Nyberg, T. (2022)
Evaluating Sequential Reasoning about Hidden Objects in Traffic
In: *ICCPS '22: Proceedings of the 13th ACM/IEEE International Conference on Cyber-Physical Systems*

Evaluating Sequential Reasoning about Hidden Objects in Traffic

Truls Nyberg*†
KTH Royal Institute of Technology
Stockholm, Sweden
trulsny@kth.se

José Manuel Gaspar Sánchez*
KTH Royal Institute of Technology
Stockholm, Sweden
jmgs@kth.se

Christian Pek
KTH Royal Institute of Technology
Stockholm, Sweden
pek2@kth.se

Jana Tumova
KTH Royal Institute of Technology
Stockholm, Sweden
tumova@kth.se

Martin Törngren
KTH Royal Institute of Technology
Stockholm, Sweden
martint@kth.se

ABSTRACT

Hidden traffic participants pose a great challenge for autonomous vehicles. Previous methods typically do not use previous observations, leading to over-conservative behavior. In this paper, we present a continuation of our work on reasoning about objects outside the current sensor view. We aim to demonstrate our recently proposed method on an autonomous platform and evaluate its reliability and real-time feasibility when using real sensor data. Showing a significant driving performance increase on a real platform, without compromising safety, would be a significant contribution to the field of autonomous driving.

KEYWORDS

Autonomous Vehicles, Hidden Traffic Participants, Traffic Occlusions, Motion Planning, Reachability Analysis, Safe Autonomy

1 INTRODUCTION

To ensure safety and gain acceptance, autonomous vehicles (AVs) must minimize the risk of causing accidents [1, 2]. For this, also hidden objects need to be considered, such as a cyclist hidden behind a large vehicle, or an object on the road occluded by a building. However, assuming unseen regions to always be occupied can be overly conservative and may even lead to the "freezing robot" problem where the AV stops and deems all future paths unsafe [3].

In the scenario depicted in Figure 1, the AV should be able to safely turn left. The dashed region behind the truck is currently not seen from the sensors on the bus, however, the system should be able to use previous observations to conclude that no traffic participant could have reached this region without severely violating the traffic rules (e.g. by making a U-turn behind the truck).

*Both authors contributed equally to this research.
†Also with Scania CV AB, Research and Development, 151 87 Södertälje, Sweden..

This work was partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation. This research has been carried out as part of the TECoSA Vinnova Competence Center for Trustworthy Edge Computing Systems and Applications at KTH Royal Institute of Technology and in addition been partly supported through InSecTT. InSecTT (www.insectt.eu) has received funding from the ECSEL Joint Undertaking (JU) under grant agreement No 876038. The JU receives support from the European Union's Horizon 2020 research and innovation programme and Austria, Sweden, Spain, Italy, France, Portugal, Ireland, Finland, Slovenia, Poland, Netherlands, Turkey". The document reflects only the author's view and the Commission is not responsible for any use that may be made of the information it contains.

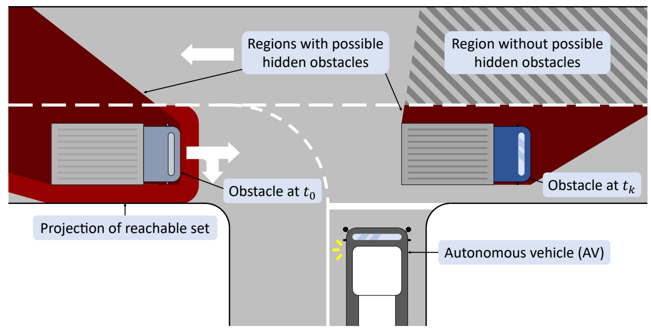

Figure 1: Initially, at t_0, the AV considers that hidden obstacles may occupy all the unseen regions (dark red) behind the truck. However, as new observations are made, unseen regions are concluded free if no hidden object may have reached there between the observations (dashed region at t_k).

Recently, several works have focused on tackling the challenge with currently unseen objects caused by occlusions and range limitations [4–9].

However, all the studies previously mentioned only consider the current observation from the ego vehicle when evaluating where hidden vehicles could be. Reachability analysis is used in [8, 9] together with a planner to generate safe trajectories in scenarios under occlusion. This is done by over-approximating all possible future occupancies of virtual obstacles placed at the current edge of the unseen region. This ensures that any possible hidden obstacle is considered, however, it may lead to considering more obstacles than necessary. By reasoning about where hidden obstacles could have reached in between each observation, this over-conservative anticipation can be reduced.

In [10], an approach was implemented to track regions that could be occupied by pedestrians. Similarly in [11], sequential reasoning, taking traffic rules into account, was presented for any possible hidden obstacle. Monte Carlo simulations showed that the method greatly improves the AVs' performance in terms of time to complete scenarios, without compromising safety. In this work, we aim to validate the real-time feasibility of this approach on a real platform. A poster showing preliminary results will be presented at the International Conference for Cyber-Physical Systems, 2022.

2 METHOD

The method used for modeling and computing possible hidden obstacles is described in detail in [11]. Iteratively, the method computes the region, $\mathcal{P}_t$, where possible hidden obstacles could be at the current time. This is done by computing the reachable set of the previous region, $\mathcal{R}(\mathcal{P}_{t-1})$, and then intersecting it with the area outside the current Field of View, FoV_t. The reachable set includes the region where any obstacle from $\mathcal{P}_{t-1}$ could have reached in one time-step given all constraints, e.g., maximum velocity or driving direction of the lane. For this reason, it is assumed that a hidden obstacle must be both outside our current field of view and within the reachable set of the previously computed region.

A schematic example of the method is seen in Figure 1. The reachable set $\mathcal{R}(\mathcal{P}_{t_0})$ is visualized in bright red, and the region with possible hidden obstacles, $\mathcal{P}_t$, is depicted in dark red. The dashed area highlights a region in the state space that is outside the current Field of View, FoV_{t_k}, yet still concluded free. The method requires a representation of the area detected as free by the sensors at every time-step, referenced as Field of View, FoV_t. The method also requires a map of the area including speed limit information and a graph description of the lane network with precise boundaries, driving directions, and other traffic rules. The set of possible hidden obstacles is computed at every new time-step. For a motion planner to utilize this information, a prediction of the possible hidden obstacles also needs to be provided. An extension of the method proposed in [11] could include constant monitoring of other vehicles' possible violations of the considered traffic rules.

3 EVALUATION

The experiments will be conducted with a modified Scania Citywide battery electrical bus, seen in Figure 2. The bus is equipped with sensors, computing units, and a fully autonomous driving software stack in which the proposed method will be implemented.

This platform provides a grid map representation of the environment. To ensure safety, we will over-approximate the reachable sets of all possible hidden obstacles between observations on this grid representation, as done in [10] for pedestrians. An investigation of real-time requirements and analysis of the method's computational complexity will also be required [12]. For motion predictions with longer time horizons, the reachability predictions will similarly be evaluated together with the provided platform's algorithms for predictions. Instead of the simple sample-based motion planner in [11], a general-purpose planner will be used, similar to the one described in [13]. Once in operation, we aim to validate that the algorithm can handle measurement uncertainties and does not introduce any delays or other artifacts that influence the stability and robustness of the system.

As in [11], we aim to compare the implemented sequential method with a baseline method that only uses the current observations. To compare the methods, experiments will first be conducted in a simulation environment of the platform and later carried out at Scania's test track in Södertälje, Sweden. High definition maps of the test track generated offline will be used, including speed limits and precise lane network description. The scenario in Figure 1 will be used to validate that our method can reduce the time for the ego vehicle to complete the scenario, without compromising safety.

Figure 2: A modified Scania Citywide bus, equipped with an experimental research platform for autonomous driving.

4 CONCLUSIONS

This work aims at evaluating a method to compute possible hidden obstacles based on previous observations and compare its performance against a baseline method on a real platform.

The potential of the approach has been shown in simulations, indicating an improvement in performance without a compromise in safety [11]. However, an implementation of the method in a real system will allow a deeper understanding of how uncertainties of physical sensors might affect the method's reliability.

REFERENCES

[1] B. Vanholme, D. Gruyer, B. Lusetti, S. Glaser, and S. Mammar, "Highly Automated Driving on Highways Based on Legal Safety," *IEEE Transactions on Intelligent Transportation Systems*, vol. 14, no. 1, pp. 333–347, 2013.

[2] Aptiv, Audi, Baidu, BMW, Continental, Daimler, FCA, Here, Infineon, Intel, and Volkswagen, "Safety First for Automated Driving," white Paper, 2019.

[3] P. Trautman and A. Krause, "Unfreezing the robot: Navigation in dense, interacting crowds," in *2010 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, pp. 797–803, 2010.

[4] M. Lee, K. Jo, and M. Sunwoo, "Collision risk assessment for possible collision vehicle in occluded area based on precise map," in *2017 IEEE 20th Int. Conf. on Intelligent Transportation Systems (ITSC)*, pp. 1–6, 2017.

[5] M.-Y. Yu, R. Vasudevan, and M. Johnson-Roberson, "Occlusion-Aware Risk Assessment for Autonomous Driving in Urban Environments," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 2235–2241, 2019.

[6] S. Brechtel, T. Gindele, and R. Dillmann, "Probabilistic decision-making under uncertainty for autonomous driving using continuous POMDPs," in *17th Int. IEEE Conf. on Intelligent Transportation Systems (ITSC)*, pp. 392–399, 2014.

[7] P. Schörner, L. Töttel, J. Doll, and J. M. Zöllner, "Predictive Trajectory Planning in Situations with Hidden Road Users Using Partially Observable Markov Decision Processes," in *2019 IEEE Intelligent Vehicles Symposium (IV)*, pp. 2299–2306, 2019.

[8] P. F. Orzechowski, A. Meyer, and M. Lauer, "Tackling Occlusions and Limited Sensor Range with Set-based Safety Verification," in *2018 21st Int. Conf. on Intelligent Transportation Systems (ITSC)*, pp. 1729–1736, 2018.

[9] M. Koschi and M. Althoff, "Set-Based Prediction of Traffic Participants Considering Occlusions and Traffic Rules," *IEEE Transactions on Intelligent Vehicles*, vol. 6, no. 2, pp. 249–265, 2021.

[10] Y. Nager, A. Censi, and E. Frazzoli, "What lies in the shadows? Safe and computation-aware motion planning for autonomous vehicles using intent-aware dynamic shadow regions," in *2019 Int. Conf. on Robotics and Automation (ICRA)*, pp. 5800–5806, 2019.

[11] J. M. Gaspar Sánchez, T. Nyberg, C. Pek, J. Tumova, and M. Törngren, "Foresee the Unseen : Sequential Reasoning about Hidden Obstacles for Safe Driving." 2022, Preprint available online, diva2:1635726, 2022.

[12] M. Törngren, "Fundamentals of Implementing Real-Time Control Applications in Distributed Computer Systems," *Real-Time Systems*, vol. 14, no. 3, pp. 219–250, 1998.

[13] S. Rosbach, V. James, S. Großjohann, S. Homoceanu, and S. Roth, "Driving with Style: Inverse Reinforcement Learning in General-Purpose Planning for Automated Driving," in *2019 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, pp. 2658–2665, 2019.