Robust Traffic Rules and Knowledge Representation for Conflict Resolution in Autonomous Driving

Kumar Manas[1,2,*], Stefan Zwicklbauer[2] and Adrian Paschke[1]

[1]Freie Universität Berlin, Department of Computer Science and Mathematics
[2]Continental AG, Germany

Abstract
Rules and world knowledge are important priors in decision making in our daily lives, and autonomous driving (AD) systems can exploit them in decision-making and conflict resolution. In this thesis, we investigate the representation and formalization of traffic rules and regulations in natural language text in order to integrate them with a machine learning-based trajectory prediction module for conflict resolution among road users. As a result, the vehicle's long-term trajectory (3 to 5 seconds horizon) can be predicted using hybrid learning, which incorporates both rules and data into the ML models. Only rule-based systems confront difficulties in depicting complicated interactions among multiple traffic scene participants. Learning-based techniques are capable of representing complicated interactions. However, they require a large amount of data, and in many circumstances, generating corner case data (e.g., accident and rule violation for exceptions) is not feasible. Recent research combining rules with data (neurosymbolic model) is an exciting research direction to leverage the best of both worlds for reasoning over traffic scenes. Apart from investigating the impact of the knowledge integration, we will work towards finding optimal traffic rules representation for the hybrid learning tasks and leverage large language models for the automated representation of traffic rules and regulation needed for downstream AD task. The hybrid learning approach will reduce the data dependency, generate the vehicle's traffic rule complaint trajectory, and make the model more generalizable even for corner cases or less representative cases in the datasets due to reasoning capability.

Keywords
Knowledge Representation, Rule Formalization, Reasoning, Natural Language Processing, Autonomous Driving

1. Introduction

Ample data and powerful hardware accelerate recent advances in Artificial Intelligence (AI). The data-driven subsymbolic model produces state-of-the-art results but is data hungry and susceptible to failure in unseen scenarios. Whereas the knowledge-driven symbolic model is data efficient but fails to model complex interactions based on statistical relationships. Recent work that combines these approaches can be a way forward to take advantage of the best of both worlds. This combination is referred to as a hybrid AI model. In the hybrid model, the idea is to achieve robustness and generalization capability of the symbolic model with the statistical relationship of the subsymbolic model. We are interested in the hybrid AI approach

RuleML+RR'22: 16th International Rule Challenge and 6th Doctoral Consortium, September 26–28, 2022, Virtual
*Corresponding author.
kumar.manas@fu-berlin.de (K. Manas)

as it has the possibility of making the model generalizable and operating in a dynamic, unseen, and open world environment. In autonomous driving domain, capturing all dangerous or common scenarios of the traffic participant's interaction is not feasible due to a large number of combinations. Also, in datasets, most of the scenarios are based on typical driving, with a low representation of failure and accident scenarios. A hybrid AI based model has the potential to help us navigate this problem by supplementing data with traffic rules or another knowledge base. Knowledge can be represented in various ways: text, mathematical equations, simulations, images, and logic. Integrating these forms of knowledge is essential, as we need to evaluate the compatibility of knowledge representation with subsymbolic model architecture. There are various possible choices for integration, like as a learning algorithm or with network architecture. Currently, there is not always a clear research consensus regarding the optimal way of knowledge representation and integration methods for a more accurate model.

The goal of this thesis is to assist in the planning of autonomous vehicle motion by ruling out an undrivable or legally infeasible trajectory based on traffic rules. For motion planning, the trajectory of the vehicle needs to be predicted. The trajectory is a state of the vehicle with respect to time along the path. A trajectory needs to consider the traffic signs, rules, and static and dynamic obstacles along the route. Our goal is to develop a knowledge representation and automated reasoning module so that trajectory prediction is more accurate and plausible and confirms traffic rules. There are two broad approaches to trajectory prediction: Traditional model based approach [1][2] and the data driven learning based approach [3][4]. Traditional model-based approaches are based on heuristics, leverage rules, and environmental and physical features (e.g., velocity, direction). They are particularly limited in modeling complex interactions among traffic participants in the traffic scene. Learning based approaches can model these complex interactions and context information with the help of a good amount of data and algorithmic design. Some algorithmic design choices can be the loss function in form of rule or rules integrated into network architecture. However, these models pay less attention to robustness and need abundant training data, which is not feasible in many cases. We will inject traffic rules into the model, dictating vehicle priority. Hybrid model can inject both rules and data into the trajectory prediction module. We will model or represent implicit and explicit traffic rules as a logical and semantic knowledge base to perform automated reasoning over the traffic scenes.

- Explicit or declarative rules are defined and standardized by authorities in the form of rule books or regulations, like the German traffic rule book (StVO) [5]. Therefore, they are often more clear and structured.
- Implicit or procedural rules are not defined in concrete terms, but they are necessary to limit chaos in traffic scenes to smooth traffic flow. The decision is based on context and surrounding information, like signaling another car to pass through as a non aggressive driver in the traffic.

Rules and knowledge need to be modeled so that these rules will later be used for reasoning task over a traffic scene. The reasoning task here is to reason about the traffic situation for the priority of traffic participants. The following section will formalize the objective and goals of this thesis work and an abstract plan to achieve the goals.

2. Objective and Research Questions

The main goal of this project is to research and answer the following questions:

How can we semantically and logically formalize, represent, and automatically reason implicit and explicit traffic rules and regulations for more robustness in trajectory planning algorithms (for autonomous vehicles)?

It is hypothesized that *integrating rules and knowledge into the learning based trajectory prediction module can improve the robustness and generate a more plausible trajectory.*

The research questions originating from the proposed objective are:

RQ1: What are the adequate semantic and logical knowledge representation methods of traffic rules and knowledge for the AD domain?
Knowledge modeling is required for rules defined for a downstream task, such as the interaction among traffic participants at intersections or overtaking. Rule representation is an essential step for modeling. In knowledge modeling, rules will be represented in the semantic and logical form. Knowledge modeling allows us to use knowledge as a base (referred to as a knowledge base) to perform reasoning or integrate it with other downstream tasks. For example, knowledge modeling of the *right before left* traffic rule in Germany.

RQ2: Can we use the language model for traffic rules and background knowledge representation so that this representation provides better or similar results than the existing knowledge representation, which requires manual intervention?
Language models have been trained on large volumes of text data to assign probabilities to sequences of words. In this thesis we use them for automated translation of natural language traffic rules into formal logical format, whereas in context of automated driving formal rules are created in hand crafted fashion for each specific context.

RQ3: How and at which stage should knowledge be infused in the hybrid model for reasoning over traffic scene?
Knowledge modeling helps us to inject knowledge into the models and perform reasoning. Prolog and Answer Set Programming are some of the logic-based approaches to perform automated reasoning. Reasoning can be performed as a separate component in the overall model pipeline or combined with knowledge representation or an existing data-driven model. The reasoning component will be used to resolve conflicts over traffic scenes, such as between two cars, in which one will have priority over another. This work will investigate how to model traffic rules and how to reason using them as priors or constraints.

RQ4: What are the possible tradeoffs and scope of improvement when we use the rule as a pre-condition for learning based trajectory prediction algorithm vs. when we use the rule as a conformity check of trajectories?

3. Relevance

This thesis aims to increase the robustness of the autonomous driving system by assisting the long-term trajectory prediction module. The long-term trajectory predicts the position of traffic participants over the next 3 to 5 seconds [6]. This long-term prediction helps to better plan the trajectory taken by an autonomous vehicle. By validating and selecting generated trajectories based on rules and commonsense knowledge, we can discard non-plausible trajectories. This removal of trajectory will increase the confidence score or probability of remaining plausible trajectory so that long-term trajectory can be planned robustly.

4. Related Work

Traffic rules are often based on specific situations and need a clear structure for models to understand, such as keeping a safe distance. We need to define the safe distance in meters. They need to be defined in a structured way so that they can be integrated into downstream trajectory prediction tasks.
There are multiple ways to represent traffic rules as knowledge. An ontology is used to represent concepts and relations of traffic situations and metadata of sensor data [7][8]. Furthermore, first order [9] or higher order logic [10] is also used to represent knowledge of traffic rules. Conversion of textual rules to logic rules (e.g., using linear temporal logic) is an active research area, as logical language removes the semantic ambiguity of textual rules. However, such conversion creates an additional burden for knowledge base creation, and many intricate details of complex traffic participant behavior are challenging to capture. Work by [11] formalized the traffic rules as temporal logic and evaluated their work on a public dataset. However, their work is limited to simple interactions such as a straight highway and ignored regulatory signs such as lane markings and other informative signs, as well as the right of way rule. Nevertheless, their work can be extended to make it more generic and scenario specific use case agnostic. [12] compiled interstate traffic rules in temporal logic and showed the creation of rules with predicates, functions, and propositions but they did not model rules for intersections and they assumed that driving lanes are separated.
The use of a large language model(LLM) in text and language understanding is quite well researched and produced state-of-the-art results. Large pre-trained models like BERT [13], T5 [14], GPT-3 [15] can be used for translation, text understanding and visual question answering. [16] worked toward converting natural language text to signal temporal logic using the Transformer [17] and [18] used the LSTM-based architecture for English sentences to first-order logic. However, to the best of the author's knowledge, these models and work have not yet been deployed for safety-critical systems such as automated formalization of traffic rules from natural language text.
The work discussed above for traffic rule formalization are not robust because they are not expandable to unseen or new scenarios where conflict needs to be resolved among traffic participants. The work of [19] attempts to integrate traffic rules and symbols in their work. Similarly, work in [20] manages static rules which can be easily mapped to constraints, it is not expandable for more complex behavioral rules with multiple agents or participants in scenarios. This work

and other work for knowledge integration achieve better performance, but they consider mostly traffic signs and symbols rather than multi-agent complex interaction based on the traffic rules. As multiple interactions with multiple traffic signs and rules require fast resolution of conflicts, rules cannot be explicitly defined for all such combinations. As a result, for unseen scenarios or situations where rule syntax cannot be directly mapped to traffic scenarios, we should combine them with data-driven models.

Introducing rules in logical format into the neural network are proposed as a solution to overcome some of the problems associated with less data and make model more robust against new unseen scenarios. Logic tensor network [21] and Logic neural network [22] are some of the work in this direction which introduced rules into the neural network. Work by [23] attempts to introduce symbolic knowledge in deep learning model. Their work achieves this by introducing logical constraints into the loss function for the classification problem.

Recent work by [24] combined a learning-based model with logical rules of traffic regulations in signal temporal logic (STL) in the context of motion planning of autonomous vehicles. In this work, symbolic knowledge is integrated as priors and is shown to be effective in generating better trajectory prediction. This work is interesting, but limited to specific rules related to the center line and safe distance and can be further researched for more rules, such as those related to controlled and uncontrolled intersections. Furthermore, this work needs to be further scaled and researched to understand the complexity of the rules, for example, in situations involving intersections of city road traffic. [25] implemented injection of knowledge into the data-driven trajectory prediction module, but their knowledge is limited to physics-constrained value and perception information is used as knowledge instead of traffic rules. Nevertheless, these two works can be interesting to explore further in terms of scaling complex rules or architecture evaluation. The thesis work will research to scale limitations of these baseline works in terms of rule integration, in such a way that the integrated model should be able to model complex traffic participants interaction. Further, they will be evaluated with the use case scenarios (e.g., overtaking and road intersections) available from open public datasets such as Commonroad [26] and with hand crafted scenarios simulating different ways of overtaking or road intersections.

5. Proposed Work and Experiment Planning

The thesis work will have two modules that closely interact: The first step is to represent knowledge, traffic rules, and regulations. After knowledge representation, they will be used to perform conflict resolution over traffic scenes. Traffic scene representation and parsing will be prerequisite for both modules. As limitation of information extracted from the traffic scene will limit our scope of rule application over traffic scene and rule book creation. Fig. 1 shows the main general ideas and components with their interconnection. The trajectory prediction module shown in the illustration is not the focus of our work, but we will use them to evaluate our work and see if we can improve their performance. We will be using an existing state-of-the-art trajectory prediction module. Modules shown with the same color indicate closer integration than others. The formalization of traffic rules from textual rules is part of the knowledge representation module, which will leverage the information generated from the traffic scene and the trajectory prediction module to perform reasoning over the traffic situation. Reasoning in our case is to first find the presence of conflicts among trajectories and then to resolve the conflicts among traffic participants. Among multiple interacting participants, their priority can be resolved at a specific instant based on rules and commonsense knowledge. More information about these modules is provided below along with a rough outline of work and planned experiments.

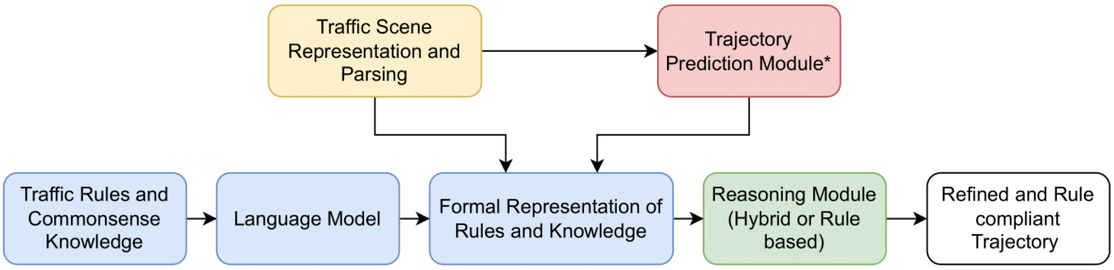

Figure 1: Architecture diagram of modules and their interactions.

5.1. Knowledge Representation and Integration

Based on traffic scenarios such as overtaking, lane merge, and intersections, we are modeling the rules required for these scenarios. Therefore, this stage will consist of knowledge representation of traffic rules (e.g., StVO) for our traffic scenarios.

- Initial evaluation for methods of knowledge integration and its role in trajectory prediction.
- Evaluate data structure for knowledge representation, which allows automated reasoning.
- Comparison of different architecture for knowledge representation and reasoning.

Current research has used more definite logical rules, but we will attempt to use natural language text directly. This component of the thesis will use a large language model (LLM) like T5 [14] and GPT-3 [15] which allow us to learn the meaning behind the traffic rules and can help generate the formal representation of traffic rules. Due to the safety requirements of autonomous driving, human verification of the formal representation of rules will be required for testing purposes. We performed some initial experiments with a language model to convert natural language text into first-order logic. These results are initial evaluations without fine-tuning these data-hungry language models using traffic and driving rules text. However, they can give us hints related to the capability of these models to formalize traffic rule knowledge.

Prompt: A two-lane expressway is a road with only one lane in each direction and usually no metallic median barrier between.

Output : P = two-lane expressway; Q = road with only one lane in each direction; R = no metallic median barrier between; $P(x) \leftrightarrow (R(x) \wedge Q(x))$

As seen in this example, our representations are technically correct. However, we had some failure cases where no fine-tuned language model failed to represent the natural text to logic due to missing quantifiers and variables or some problems with the wrong assignment. Also,

it can be seen that there is a scope for more abstract predicate generation from the natural language text. We are in the process of fine-tuning the language models for the autonomous driving domain, as then they can provide more accurate representation due to more specific training targeted to a domain.
This segment of work is connected to RQ1 and RQ2 regarding representation types and the use of language model for the same.

5.2. Traffic Scene Parsing and Trajectory Prediction Module

We would like to explain some terminology used in this paper. *Path* tells us how to reach point A to B without considering the interaction that occurs along the route and speed. *Trajectory* has a notion of time and speed, and it is a vehicle state with respect to time. In the trajectory prediction for traffic participants, usually, multiple trajectories are predicted for a single vehicle. This is because there are multiple possibilities due to the uncertainty attached to the other traffic participant's behavior. *Traffic Scene* consists of stationary and moving elements in the scene; they usually consist of road segments and traffic participants like vehicles, pedestrians, and traffic control infrastructure. It also contains the relationship and the relation between the elements mentioned above. *Scenario* is the temporal development of scenes in a sequence. To apply our rules, we need to parse the traffic scene so that our logical system can relate to the traffic scene before its usage. Many AD datasets provide an API to access the map data and respective locations of traffic participants. So, we can use them to extract and parse the information about the scene. Another module shown in Fig. 1 is the trajectory prediction module, which generates the trajectory of vehicles based on their past behavior, velocity, acceleration, etc. These two modules output will be combined with formalized knowledge mentioned in Sec. 5.1 to be processed by the reasoning module.

5.3. Reasoning Over Traffic Scene

Represented knowledge will be used for reasoning over traffic scenes. Reasoning needs to be performed when conflicts are detected among the trajectories. Conflicts are resolved for trajectories generated by the trajectory prediction module against the traffic rule representation from the earlier module. So in a sense, the rules will act like prior for more robust trajectory generation. The reasoning module could consist of Answer Set Programming(ASP), ProbLog, or neurosymbolic approaches like Logic tensor network [21] or DeepProbLog [27]. Reasoners like ASP and its variants are powerful and data efficient, but they face issues when uncertainty is involved, or rules cannot be directly applied due to incompleteness. Therefore, we will evaluate the neurosymbolic model to perform reasoning based on rules and data. These evaluations can provide us with an opportunity to answer our RQ3, regarding where we should infuse our knowledge. Subsection 5.3.1 explains it more thoroughly. Fig. 2 shows this idea of the hybrid model, where the AD dataset and formalized knowledge as rules will be combined. This hybrid or neurosymbolic model integrates the symbolic representation homogenously into the neural network, so that the model can learn from data and rules simultaneously. We hope that by such a combination we can get the best of both the symbolic world, which is data-efficient and can handle unseen situations, and the subsymbolic world, which manages complex relationships and interactions. Right now, rules in first order logic (FOL) form is our initial choice, considering the recent development of the logic-based neural network [21] [22], which are more adapted toward FOL. Furthermore, FOL is more expressive compared to horn or propositional logic.

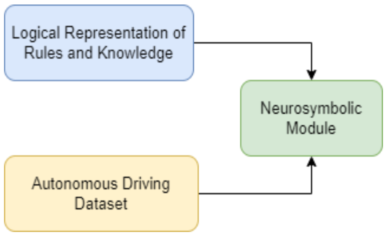

Figure 2: Illustration of the hybrid model concept.

5.3.1. Scaling Representation and Reasoning

We wanted to make knowledge representation and reasoning as generic as possible. This is currently missing from most recent work in this domain. During this phase, we will adapt our work to this purpose (generalization capability) and evaluate it on increasingly complex scenarios involving more traffic participants or multiple rules that apply at the same time apart from earlier mentioned use cases. This includes evaluating new unseen scenarios, where the model needs to perform reasoning based on the rules. We believe that integrating a rule-based module with a data-driven learning-based module is the way to achieve this robustness and plausible trajectory. Neurosymbolic approaches like logic neural network [28] leverage first order rules to perform reasoning tasks, and it would be interesting to evaluate the use case of trajectory prediction with modifications. These neurosymbolic models have not yet been evaluated for logical language having a temporal aspect such as signal temporal logic (STL). A neural network with STL can be evaluated on the basis of the initial results obtained from the FOL logic and the neural network. Additionally, a subsymbolc neural network is believed to perform a faster reasoning process and decision space search than a logical symbolic system based reasoner. In the learning-based approach, some attempts were made to use rules as a loss function to generate rule complaint trajectory. Ways to integrate a loss function that covers a broad range of traffic rules is an open question, and we will work toward this. This evaluation is also connected to our RQ4, where we evaluate trade-offs and performance in various scenarios.

5.4. Evaluation of the Pipeline

We will evaluate our work with public datasets and specific hand-made scenarios. In this step, we will make use of the trajectory prediction algorithm to evaluate the impact of our work. Once we are more advanced in our work, we will refine and formalize this step, as we will then have a better idea of the overall architecture. Currently, we can think of using scenario-based testing, a safe distance from another vehicle, and the occurrence of rule violation as a starting point of testing coupled with typical trajectory prediction evaluation criteria such as the final displacement error (FDE), the mean absolute error (MAE). These criteria generally measure the deviation of ground truth trajectory values with the predicted trajectory provided by models. For automated knowledge modeling assessment, we use the trajectory prediction metric when using formal logic traffic rules for trajectory prediction compared to the metric when using handcrafted traffic rules, based on the above metric. Evaluation of neurosymbolic model can be achieved by designing scenarios where data are not available and rule by itself perform correct decision making and similarly we can evaluate the scenarios for which exact rule is not available but similar interaction among traffic participants can be found in the dataset. Argoverse [29], Interaction [30] nuScenes [31] are some of the public autonomous driving dataset. These datasets can be used for the evaluation of our work. To evaluate specific use case scenario, we can use a driving simulator like CARLA [32] to create a specific scenario and evaluate the performance of our model.

6. Discussion and Open Questions

Trajectory prediction and planning is a well-researched area. However, there are still a lot of open research areas related to long-term trajectory prediction and planning due to the uncertainty involved in sensor perception to maneuver as road movement is dynamic and evolving. Recent advances of high definition maps make it easier to pre-calculate routes and plans based on road signs and traffic lights with static obstacles, but resolving real time conflict with dynamic objects is a difficult task. We are in the first year of our work and we hope to pose some open questions and discussions that can benefit us.

1. Since formal rule-based methods are deterministic and creating them is time consuming, we want to discuss a possible alternative that might be deterministic, and at the same time, they should not create incompleteness issues. In our view, rules can run into a problem when due to a missing or uncertain predicate or parameter they can be unsolvable.
2. What can be possible approaches to knowledge integration in a learning-based model? For example, some possible ways are: as a loss function, as a conformity check, or in the model architecture. It would be insightful to talk about them together with other possible approaches.
3. When we use rules as a constraint in neural network-based trajectory prediction, how do we tackle multiple rules in the same model? What about the priority among the rules in the neural network? We want to discuss such feasibility with experts.

We are highly interested in discussing these questions regarding our core research proposal and ideas as we are trying to combine both worlds of AI and in highly application oriented domains with real-time and safety requirements.

Acknowledgments

The authors acknowledge the funding provided by the German Federal Ministry for Economic Affairs and Energy within the project "KI Wissen – Automotive AI powered by Knowledge" and Continental AG for the project. Special thanks go to Daniel Bär for feedback, idea generation, and also to Raffael Schön for discussions.

References

[1] A. Houenou, P. Bonnifait, V. Cherfaoui, W. Yao, Vehicle trajectory prediction based on motion model and maneuver recognition, in: 2013 IEEE/RSJ International Conference on Intelligent Robots and Systems, 2013, pp. 4363–4369. doi:10.1109/IROS.2013.6696982.

[2] W. Schwarting, A. Pierson, J. Alonso-Mora, S. Karaman, D. Rus, Social behavior for autonomous vehicles, Proceedings of the National Academy of Sciences 116 (2019) 24972–24978. doi:10.1073/pnas.1820676116.

[3] A. Alahi, K. Goel, V. Ramanathan, A. Robicquet, L. Fei-Fei, S. Savarese, Social lstm: Human trajectory prediction in crowded spaces, in: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016, pp. 961–971. doi:10.1109/CVPR.2016.110.

[4] M. Bansal, A. Krizhevsky, A. S. Ogale, Chauffeurnet: Learning to drive by imitating the best and synthesizing the worst, CoRR abs/1812.03079 (2018). URL: http://arxiv.org/abs/1812.03079. arXiv:1812.03079.

[5] Bundesministeriums der justiz und fur verbraucherschutz, strassenverkehrs-ordnung (stvo), 2013. URL: https://www.gesetze-im-internet.de/stvo_2013/StVO.pdf.

[6] R. Chandra, T. Guan, S. Panuganti, T. Mittal, U. Bhattacharya, A. Bera, D. Manocha, Forecasting trajectory and behavior of road-agents using spectral clustering in graph-lstms, CoRR abs/1912.01118 (2019). URL: http://arxiv.org/abs/1912.01118. arXiv:1912.01118.

[7] L. Zhao, R. Ichise, Y. Sasaki, Z. Liu, T. Yoshikawa, Fast decision making using ontology-based knowledge base, in: 2016 IEEE Intelligent Vehicles Symposium (IV), 2016, pp. 173–178. doi:10.1109/IVS.2016.7535382.

[8] M. Buechel, G. Hinz, F. Ruehl, H. Schroth, C. Gyoeri, A. Knoll, Ontology-based traffic scene modeling, traffic regulations dependent situational awareness and decision-making for automated vehicles, in: 2017 IEEE Intelligent Vehicles Symposium (IV), 2017, pp. 1471–1476. doi:10.1109/IVS.2017.7995917.

[9] A. Karimi, P. S. Duggirala, Formalizing traffic rules for uncontrolled intersections, in: 2020 ACM/IEEE 11th International Conference on Cyber-Physical Systems (ICCPS), 2020, pp. 41–50. doi:10.1109/ICCPS48487.2020.00012.

[10] A. Rizaldi, M. Althoff, Formalising traffic rules for accountability of autonomous vehicles, in: 2015 IEEE 18th International Conference on Intelligent Transportation Systems, 2015, pp. 1658–1665. doi:10.1109/ITSC.2015.269.

[11] K. Esterle, L. Gressenbuch, A. Knoll, Formalizing traffic rules for machine interpretability, in: 2020 IEEE 3rd Connected and Automated Vehicles Symposium (CAVS), 2020, pp. 1–7. doi:10.1109/CAVS51000.2020.9334599.

[12] S. Maierhofer, A.-K. Rettinger, E. C. Mayer, M. Althoff, Formalization of interstate traffic rules in temporal logic, in: 2020 IEEE Intelligent Vehicles Symposium (IV), 2020, pp. 752–759. doi:10.1109/IV47402.2020.9304549.

[13] J. Devlin, M. Chang, K. Lee, K. Toutanova, BERT: pre-training of deep bidirectional transformers for language understanding, CoRR abs/1810.04805 (2018). URL: https://arxiv.org/abs/1810.04805.

[14] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, P. J. Liu, Exploring the limits of transfer learning with a unified text-to-text transformer, Journal of Machine Learning Research 21 (2020) 1–67. URL: http://jmlr.org/papers/v21/20-074.html.

[15] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, D. Amodei, Language models are few-shot learners, CoRR abs/2005.14165 (2020). URL: https://arxiv.org/abs/2005.14165.

[16] J. He, E. Bartocci, D. Nickovic, H. Isakovic, R. Grosu, From english to signal temporal logic, CoRR abs/2109.10294 (2021). URL: https://arxiv.org/abs/2109.10294. arXiv:2109.10294.

[17] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, I. Polosukhin, Attention is all you need, Advances in neural information processing systems 30 (2017). URL: https://arxiv.org/abs/1706.03762.

[18] H. Singh, M. Aggarwal, B. Krishnamurthy, Exploring neural models for parsing natural language into first-order logic, CoRR abs/2002.06544 (2020). URL: https://arxiv.org/abs/2002.06544. arXiv:2002.06544.

[19] A. Best, S. Narang, D. Barber, D. Manocha, Autonovi: Autonomous vehicle planning with dynamic maneuvers and traffic constraints, CoRR abs/1703.08561 (2017). URL: http://arxiv.org/abs/1703.08561. arXiv:1703.08561.

[20] Z. Ajanovic, B. Lacevic, B. Shyrokau, M. Stolz, M. Horn, Search-based optimal motion planning for automated driving, in: 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2018, pp. 4523–4530. doi:10.1109/IROS.2018.8593813.

[21] S. Badreddine, A. S. d'Avila Garcez, L. Serafini, M. Spranger, Logic tensor networks, Artif. Intell. 303 (2022) 103649. URL: https://doi.org/10.1016%2Fj.artint.2021.103649.

[22] R. Riegel, A. Gray, F. Luus, N. Khan, N. Makondo, I. Y. Akhalwaya, H. Qian, R. Fagin, F. Barahona, U. Sharma, S. Ikbal, H. Karanam, S. Neelam, A. Likhyani, S. Srivastava, Logical neural networks, 2020. URL: https://arxiv.org/abs/2006.13155. doi:10.48550/ARXIV.2006.13155.

[23] J. Xu, Z. Zhang, T. Friedman, Y. Liang, G. V. d. Broeck, A semantic loss function for deep learning with symbolic knowledge, 2017. URL: https://arxiv.org/abs/1711.11157. doi:10.48550/ARXIV.1711.11157.

[24] X. Li, G. Rosman, I. Gilitschenski, J. A. DeCastro, C. I. Vasile, S. Karaman, D. Rus, Differentiable logic layer for rule guided trajectory prediction, in: J. Kober, F. Ramos, C. J. Tomlin (Eds.), 4th Conference on Robot Learning, CoRL 2020, 16-18 November 2020, Virtual Event / Cambridge, MA, USA, volume 155 of *Proceedings of Machine Learning Research*, PMLR, 2020, pp. 2178–2194. URL: https://proceedings.mlr.press/v155/li21b.html.

[25] M. Bahari, I. Nejjar, A. Alahi, Injecting knowledge in data-driven vehicle trajectory predictors, Transportation research part C: emerging technologies 128 (2021) 103010. URL: https://www.sciencedirect.com/science/article/pii/S0968090X21000425.

[26] M. Althoff, M. Koschi, S. Manzinger, Commonroad: Composable benchmarks for motion planning on roads, in: 2017 IEEE Intelligent Vehicles Symposium (IV), 2017, pp. 719–726. doi:10.1109/IVS.2017.7995802.

[27] R. Manhaeve, S. Dumancic, A. Kimmig, T. Demeester, L. D. Raedt, Deepproblog: Neural probabilistic logic programming, CoRR abs/1805.10872 (2018). URL: http://arxiv.org/abs/1805.10872. arXiv:1805.10872.

[28] R. Riegel, A. G. Gray, F. P. S. Luus, N. Khan, N. Makondo, I. Y. Akhalwaya, H. Qian, R. Fagin,

F. Barahona, U. Sharma, S. Ikbal, H. Karanam, S. Neelam, A. Likhyani, S. K. Srivastava, Logical neural networks, CoRR abs/2006.13155 (2020). URL: https://arxiv.org/abs/2006.13155. arXiv:2006.13155.

[29] M.-F. Chang, J. Lambert, P. Sangkloy, J. Singh, S. Bak, A. Hartnett, D. Wang, P. Carr, S. Lucey, D. Ramanan, et al., Argoverse: 3d tracking and forecasting with rich maps, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2019, pp. 8748–8757. doi:10.1109/CVPR.2019.00895.

[30] W. Zhan, L. Sun, D. Wang, H. Shi, A. Clausse, M. Naumann, J. Kummerle, H. Konigshof, C. Stiller, A. de La Fortelle, M. Tomizuka, Interaction dataset: An international, adversarial and cooperative motion dataset in interactive driving scenarios with semantic maps, arXiv preprint arXiv:1910.03088 (2019). URL: https://arxiv.org/abs/1910.03088.

[31] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, O. Beijbom, nuscenes: A multimodal dataset for autonomous driving, 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2020) 11618–11628. doi:10.1109/CVPR42600.2020.01164.

[32] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, V. Koltun, CARLA: An open urban driving simulator, in: Proceedings of the 1st Annual Conference on Robot Learning, 2017, pp. 1–16. URL: https://arxiv.org/abs/1711.03938.