press, 2017, pp. 457–463.
[24] K.Pudenz, “Volkswagen zeigt Start-Stopp 2.0 und Freilauf-Motor-Aus,” [Online]. Available at: https://www.springerprofessional.de/automobil—motoren/volkswagen-zeigt-start-stopp-2-0-und-freilauf-motor-aus/6585236 [accessed: Jun. 03, 2019].
[25] K. R. Varshney, “Engineering safety in machine learning,” in *2016 Information Theory and Applications Workshop (ITA): Jan. 31, 2016-Feb. 5, 2016*, IEEE, Piscataway, NJ, USA, IEEE, 2016.
[26] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms*, 3rd ed. Cambridge, MA, USA, MIT Press, 2009.
[27] F. Al-Khoury, “Safety of machine learning systems in autonomous driving,” Master thesis, Stockholm, Schweden, School of Industrial Engineering and Management, KTH Royal Institute of Technology, 2017.
[28] I. N. D. Silva, D. H. Spatti, R. Andrade Flauzino, L. H. B. Liboni, and S. F. D. Reis Alves, *Artificial Neural Networks: A Practical Course*, Cham, Schweiz, Springer, 2017.
[29] Kraftfahrt-Bundesamt, “Verkehr in kilometern: inländerfahrleistung,” [Online]. Available at: https://www.kba.de/DE/Statistik/Kraftverkehr/VerkehrKilometer/verkehr_in_kilometern_node.html [accessed: Jan. 03, 2020].
[30] Bundestanstalt für Straßenwesen, “Ergebnisse Fahrleistungserhebung 2014: erhebung der Inländerfahrleistung und der Inlandsfahrleistung,” [Online]. Available at: https://www.bast.de/BASt_2017/DE/Verkehrssicherheit/Fachthemen/u2-fahrleistung-2014/u2-Fahrleistung-2014-ergebnisse.html [accessed: Jan. 03, 2020].
[31] T. Fawcett, “An introduction to ROC analysis,” *Pattern Recogn. Lett.*, vol. 27, no. 8, pp. 861–874, 2006.
[32] L. Zhu, S. Li, Y. Li, et al, “Study on driver’s braking intention identification based on functional near-infrared spectroscopy,” *J. Intell. Connected Veh.*, vol. 1, Emerald Publishing Limited, pp. 107–113, 2018.
[33] D. Tran, W. Sheng, L. Liu, and M. Liu, “A hidden markov model based driver intention prediction system,” in *IEEE International Conf. on CYBER Technology in Automation, Control, and Intelligent Systems (CYBER): 8 - 12 June 2015, Shenyang, China, IEEE*, Piscataway, NJ, USA, IEEE, 2015, pp. 115–120. 2015.

[34] M. Garcia Ortiz, J. Fritsch, F. Kummert, and A. Gepperth, "Behavior prediction at multiple time-scales in inner-city scenarios," in *IEEE Intelligent Vehicles Symposium (IV), IEEE*, Piscataway, NJ, USA, IEEE, 2011, pp. 1068–1073.
[35] H. I. Fawaz, G. Forestier, J. Weber, L. Idoumghar, and P.-A. Muller, "Transfer learning for time series classification," in *IEEE International Conf. on Big Data: Proceedings: Dec 10 - Dec 13, 2018*, Seattle, WA, USA, N. Abe, Ed., Piscataway, NJ, USA, IEEE, 2018, pp. 1367–1376.

## Bionotes

**Johannes Reschke**
Development Functions Light, AUDI AG, Ingolstadt, Germany
**johannes.reschke@audi.de**

Johannes Reschke studied Electrical Engineering and Applied Research in Engineering Sciences at the Ostbayerische Technische Hochschule Regensburg, Germany. He is currently working on his PhD thesis at the Light Technology Institute at Karlsruhe Institute of Technology, Germany. In parallel, he is a development engineer in the Department of Development Lighting Innovations/Functions at Audi, Germany. His research focuses on vehicle-pedestrian-communication and machine learning.

**Stephan Berlitz**
Lighting Innovations, AUDI AG, Ingolstadt, Germany

Stephan Berlitz studied Electrical Engineering at the Technical University of Munich, Germany. Afterwards, he started his professional carrier in the Development and Technical Coordination for Tail Lights at Schefenacker in Esslingen, Germany. He worked in the Department of Innovations Lighting at Audi, Germany, from 2001 to 2005. Since then, he is head of Lighting Innovations/Functions and therefore responsible for leading light innovations at Audi.

**Cornelius Neumann**
Light Technology Institute (LTI), Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany

Cornelius Neumann studied Physics and Philosophy at the University of Bielefeld, Germany. After his PhD, he worked for the automotive supplier Hella in the advanced development for Automotive Lighting. During his time at Hella, he was responsible for signal lighting, LED application and acted as a director of the L-LAB, a laboratory for lighting and mechatronics in public private partnership with the University of Paderborn, Germany. In 2009, he became professor for Optical Technologies in Automotive and General Lighting and one of the two directors of the Light Technology Institute at the Karlsruhe Institute of Technology, Germany.