Counterfactual Policy Evaluation for Decision-Making in Autonomous Driving

Patrick Hart[1] and Alois Knoll[2]

Abstract— **Learning-based approaches, such as reinforcement and imitation learning are gaining popularity in decision-making for autonomous driving. However, learned policies often fail to generalize and cannot handle novel situations well. Asking and answering questions in the form of "Would a policy perform well if the other agents had behaved differently?" can shed light on whether a policy has seen similar situations during training and generalizes well. In this work, a counterfactual policy evaluation is introduced that makes use of counterfactual worlds — worlds in which the behaviors of others are non-actual. If a policy can handle all counterfactual worlds well, it either has seen similar situations during training or it generalizes well and is deemed to be fit enough to be executed in the actual world. Additionally, by performing the counterfactual policy evaluation, causal relations and the influence of changing vehicle's behaviors on the surrounding vehicles becomes evident. To validate the proposed method, we learn a policy using reinforcement learning for a lane merging scenario. In the application-phase, the policy is only executed after the counterfactual policy evaluation has been performed and if the policy is found to be safe enough. We show that the proposed approach significantly decreases the collision-rate whilst maintaining a high success-rate.**

I. INTRODUCTION

Learning-based approaches become ever more prevalent in the decision-making community for autonomous driving. Often, deep neural networks (DNNs) are used in various approaches to learn behaviors or to predict other traffic participants. Popular approaches for learning behaviors include reinforcement learning [1, 2], imitation learning [3–5], and other methodologies [6, 7]. These approaches often share one common characteristic — they use DNNs as function approximators for their policy. However, DNNs often tend to overfit [8, 9] or do not generalize well to novel situations [10, 11]. Executing learned policies in safety-critical applications, such as autonomous driving without any estimate of their performance poses significant threats. For example, would the learned policies still perform well if the leading vehicle would suddenly accelerate or decelerate?

In this work, a counterfactual policy evaluation (CPE) is introduced that evaluates the performance of a learned behavior policy π prior to its execution in the actual world using counterfactual worlds. In a counterfactual world, the behavior policy of at least one other vehicle is replaced and, thus, non-actual behaviors are introduced. Using these worlds, counterfactual reasoning can be performed to answer

[1]fortiss GmbH, An-Institut Technische Universität München, Munich, Germany
[2]Alois Knoll is with Robotics, Artificial Intelligence and Real-time Systems, Technische Universität München, Munich, Germany
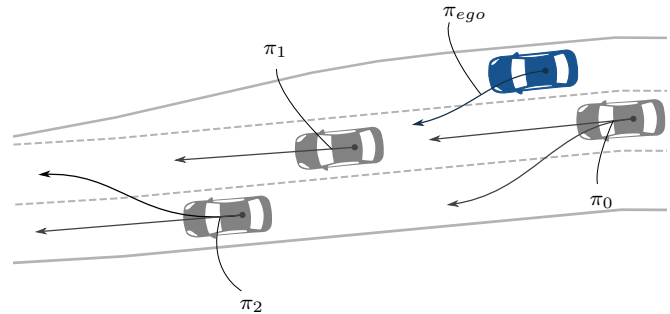

Fig. 1: The behaviors $\pi_{1:N}$ of other traffic participants are not observable by the ego vehicle (depicted in blue) resulting in an uncertain environment. Vehicles can suddenly change their driving behavior, such as braking or changing the lane. Would the policy still perform well and be safe if the other vehicle changed their behavior?

questions, such as: Would a policy perform well if the other vehicles had behaved as in the counterfactual world? Performing counterfactual reasoning is especially important in applications, such as autonomous driving, where the behavior of others is not directly observable and, thus, uncertain.

To perform CPE, counterfactual worlds are generated by replacing the behavior policies of vehicles surrounding the ego vehicle. To determine and better estimate the direct effect of replacing behavior policies on the policy under test (PUT), the replaced behaviors are chosen to be independent of their surroundings — they do not interact with or react to their surroundings. By simulating the counterfactual worlds into the future and by evaluating the evolution of these, we can obtain estimates on how the PUT would perform if the other vehicles behave as in the counterfactual worlds. If the PUT has seen similar situations as in the counterfactual worlds during training or if it generalizes well, it should be able to handle all counterfactual worlds sufficiently well. Thus, in this work, the PUT is only deemed safe enough to be executed if there are no or only a few collisions in the counterfactual worlds. We define an upper collision-rate threshold that defines whether the PUT shall be executed in the actual world posterior to its CPE. Additionally, by using CPE, we can extract causal relations and make the influence of changing other vehicle's behavior policies become evident. Especially when using black-box approaches, such as DNNs, causal relations and how the PUT reacts to changes in its surroundings are important to observe. For example, if the PUT just had memorized certain situations during

training it would not react to changes in its environment and most probably would fail to generalize to and handle novel situations.

To validate the proposed approach and to empirically show the merits of performing CPE, we use a lane merging scenario. First, we learn a policy π_{ego} using an actor-critic reinforcement learning algorithm. In the application-phase, the learned policy π_{ego} becomes the PUT and is only being executed after the CPE has been performed and is deemed safe enough to be executed. Otherwise, the ego vehicle is controlled by a conventional lane-following behavior policy. We show that performing CPE prior to the execution in the actual world of the learned policy significantly decreases the collision-rate whilst maintaining high performance in terms of reaching the vehicle's goal.

This work is further structured as follows: Section II gives an overview of related work in counterfactual reasoning. Section III outlines the proposed approach in detail. In Section IV, the results and findings are presented, and finally, a conclusion is given.

II. RELATED WORK

In this section, we provide a brief overview of counterfactual reasoning and counterfactual decision-making.

The term counterfactual reasoning has predominantly been coined by Judea Pearl [12] who also introduced a mathematical framework (do-calculus) for counterfactual reasoning [13]. Do-calculus aims to describe the human capacity to reason about counterfactual outcomes of past experiences with the goal of "mining worlds that could have been". In safety-critical applications, such as autonomous driving do-calculus often is the only option to obtain certain distributions. For example, the probability distribution $P(\text{collision}|\pi_{other} = \pi_0)$ cannot directly be collected in the actual world. Having a joint probability distribution, causal assumptions, and using do-calculus this distribution can be inferred a-posteriori. In this work, contrary to using a joint probability distribution and causal assumptions, we can directly collect these probability distributions in simulation and do not have to infer these.

Another approach in the literature is counterfactual explanations where causal situations in the form of "If X had not occurred, Y would not have occurred" are formulated and answered [14]. In the context of autonomous driving, this is can be re-formulated as e.g. "If the other vehicle had not accelerated, a collision with the ego-vehicle would not have occurred". Using the in this work defined counterfactual worlds, we can infer which event lead to a certain outcome – e.g. if a collision could be avoided if the other vehicle had decelerated or accelerated.

In literature, there exist various approaches that make use of counterfactuals for decision-making. Buesing *et al.* [15] propose a counterfactual guided-policy search reinforcement learning algorithm. In their work, they leverage a model to consider alternative outcomes and, thus, increase the algorithms sample efficiency on required experiences. Isele [16] consider counterfactual behaviors in the prediction by
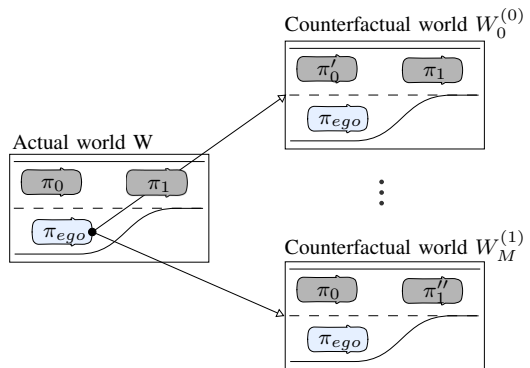

Fig. 2: Counterfactual worlds $W_i^{(\cdot)}$ are created by mirroring the actual world W and by exchanging the vehicle's policies. In the counterfactual world $W_0^{(0)}$ the policy of vehicle v_0 has been exchanged and in counterfactual world $W_M^{(1)}$ the policy of vehicle v_1.

spanning an intention tree. Counterfactual regret minimization (CRM) has been widely used for finding best response strategies in multi-player games [17, 18]. In CRM, the choice of future actions is based on the regret the agent accumulates of not having chosen specific actions and self-play.

III. APPROACH

This section describes the counterfactual policy evaluation (CPE) algorithm in detail. Using CPE, the performance of executing a policy prior to execution in the actual world is being evaluated using counterfactual reasoning. First, we introduce all necessary tools and notation as a basis and then outline the full structure of the algorithm.

In each world W, there is a set of vehicles $V = (v_0, v_j, \ldots, v_N)$ with N being the number of vehicles. Each vehicle v_j has a policy π_j that specifies its behavior. The full parameterization of a world W is given by a set of vehicles paired with their respective policies π_j. A policy π_j outputs an action a for each world state W. A full world state is given by $W = ((v_0, \pi_0), (v_{ego}, \pi_{ego}), \ldots, (v_N, \pi_N))$ with one of the N vehicles being the ego vehicle v_{ego} and that is being controlled by the ego policy π_{ego}. The ego vehicle's policy π_{ego} is defined to be the policy under test (PUT). We neglect the map and environment information in the notation as these are assumed to be the same in all worlds.

In a counterfactual world, the policy of at least one vehicle is replaced by an independent policy. To determine the direct effect of replacing an individual policy on its surrounding, the replaced behavior policies are independent – they do not react to or interact with their surroundings. For example, policies that accelerate or decelerate independent of changes in their surroundings. The behavior policies of the K nearest vehicles are being replaced as we assume that vehicles far away do not have a large direct effect on the PUT. These independent behavior policies are assigned from a pre-defined policy pool $\pi_0, \ldots, \pi_M \sim \mathcal{P}$ to vehicles other than the ego vehicle resulting in $K \times M$ counterfactual

worlds. A counterfactual world where the policy of vehicle v_j has been replaced by the independent policy π_i is denoted as $W_i^{(j)}$. We introduce a short form notation $W_{\mathcal{P}}^{(j)}$ for the set of counterfactual worlds that are generated by replacing the behavior policy of vehicle v_j M times. These worlds are called counterfactual worlds, as we assume non-actual behavior policies for other vehicles. This allows us to answer counterfactual questions, such as: How would the PUT perform if the other vehicles had behaved as in the counterfactual world?

To perform counterfactual reasoning, the set of counterfactual worlds $W_{\mathcal{P}}^{(j)}$ is forward simulated for a pre-defined time-horizon T_{ct} starting at the current world time t_w. From the forward simulation trajectories $\tau_{\mathcal{P}}^{(j)}$ of tuples in the form of $\langle observation, terminal, reward, info \rangle$ are collected within the time-range $[t_w, t_w + T_{ct}]$ where the policy of vehicle's v_j has been replaced by the independent behavior policy π_i.

Using the recorded future trajectories $\tau_{\mathcal{P}}^{(j)}$ of the counterfactual worlds, conditional probabilities $P(X|W_{\mathcal{P}}^{(j)})$ for the PUT can be obtained. The conditional probability $P(X|W_{\mathcal{P}}^{(j)})$ for exchanging vehicle's v_j policy is given by

$$P(X|W_{\mathcal{P}}^{(j)}) = \frac{1}{M} \sum_{i=0}^{M} \text{I}_{-}(\tau_i^{(j)}) \tag{1}$$

with a boolean indicator function $\text{I}_{-}(x)$ returning whether a pre-defined event, such as a colission has occured in the trajectory $\tau_i^{(j)}$.

In safety-critical applications, such as autonomous driving, we especially are interested in the average conditional collisions probability over all counterfactual worlds denoted as $P(C|W_{\mathcal{P}}^{(\cdot)})$. For convenience of notation, we define the average conditional collision probability as

$$P_{\text{C}} := \frac{1}{K} \sum_{j=0}^{K} P(C|W_{\mathcal{P}}^{(j)}) \tag{2}$$

with N being the number of vehicles.

We use the average conditional collision probability P_{C} of the counterfactual worlds as decision criterion on whether the PUT should be executed in the actual world W. We define an upper collision threshold $\rho_{\max}$ up to which the PUT is being executed. If P_{C} is larger, the ego vehicle is controlled by a lane-following model, such as the intelligent driver model (IDM) instead [19]. We assume that the lane-following model is collision-free when driving on a single lane. The full CPE algorithm is outlined in Algorithm 1.

Additionally, due to the replaced behavior policies being indepent of their surroundings, causal relations between the vehicle's policies can be extracted. In this work, we use the mean average displacement (MAD) to measure the distance of the vehicle's states in the counterfactual worlds $W_{\mathcal{P}}^{(j)}$ to the actual world W. The greater the MAD in a counterfactual world is compared to the actual world, the greater the influence of changing the other vehicle's policy is on the PUT. This indicates how well and if the PUT reacts to its

---

Algorithm 1 Counterfactual Policy Evaluation (CPE)

function CPE(World W, PolicyPool $\mathcal{P}$)
    $v_{0:K}$ = NearbyVehicles(W)
    trajectories = []
    for v in $v_{0:K}$ **do**
        $W_{\mathcal{P}}^{(v)}$ = GetCounterfactualWorlds(W, v, $\mathcal{P}$)
        $\tau_{\mathcal{P}}^{(v)}$ = ForwardSimulate($W_{\mathcal{P}}^{(v)}$, T_{ct})
        trajectories.append($\tau_{\mathcal{P}}^{(v)}$)
    end for
    return ToBeExecuted(trajectories, ρ_{max})
end function

---

surrounding and changes that occur in it. For example, if the leading vehicle of the ego vehicle brakes rapidly and the PUT does not react at all, this could be an indicator that the policy does not generalize well and that it did not see similar behaviors during training.

In the next section, we present results and an evaluation showing that the CPE algorithm significantly descreases the colllision-rate and increases the performance of executing a learned policy in the actual world. Additionally, we conduct studies on varying the maximum allowed collision-rate ρ_{max}.

IV. Experiments

In this section, experiments are performed and an evaluation of the counterfactual policy evaluation (CPE) is presented using a highway merging senario. First, we introduce the used simulation framework and the scenario. Next, we outline the actor-critic reinforcement learning approach used for learning lane merging policies. And finally, we apply the CPE to the learned policy and present results and findings.

A. Simulation and Scenario

We use the semantic simulation framework BARK[1] and its machine learning extension BARK-ML[2] for all simulations and to learn the ego vehicle's policy π_{ego}. To demonstrate and evaluate the approach, we chose a lane merging scenario where the ego vehicle's goal is to merge onto the highway. The scenarios are generated using a sampling strategy that samples the initial conditions and the behavioral parameters of the other vehicles according to a pre-defined distribution. All other vehicles are controlled by the ruled-based Mobil moddel. The other vehicles are controlled by a rule-based model — the Mobil model [20]. Further simulation parameters are outlined in the Appendix. The simulation is episodic and ends either if the goal has been reached, a collision has occurred, or the maximum number of steps has been reached. The goal of the ego vehicle is defined on the left lane using a polygonal area (depicted in blue in Figure 3) and the vehicle has to reach this area within a velocity range of $[5m/s, 16m/s]$ and the vehicle angle θ in the range of $[-0.05rad, 0.05rad]$.
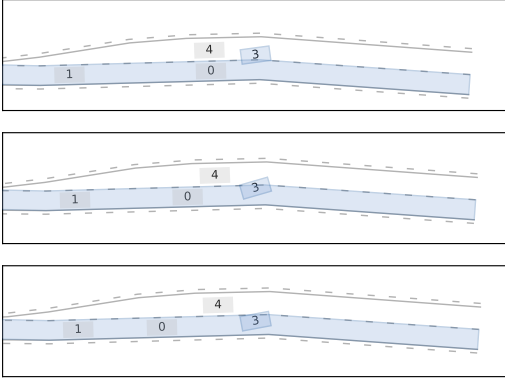

Fig. 3: Three counterfactual worlds $W_i^{(0)}$ are shown where the policy of vehicle v_0 has been replaced. Initially, all these worlds had the same state but evolved differently over time. On the top, vehicle v_0 decelerates, in the middle, vehicle v_0 drives with a constant velocity, and, on the bottom, vehicle v_0 accelerates. The ego vehicle v_3 (depicted in blue) is controlled by the learned policy $\pi_\psi(a|s)$. This figure shows the counterfactual world at time $t_w + 1s$.

Fig. 4: The heatmap shows the mean average displacement of the vehicle's states over the counterfactual worlds $W_{\mathcal{P}}^{(j)}$. Vehicle v_3 is controlled by the learned policy and the heatmap shows it being influenced by all surrounding vehicles. The other vehicles are controlled by a lane-following model and, thus, are only influenced by their leading vehicles on the same lane.

B. Learning Policies

We base learning our policy on [21] and use the soft actor-critic (SAC) [22] algorithm to learn a stochastic policy $\pi_\chi(a|s)$ that is parameterized by the neural network parameters χ for changing lanes. The action space is comprised of the acceleration a and the steering-rate δ. To generate an observation for a single time-step, we use a nearest state observer that concatenates the states of $l = 5$ nearest vehicles (cartesian coordinates, velocity, and vehicle angle). The ego vehicle's state is always in the first position in the observation and the other vehicles follow sorted based on their distance to the ego vehicle. If there are less than five objects the rest of the observation s is filled with zeros. The reward signal r_t for time t is composed of several terms. Reaching the goal is rewarded with a positive reward of 10 and a collision is penalized with a reward of -10. To accelerate the learning process, we further introduce a guiding reward signal that uses the $L2$ norm to the defined goal state (distance, deviation of the vehicle angle, and the deviation to the goal velocity). To achieve more smooth and comfortable driving behaviors, we additionally put a sqared penality term on the actions. All rewards besides the goal and collision term are a factor of 10 smaller to prioritize to have no collisions and reaching the goal state. In this work, we output a normal distribution $\underline{a} \sim \mathcal{N}(\underline{\mu}, \underline{\sigma})$ for the policy π with $\underline{\mu}$ being the means and $\underline{\sigma}$ the standard deviations for the actions. During training, this stochastic policy is being sampled to explore the configuration space and during application, the mean of each action is used. The hyperparameters used for the SAC algorithm are outlined in the Appendix.

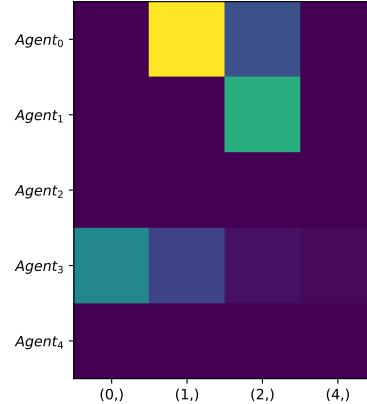After $250,000$ thousand training episodes, a success-rate of 32% and a collision-rate of 48% is achieved. Even after fine-tuning the hyperparameters and a longer training time, collisions persist. Many machine learning approaches cannot guarantee or achieve a collision-rate of zero. Therefore, they are considered to be infeasible for safety-critical situations, such as autonomous driving. To decrease the collision-rate and to obtain better estimates on the performance of a policy π, this work introduces a counterfactual policy evaluation.

C. Counterfactual Policy Evaluation (CPE)

For the CPE, we use a behavior policy pool $\mathcal{P}$ with a constant acceleration policy $\mathcal{B}$. In this work, we use three different parameterizations for the constant acceleration policy $\mathcal{B}$ having the accelerations of $a = [-2, 0, 2]$ and resulting in three behavior policies. Thus, resulting in three counterfactual worlds $(W_0^{(j)}, W_1^{(j)}, W_2^{(j)})$ for each vehicle v_j, where the replaced vehicle decelerates, drives with constant velocity, and accelerates. Further, we simulate each counterfactual world for a pre-defined horizon $T_{ct} = 1s$ from the current world time t_w.

Figure 3 shows the evolution of the counterfactual worlds $W_i^{(0)}$ at the time $t = t_w + 1s$ where the policy of vehicle v_0 has been replaced. Initially, all worlds had the same initial state but then evolved differently due to the policy of vehicle v_0 being replaced by the policies from the behavior policy pool $\mathcal{P}$. Vehicle v_3 is controlled by the learned SAC policy $\pi_\chi(a|s)$ and is also the policy under test (PUT). In Figure 3 on the top, in counterfactual world $W_0^{(0)}$ vehicle v_0 decelerates, in the middle, vehicle v_0 drives with a constant velocity, and on the bottom vehicle v_0 accelerates. In each counterfactual world, the evolution over time and the reaction of the PUT to the environment differs significantly.

As can be seen, the learned policy $\pi_\chi(a|s)$ is able to handle all of the counterfactual worlds well and does not cause any collisions. Thus, resulting in the conditional collision
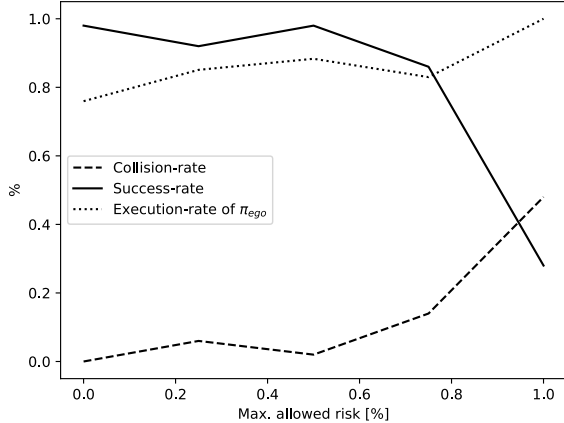
Fig. 5: The collison- and success-rate of the actual world plotted over the maximum collision-rate threshold ρ_{max}. If the collision-rate averaged over all counterfactual worlds is below the maximum collision-rate ρ_{max}, the learned policy is executed in the actual world. Otherwise, the ego vehicle is controlled by a lane-following model.

probability of $P_C = 0$.

Figure 4 shows the influence heatmap for the scenario shown in Figure 3 using the mean absolute displacement (MAD) comparing the vehicle's states of the counterfactual worlds to the actual world. The x-axis represents the counterfactual worlds $W_{\mathcal{P}}^{(\cdot)}$ in which policies have been exchanged and the y-axis shows the deviation using the MAD of the vehicle's states to the actual world W. For example, the second column in Figure 4 represents the counterfactual world $W_{\mathcal{P}}^{(1)}$ and the second row shows the MAD for the vehicle v_1 over the counterfactual worlds $W_{\mathcal{P}}^{(j)}$. As can be seen, replacing the policy of vehicle v_1 has a large influence on vehicle v_0. This is expected, as vehicle v_0 is controlled by the IDM model and v_1 is the leading vehicle. Vehicle v_2 is the leading vehicle of v_1 and, thus, has a large impact on vehicle v_1. Vehicle v_3 is controlled by the learned policy $\pi_\chi(a|s)$. As can be seen, all other vehicles have an influence on the learned policy of vehicle v_3. This is an important quantity to observe as this implies that the ego vehicle at least takes all its surrounding vehicles into account. Vehicle v_4 is controlled by the IDM model and, thus, is not being influenced by any other vehicles as there is not leading vehicle in the lane.

Figure 5 shows the collision-, execution-, and success-rate of the ego vehicle's policy π_{ego} in the actual world W plotted over the maximum allowed collision-rate ρ_{max} of the counterfactual worlds. If the average conditional collision-rate P_C is lower than the defined collision-threshold ρ_{max}, the learned policy $\pi_\chi(a|s)$ of the ego vehicle is executed. Otherwise, the ego vehicle is controlled by the Mobil model. We evaluate each of the maximum allowed collision-rates using 250 episodes. The success-rate is initially high as the policy is only executed if ρ_{max} is small and, thus, initially no collisions are caused when executing the learned policy.

With a higher maximum allowed collision rate ρ_{max}, the execution-rate of the learned policy goes up and the number of collisions increases in the actual world. This also leads to a significant drop in the success-rate as the goal is not reached as often. As can be seen, the best results are achieved using a maximum allowed collision-rate $\rho_{max} = 0$. At $\rho_{max} = 0$, the ego vehicle's policy π_{ego} is executed 76% of the time, the goal is reached with 99.8%, and the policy has a collision-rate of 0%.

Using the counterfactual policy evaluation prior to the execution of the policy in the actual world decreases the collision-rate significantly. Even though the policy is not as often executed, the goal is reached almost every single time as can be seen in Figure 5. Especially for black-box approaches, such as DNNs, the CPE algorithm provides insights on the performance, generalization, and safety prior to execution. Thus, increasing the applicability of these approaches in safety-critical applications. It can be seen whether a policy generalizes well judging on how well it can handle the counterfactual worlds.

V. CONCLUSION

This work introduces a counterfactual policy evaluation (CPE) algorithm that evaluates a policy prior to its execution in the real world. To perform such an evaluation, counterfactual worlds are introduced that evolve differently from the actual world due to exchanged behavior policies of the other vehicles. Using the evolutions of these worlds, we can reason on how well a policy performs and if it generalizes well enough to handle all counterfactual worlds. A maximum allowed collision-rate threshold is used that defines if a policy should be executed in the actual world. If the collision-rate is above the threshold, the policy is not executed in the actual world and the ego vehicle uses a lane-following model instead. Besides that, the CPE also makes causal relations between the vehicle visible and the influence of the vehicles of each other become evident. Especially, when using deep neural network it is important to see what and which vehicles the policy takes into account. Overall, the proposed method increases the applicability and understanding of policies in decision-making for autonomous vehicles significantly. In future work, the policies of multiple vehicles could be replaced to obtain even deeper insights on the causal relations between the vehicles. And the field of counterfactual explanations – also in the view of assurance cases – could be investigated further.

APPENDIX

Parameters used for learning the ego vehicle's behavior policy and other parameters used in the simulation:

```
"BehaviorIDM": {
  "MaxVelocity": 30.0,
  "MinimumSpacing": 2.0,
  "DesiredTimeHeadway": [1.0, 5.0],
  "MaxAcceleration": 2.5,
  "DesiredVelocity": 15.0,
  "ComfortableBrakingAcceleration": 1.6,
  "MinVelocity": 0.0,
  "Exponent": 4
},
"BehaviorSACAgent": {
  "ActorFcLayerParams": [512,256,256],
  "CriticJointFcLayerParams": [512,256,2
    56],
  "ActorLearningRate": 0.0003,
  "CriticLearningRate": 0.0003,
  "AlphaLearningRate": 0.0003,
  "TargetUpdateTau": 0.05,
  "TargetUpdatePeriod": 3,
  "Gamma": 0.995,
  "RewardScaleFactor": 1.0,
  "ReplayBufferCapacity": 10000,
  "BatchSize": 512,
},
"step_time": 0.2
```

REFERENCES

[1] S. Shalev-Shwartz, S. Shammah, and A. Shashua, *Safe, multi-agent, reinforcement learning for autonomous driving*, 2016.

[2] P. Hart and A. Knoll, *Graph neural networks and reinforcement learning for behavior generation in semantic environments*, 2020.

[3] J. Ho and S. Ermon, "Generative adversarial imitation learning," in *Advances in neural information processing systems*, 2016, pp. 4565–4573.

[4] Y. Li, J. Song, and S. Ermon, "Infogail: Interpretable imitation learning from visual demonstrations," in *Advances in Neural Information Processing Systems*, 2017, pp. 3812–3822.

[5] M. Bansal, A. Krizhevsky, and A. Ogale, "Chauffeurnet: Learning to drive by imitating the best and synthesizing the worst," *arXiv preprint arXiv:1812.03079*, 2018.

[6] M. Bojarski, D. Del Testa, D. Dworakowski, *et al.*, "End to end learning for self-driving cars," *arXiv preprint arXiv:1604.07316*, 2016.

[7] F. Codevilla, M. Miiller, A. López, V. Koltun, and A. Dosovitskiy, "End-to-end driving via conditional imitation learning," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2018, pp. 1–9.

[8] C. Zhang, O. Vinyals, R. Munos, and S. Bengio, *A study on overfitting in deep reinforcement learning*, 2018.

[9] X. Song, Y. Jiang, S. Tu, Y. Du, and B. Neyshabur, *Observational overfitting in reinforcement learning*, 2019.

[10] K. Kawaguchi, L. P. Kaelbling, and Y. Bengio, "Generalization in deep learning," *arXiv preprint arXiv:1710.05468*, 2017.

[11] K. Cobbe, O. Klimov, C. Hesse, T. Kim, and J. Schulman, *Quantifying generalization in reinforcement learning*, 2018.

[12] J. Pearl and D. Mackenzie, *The Book of Why: The New Science of Cause and Effect*. Basic Books, 2018.

[13] J. Pearl, *The do-calculus revisited*, 2012.

[14] C. Molnar, *Interpretable Machine Learning*. Lulu. com, 2020.

[15] L. Buesing, T. Weber, Y. Zwols, *et al.*, "Woulda, coulda, shoulda: Counterfactually-guided policy search," *7th International Conference on Learning Representations, ICLR 2019*, pp. 1–15, 2019.

[16] D. Isele, "Interactive Decision Making for Autonomous Vehicles in Dense Traffic," *2019 IEEE Intelligent Transportation Systems Conference, ITSC 2019*, pp. 3981–3986, 2019.

[17] N. Brown, A. Lerer, S. Gross, and T. Sandholm, "Deep counterfactual regret minimization," in *International conference on machine learning*, 2019, pp. 793–802.

[18] T. W. Neller and M. Lanctot, "An introduction to counterfactual regret minimization," in *Proceedings of Model AI Assignments, The Fourth Symposium on Educational Advances in Artificial Intelligence (EAAI-2013)*, vol. 11, 2013.

[19] M. Treiber, A. Hennecke, and D. Helbing, "Congested Traffic States in Empirical Observations and Microscopic Simulations," *Phys. Rev. E*, vol. 62, no. 2, pp. 1805–1824, 2000.

[20] M. Treiber and D. Helbing, "Mobil: General lane-changing model for car-following models," *Disponível Acesso Dezembro*, 2016.

[21] P. Hart, L. Rychly, and A. Knol, "Lane-Merging Using Policy-based Reinforcement Learning and Post-Optimization," 2020.

[22] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, "Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor," in *Proceedings of the 35th International Conference on Machine Learning, {ICML} 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, J. G. Dy and A. Krause, Eds., ser. Proceedings of Machine Learning Research, vol. 80, PMLR, 2018, pp. 1856–1865.