Proceedings Paper:

Qasemabadi, A.N., Mozaffari, S., Rezaei, M. et al. (2026) Uncertainty-Aware Lane Change Prediction for Autonomous Driving Using LSTM Networks. In: 2025 IEEE Canadian Conference on Electrical and Computer Engineering (CCECE). 2025 IEEE Canadian Conference on Electrical and Computer Engineering (CCECE), 26-29 May 2025, Vancouver, BC, Canada. Institute of Electrical and Electronics Engineers (IEEE), pp. 497-501. ISSN: 0840-7789.

https://doi.org/10.1109/ccece64018.2025.11364403

Uncertainty-Aware Lane Change Prediction for Autonomous Driving Using LSTM Networks

Armin Nejadhossein Qasemabadi[1], Saeed Mozaffari[2], Mahdi Rezaei[3], and Shahpour Alirezaee[2,*]
[1]Electrical and Computer Engineering Department, University of Windsor, Windsor, Canada
nejadho@uwindsor.ca
[2]Mechanical, Automotive, and Material Engineering Department, University of Windsor, Windsor, Canada
{saeed.mozaffari, s.alirezaee}@uwindsor.ca
[3]Institute for Transport Studies, University of Leeds, Leeds, UK
m.rezaei@leeds.ac.uk

***Abstract*—This paper addresses the challenge of accurately predicting lane change maneuvers in autonomous driving while quantifying the prediction uncertainty. We present an LSTM-based model that predicts three classes: lane keeping, left lane change, and right lane change. By considering the position, velocity, and acceleration of multiple vehicles, we enhance the accuracy of our predictions, achieving an overall accuracy of 97.10% compared to our base LSTM model. Our analysis of different uncertainty sources, including Gaussian noise injection, Dropout noise, and deep ensembles, reveals varying levels of uncertainty, with mean standard deviations ranging from 0.0084 to 0.1094. This highlights the model's sensitivity to different types of perturbations and the importance of uncertainty quantification for robust lane change prediction. This contribution enables more informed decision-making in ADAS and paves the way for safer autonomous driving.**

I. INTRODUCTION

Advanced Driver Assistance Systems (ADAS) are revolutionizing the automotive industry by enhancing safety and driving comfort. These systems represent a significant step towards the realization of fully autonomous vehicles, capable of navigating and operating without human intervention. This evolution promises to reduce accidents, improve traffic flow, and enhance overall transportation efficiency. ADAS technologies leverage a sophisticated fusion of cutting-edge sensor technologies and artificial intelligence algorithms to perceive and interpret the complex driving environment [1], [2].

These sensor technologies, including cameras, LiDAR, and Radar, act as the eyes and ears of the vehicle, gathering crucial information about the road, other vehicles, and pedestrians. The data collected from these sensors is then processed by sophisticated artificial intelligence algorithms, enabling features such as collision avoidance, adaptive cruise control, and lane-keeping assistance. These features not only enhance driver safety by providing warnings and assisting in critical situations but also improve the overall driving experience by automating routine tasks and reducing driver workload.

Recent research has focused on developing robust and accurate lane change prediction models to further enhance the capabilities of ADAS [3], [4]. These models aim to anticipate the intentions of drivers to change lanes, enabling the vehicle to make informed decisions and take appropriate actions to ensure safety and efficiency. Accurate and timely lane change prediction can prevent accidents, improve traffic flow, and enhance the overall driving experience.

However, the realization of reliable lane change prediction hinges on addressing a critical challenge: uncertainty. Uncertainty in lane change prediction arises from various sources. Sensors, for instance, are not perfect and can be affected by noise, occlusions, and adverse weather conditions, leading to inaccurate or incomplete perception of the environment. Furthermore, drivers can exhibit unpredictable behaviors, making it difficult to anticipate their intentions and maneuvers. The real-world driving environment itself is inherently complex and dynamic, with constantly changing conditions that can be difficult to model and predict accurately.

Therefore, it is imperative to incorporate uncertainty quantification into lane change prediction models to ensure robust and trustworthy decision-making in the face of unforeseen circumstances. By explicitly modeling and quantifying uncertainty, these models can provide more reliable predictions, enabling ADAS to make more informed and cautious decisions, leading to safer and more efficient autonomous driving systems.

II. UNCERTAINTY MEASUREMENT

Uncertainty quantification plays a pivotal role in developing robust and reliable systems that interact with the real world, including autonomous driving systems. Uncertainty represents the lack of complete knowledge about a particular event, measurement, or prediction, and can arise from various sources. In the realm of autonomous driving, uncertainty can stem from limitations in sensor data (noise, occlusions, environmental factors), the inherent unpredictability of events and behaviors (such as pedestrian movements), and the complexity of real-world environments with their dynamic interactions. By understanding and quantifying these uncertainties, we can equip systems with the ability to make informed decisions even when faced with incomplete or noisy information, which is essential for ensuring safety and building trust.

To further understand the nature of uncertainty, it is often helpful to categorize it into different types. Aleatoric uncertainty stems from the inherent randomness and variability in the environment [5]. It captures the unpredictable nature of events, such as sudden vehicle movements or unexpected changes in weather conditions. Epistemic uncertainty arises from limitations in our knowledge and models of the world [6]. It reflects our lack of understanding or incomplete data about certain aspects of the environment and can be reduced by improving our models or gathering more data. For example, in time series forecasting, epistemic uncertainty might arise from limited historical data or inadequate model architectures. Techniques like knowledge distillation [7] can help transfer knowledge between models, potentially reducing epistemic uncertainty.

In many applications, it is desirable to estimate the total uncertainty, which combines both aleatoric and epistemic components. This can be achieved through various approaches that consider the variability in predictions across different models or data instances. One such approach [8] quantifies total uncertainty by considering the individual uncertainties and the deviations from the overall mean prediction:

$$\sigma_\epsilon^2 = \frac{1}{K} \sum_{\theta \in \mathcal{E}} \sigma_\theta^2 + \frac{1}{K-1} \sum_{\theta \in \mathcal{E}} (\mu_\theta - \mu_{\mathcal{E}})^2 \tag{1}$$

This equation calculates the total uncertainty (σ_ϵ^2) by combining the individual uncertainties (σ_θ^2) and the deviations from the overall mean prediction ($\mu_\theta - \mu_{\mathcal{E}}$).

To quantify the uncertainty associated with lane change predictions, we need to address both aleatoric and epistemic uncertainty. In real-world autonomous driving scenarios, aleatoric uncertainty stems from sensor noise and data acquisition limitations, including measurement uncertainties, imperfections, and errors in vehicle state estimation. Additionally, temporary sensor failures or missing data due to occlusions and environmental conditions contribute to the overall uncertainty. Model uncertainty arises from limitations in our understanding of the complex driving environment and the variability in possible model parameters that could explain the training data. This comprehensive understanding of different uncertainty sources is crucial in safety-critical applications like autonomous driving, where even small uncertainties can have significant consequences.

III. RELATED WORK

Uncertainty-aware approaches have been gaining traction in the field of autonomous driving, with researchers exploring methods to incorporate uncertainty information into various aspects of ADAS, including perception, planning, and control.

In object detection, [9] compares different uncertainty estimation approaches for deep learning components in autonomous vehicles, highlighting the importance of both aleatoric and epistemic uncertainty for reliable perception and addressing computational complexity challenges. [10] proposes a method for generating universal adversarial perturbations with strong transferability across models, emphasizing the need for robust DNN-based perception systems in safety-critical applications.

For path planning, [11] introduces an RNN-based prediction method using deep ensembles to estimate uncertainty in obstacle vehicle trajectories, enabling autonomous vehicles to anticipate and adapt to potential path deviations. In control systems, [12] presents PNNUAD, incorporating uncertainty from perception neural networks into decision-making, while [13] develops a DQN-based approach for lane change decisions that accounts for uncertainty in surrounding vehicle behaviors.

Recent surveys in this field have provided comprehensive overviews of uncertainty quantification methods and their applications. For instance, [14] discusses different types of uncertainty, their sources, and their impact on autonomous system performance, while highlighting future research directions in this critical area.

IV. PROPOSED METHOD

To address the challenges of accurate and reliable lane change prediction, we propose an LSTM-based model that incorporates uncertainty quantification. Our model leverages a combination of Gaussian noise injection, Dropout noise, and deep ensembles to capture different facets of uncertainty inherent in the prediction task.

A. LSTM Network

Recurrent Neural Networks (RNNs) are a type of artificial neural network designed to process sequential data, with a unique structure that maintains a hidden state, or memory, capturing information from previous inputs in the sequence. While RNNs effectively model temporal dependencies in tasks such as natural language processing, speech recognition, and time series forecasting, Long Short-Term Memory (LSTM) networks address their vanishing gradient problem through a specialized gated cell architecture [15]. LSTMs selectively remember or forget information over long periods using three main gates: the input gate (i_t), forget gate (f_t), and output gate (o_t), which regulate the flow of information into and out of the LSTM cell. The cell state (c_t) is updated based on these gates and the input (x_t) at each time step, and the hidden state (h_t) is calculated from the cell state, enabling the network to learn complex temporal patterns. These interactions are formally defined as:

$$i_t = \sigma(W_{xi} x_t + W_{hi} h_{t-1} + b_i) \tag{2}$$

$$f_t = \sigma(W_{xf} x_t + W_{hf} h_{t-1} + b_f) \tag{3}$$

$$o_t = \sigma(W_{xo} x_t + W_{ho} h_{t-1} + b_o) \tag{4}$$

$$c_t = f_t \odot c_{t-1} + i_t \odot \tanh(W_{xc} x_t + W_{hc} h_{t-1} + b_c) \tag{5}$$

$$h_t = o_t \odot \tanh(c_t) \tag{6}$$

where h_{t-1} is the hidden state from the previous time step, W_{xi}, W_{hi}, etc. are weight matrices, b_i, b_f, etc. are biases, σ is the sigmoid activation function, and $\odot$ represents element-wise multiplication. These equations illustrate the complex interactions between the gates and the cell state in an LSTM network, enabling it to effectively capture long-term dependencies and learn intricate temporal patterns in sequential data.

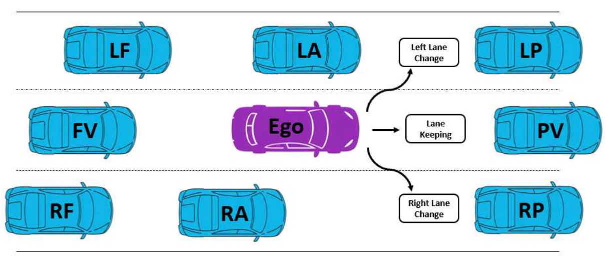

Fig. 1. Lane change scenarios and vehicle labels relative to the ego vehicle (purple). Surrounding vehicles are shown in blue. Lane change maneuvers are categorized into three classes: left lane change, lane keeping, and right lane change.

For lane change prediction, our LSTM network processes vehicle trajectory data in a specific temporal window. The input consists of 75 consecutive frames (timesteps) before the potential lane change, where each frame contains 54 features describing the ego vehicle and its surroundings. These features include the ego vehicle's state, represented by its position (x, y), velocity (xVelocity, yVelocity), and acceleration (xAcceleration, yAcceleration), as well as information about surrounding vehicles, including their positions relative to the ego vehicle (preceding, following, left preceding, left alongside, left following, right preceding, right alongside, right following) and their respective states. The prediction process involves analyzing the vehicle's behavior in the input window [LCF-125, LCF-50], where LCF (Lane Change Frame) represents the frame where a lane change occurs, and predicts the maneuver type in the subsequent window [LCF-50, LCF+25] as either no change (0), left change (1), or right change (2) (These scenarios are shown in Figure 1).

Our network comprises two LSTM layers with 64 hidden units each, followed by a fully connected layer for final prediction. The data is balanced across classes and split by vehicle IDs (70/10/20 for train/validation/test) to prevent data leakage. As detailed in Table I, the model is trained using the Adam optimizer with a learning rate of 0.001, batch size of 32, and cross-entropy loss for a maximum of 50 epochs.

TABLE I
LSTM Network Properties

Property	**Value**
Input Sequence Length	75 timesteps
Number of Features	54
LSTM Layers	2
Hidden Units per Layer	64
Optimizer	Adam
Learning Rate	0.001
Batch Size	32
Maximum Epochs	50
Loss Function	Cross-entropy

B. *Uncertainty Quantification Methods*

To assess the uncertainty associated with our lane change predictions, we employ three complementary methods: Gaussian noise injection, Dropout noise, and deep ensembles. Gaussian noise injection simulates the inherent noise in sensor measurements and state estimation that occurs in real-world driving scenarios. Dropout noise simulates the effects of missing or unreliable data due to temporary sensor failures or occlusions. Deep ensembles, comprising multiple models with different random initializations, capture the variability in model parameters and provide insights into the model's epistemic uncertainty [6].

V. Experimental Results

This section presents the experimental results of our proposed lane change prediction model, including the evaluation of its performance on the highD dataset and the analysis of uncertainty quantification methods.

A. *Dataset*

In this study, we utilize the highD dataset [16], a comprehensive collection of naturalistic vehicle trajectories captured on German highways using a drone. The dataset's aerial perspective eliminates occlusions commonly encountered with onboard sensors, enabling accurate tracking of all vehicles within the recorded area. This comprehensive view of the traffic environment allows us to capture the intricate interactions between vehicles and their surroundings, which is crucial for understanding the factors that influence lane change decisions.

Specifically, the highD dataset comprises over 11 hours of recordings captured at 25 frames per second, resulting in a total of 110,000 frames. These recordings cover six different highway sections with varying traffic densities and road geometries. The dataset includes over 11,500 vehicles, with detailed trajectories and annotations for each vehicle. This rich dataset allows us to investigate the impact of various factors, such as vehicle type, surrounding traffic conditions, and road geometry, on lane change behavior and uncertainty.

B. *Base Model Performance*

The LSTM model demonstrates strong performance across all three classes in the lane change prediction task. As shown in Table II, the model achieves an overall accuracy of 97.10% with balanced performance across classes. For the No Change class, the model achieves 95.38% accuracy with high precision (96.15%) and recall (95.38%), resulting in an F1-score of 95.76%. The Left Change predictions show slightly better performance with 97.62% accuracy and 97.47% precision, yielding an F1-score of 97.54%. The Right Change class exhibits the best performance with 98.26% accuracy and 97.64% precision, achieving an F1-score of 97.95%.

The class distribution in the test set is well-balanced, with 628 No Change, 671 Left Change, and 632 Right Change instances. This balanced distribution, combined with the high performance metrics across all classes, indicates that the model has successfully learned to distinguish between different lane

TABLE II
PERFORMANCE METRICS OF BASE LSTM MODEL

Class	Accuracy	Precision	Recall	F1-Score
No Change	0.9538	0.9615	0.9538	0.9576
Left Change	0.9762	0.9747	0.9762	0.9754
Right Change	0.9826	0.9764	0.9826	0.9795
Overall Accuracy: 0.9710	Macro F1: 0.9709			

change behaviors without class-specific biases. The confusion matrix visualization (Figure 2) further demonstrates the model's robust performance, with minimal confusion between classes and strong diagonal dominance.

C. Gaussian Noise Injection

Our analysis begins with Gaussian noise injection, where we add random perturbations drawn from a normal distribution with zero mean and standard deviation $\sigma = 0.1$ to the input features. This standard deviation value was chosen to simulate realistic sensor measurement variations while maintaining the signal's integrity. The results show that our model exhibits strong robustness to this level of noise, with mean standard deviations of 0.0168, 0.0097, and 0.0084 for No Change, Left Change, and Right Change predictions respectively. These low uncertainty values indicate that the model's predictions remain stable under typical sensor noise conditions, suggesting reliable performance in real-world scenarios where sensor measurements may contain noise.

D. Dropout Noise

The dropout analysis, performed with a dropout probability $p = 0.1$ (meaning each input feature has a 10% chance of being zeroed out during inference), reveals higher uncertainty levels compared to Gaussian noise injection, with mean standard deviations of 0.1094, 0.0668, and 0.0805 for the three classes. This increased uncertainty suggests that the model is more sensitive to missing or corrupted input features than to noisy measurements. The higher standard deviations,

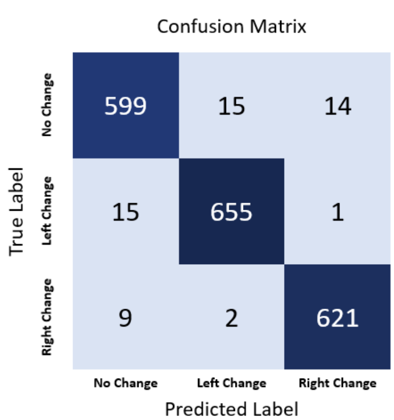

Fig. 2. Confusion matrix showing the distribution of predictions across the three classes. The diagonal elements represent correct predictions, while off-diagonal elements show misclassifications.

TABLE III
UNCERTAINTY ANALYSIS RESULTS

Method	Statistic	No Change	Left Change	Right Change
Gaussian	Mean	0.0168	0.0097	0.0084
	Maximum	0.5189	0.5207	0.4249
	Deviation	0.0539	0.0379	0.0415
Dropout	Mean	0.1094	0.0668	0.0805
	Maximum	0.5379	0.5366	0.5443
	Deviation	0.1398	0.1207	0.1367
Ensemble	Mean	0.0449	0.0305	0.0204
	Maximum	0.4851	0.4879	0.4717
	Deviation	0.0892	0.0788	0.0583

particularly for the No Change class (0.1094), indicate that the model relies on a complete feature set for confident predictions.

E. Deep Ensemble

The ensemble method, consisting of $M = 5$ independently trained LSTM models with identical architectures but different random initializations, demonstrates moderate uncertainty levels, with mean standard deviations of 0.0449, 0.0305, and 0.0204 for No Change, Left Change, and Right Change respectively. These values, falling between those of Gaussian and Dropout methods, suggest that different model initializations produce consistent but not identical predictions. The relatively lower uncertainty for Right Change predictions (0.0204) indicates particularly stable ensemble agreement for this class.

F. Comparative Analysis

To provide a comprehensive view of model uncertainty across different methods, we present a comparative analysis of all three approaches. Figure 3 visualizes the uncertainty distributions, while Table III presents detailed statistics.

Each uncertainty quantification method reveals distinct characteristics of our model's behavior. Gaussian noise injection demonstrates the lowest overall uncertainty (mean STD: 0.0084 - 0.0168), suggesting robust predictions under input perturbations. In contrast, Dropout shows the highest uncertainty levels (mean STD: 0.0668 - 0.1094), indicating greater sensitivity to feature removal. The ensemble method presents moderate uncertainty (mean STD: 0.0204 - 0.0449), reflecting consistent predictions across model initializations.

A notable pattern across all methods is that the No Change class consistently shows higher uncertainty levels, while Right Change predictions typically demonstrate lower uncertainty. This suggests that the model is more cautious in predicting lane-keeping behavior, possibly due to the subtle differences between actual lane changes and minor lateral movements. The maximum standard deviations are relatively consistent across methods (ranging from 0.4249 to 0.5443), indicating similar bounds on prediction uncertainty under extreme conditions.

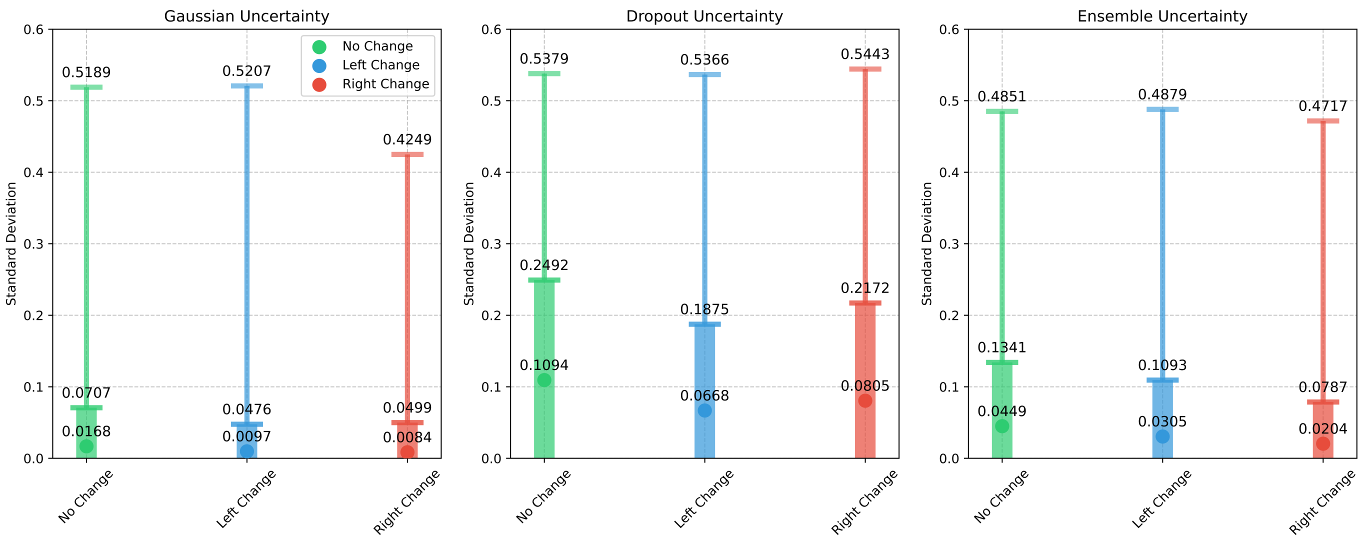

Fig. 3. Comprehensive uncertainty analysis across methods and classes. The thick vertical bars represent mean±std ranges, thin lines extend to maximum values, and dots indicate mean values. Numerical values show maximum, mean+std, and mean from top to bottom.

VI. CONCLUSION

This paper proposed an LSTM-based model for lane change prediction that quantifies uncertainty. The model achieved high accuracy (97.10%) on the highD dataset. Our analysis of Gaussian noise injection, Dropout noise, and deep ensembles revealed varying levels of uncertainty, with mean standard deviations of 0.0084 to 0.1094, respectively. This highlights the model's sensitivity to different types of input perturbations and underscores the importance of uncertainty quantification for robust lane change prediction, enabling more informed decision-making in ADAS and safer autonomous driving.

REFERENCES

[1] V. Leonhardt and G. Wanielik, "Recognition of lane change intentions fusing features of driving situation, driver behavior, and vehicle movement by means of neural networks," in *Deep neural networks and data for automated driving*. Springer, 2017, pp. 141–157.

[2] M. Bonyani, M. Rahmanian, S. Jahangard, and M. Rezaei, "Dipnet: Driver intention prediction for a safe takeover transition in autonomous vehicles," *IET Intelligent Transport Systems*, vol. 17, pp. 1769–1783, 2023.

[3] A. N. Qasemabadi, S. Mozaffari, M. Rezaei, M. Ahmadi, and S. Alirezaee, "A novel model for driver lane change prediction in cooperative adaptive cruise control systems," in *2023 International Symposium on Signals, Circuits and Systems (ISSCS)*. IEEE, 2023, pp. 1–4.

[4] A. Shamsoshoara, S. B. Salih, and P. Aghazadeh, "Swaptransformer: Highway overtaking tactical planner model via imitation learning on osha dataset," in *2024 IEEE 17th International Conference on Automation Science and Engineering (CASE)*. IEEE, 2024, pp. 1–6.

[5] M. Abdar, F. Pourpanah, S. Hussain, D. Rezazadegan, L. Liu, M. Ghavamzadeh, P. Fieguth, X. Cao, A. Khosravi, U. R. Acharya *et al.*, "A review of uncertainty quantification in deep learning: Techniques, applications and challenges," *Information Fusion*, vol. 76, pp. 243–297, 2021.

[6] J. Gawlikowski, C. R. N. Tassi, M. Ali, J. Lee, M. Humt, J. Feng, A. Kruspe, R. Triebel, P. Jung, R. Roscher *et al.*, "A survey of uncertainty in deep neural networks," *Artificial Intelligence Review*, vol. 56, no. 12, pp. 51 513–51 589, 2023.

[7] M. Hemmatian, A. Shahzadi, and S. Mozaffari, "Uncertainty-based knowledge distillation for bayesian deep neural network compression," *International Journal of Approximate Reasoning*, vol. 175, p. 109301, 2024.

[8] R. Maulik, R. Egele, K. Raghavan, and P. Balaprakash, "Quantifying uncertainty for deep learning based forecasting and flow-reconstruction using neural architecture search ensembles," *Physica D: Nonlinear Phenomena*, vol. 454, p. 133852, 2023.

[9] F. Arnez, H. Espinoza, A. Radermacher, and F. Terrier, "A comparison of uncertainty estimation approaches in deep learning components for autonomous vehicle applications," *arXiv preprint arXiv:2006.15172*, 2020.

[10] A. S. Hashemi, A. Bär, S. Mozaffari, and T. Fingscheidt, "Improving transferability of generated universal adversarial perturbations for image classification and segmentation," *Deep Neural Networks and Data for Automated Driving*, pp. 171–196, 2022.

[11] K. Min, D. Kim, J. Park, and K. Huh, "Rnn-based path prediction of obstacle vehicles with deep ensemble," *IEEE Transactions on Vehicular Technology*, vol. 68, no. 10, pp. 10 252–10 258, 2019.

[12] J. Liu, H. Wang, L. Peng, Z. Cao, D. Yang, and J. Li, "Pnnuad: perception neural networks uncertainty aware decision-making for autonomous vehicle," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 12, pp. 24 355–24 366, 2022.

[13] A. Alizadeh, M. Moghadam, Y. Bicer, N. K. Ure, U. Yavas, and C. Kurtulus, "Automated lane change decision making using deep reinforcement learning in dynamic and uncertain highway environment," in *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*. IEEE, 2019, pp. 975–980.

[14] R. Michelmore *et al.*, "Quantification of uncertainty and its applications to complex domain for autonomous vehicles perception system," *PhD diss., Université Clermont Auvergne*, 2020.

[15] S. Siami-Namini, N. Tavakoli, and A. S. Namin, "The performance of lstm and bilstm in forecasting time series," *2019 IEEE International Conference on Big Data (Big Data)*, pp. 3285–3292, 2019.

[16] R. Krajewski, J. Bock, L. Kloeker, and L. Eckstein, "The highd dataset: A drone dataset of naturalistic vehicle trajectories on german highways for validation of highly automated driving systems," in *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2018, pp. 2118–2125.