

# Driving Style Encoder: Situational Reward Adaptation for General-Purpose Planning in Automated Driving

Sascha Rosbach<sup>1,2</sup>, Vinit James<sup>1</sup>, Simon Großjohann<sup>1</sup>, Silviu Homoceanu<sup>1</sup>, Xing Li<sup>1</sup> and Stefan Roth<sup>2</sup>

**Abstract**—General-purpose planning algorithms for automated driving combine mission, behavior, and local motion planning. Such planning algorithms map features of the environment and driving kinematics into complex reward functions. To achieve this, planning experts often rely on linear reward functions. The specification and tuning of these reward functions is a tedious process and requires significant experience. Moreover, a manually designed linear reward function does not generalize across different driving situations. In this work, we propose a deep learning approach based on inverse reinforcement learning that generates situation-dependent reward functions. Our neural network provides a mapping between features and actions of sampled driving policies of a model-predictive control-based planner and predicts reward functions for upcoming planning cycles. In our evaluation, we compare the driving style of reward functions predicted by our deep network against clustered and linear reward functions. Our proposed deep learning approach outperforms clustered linear reward functions and is at par with linear reward functions with a-priori knowledge about the situation.

## I. INTRODUCTION

Automated driving in urban environments requires intelligent decision making that scales over a variety of traffic situations. A scalable approach needs to be able to address both structured and un-structured traffic as experienced in urban environments. In model-based planning, a semantic description of the environment is encoded in the form of static and kinematic features. The planning system has to translate this semantic description of the environment into safe and human-acceptable actions. The underlying planning algorithm often relies on manually-tuned linear reward functions that encode the relevance of the predefined features. Tuning such reward functions is a tedious task and is usually performed by motion planning experts. As a result, the reward function is often considered to be a static external signal that does not depend on the driving situation [1]. Manual tuning of the reward function becomes infeasible if a planning system is applied at scale and has to adopt a variety of driving styles. In this paper, we propose a deep learning approach in which a neural network dynamically predicts reward functions based on constantly changing static and kinematic features of the environment.

Human driving demonstrations enable the application of inverse reinforcement learning (IRL) for finding the under-

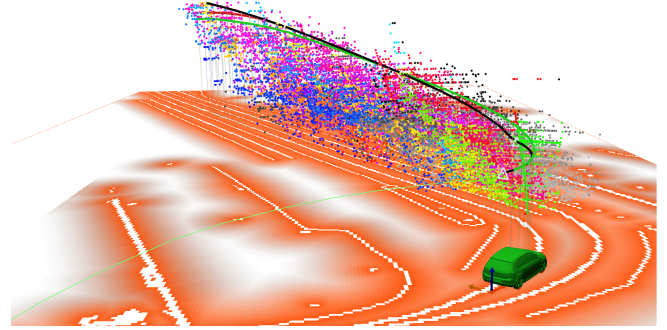


Fig. 1: This figure illustrates our planner for automated driving that samples policies for our path integral (PI) IRL formulation. The visualized state space is color-coded based on the state-action values. The z-axis corresponds to the velocity, whereas the ground plane depicts the spatial feature maps such as distance transformed lane centers and road boundaries. Three color-coded policies are visualized, namely the optimal policy (black), the odometry of a human demonstration (green), and the projection of the demonstration into the state space (red).

lying reward functions [2], [3]. In this work, we utilize this methodology to learn situation-dependent reward functions for our planner [4]–[6]. Our planner samples a large set of actions generating distributions of path integral (PI) features and actions. Unlike related work in deep inverse reinforcement learning (DIRL) that focuses on spatial reward functions [7], [8], our approach generates driving styles by incorporating vehicle kinematics. This is done by integrating our deep learning approach into a model-predictive control (MPC) planning algorithm. Sampled driving policies of the MPC are used as inputs to our neural network. The network learns a representation of the driving situation by matching distributions of features and actions to reward functions based on the maximum entropy principle.

The main contributions of this paper are threefold: First, we formulate a deep IRL methodology that predicts situation-dependent reward functions based on PI features and actions. Second, we propose a neural network architecture that utilizes one-dimensional convolutions over PI features and actions of sampled driving policies to learn a representation of the statics and kinematics of the situation. Third, we show the feasibility of our proposed approach in real-world traffic environments and compare our method to linear and latent maximum entropy IRL.

<sup>1</sup>The authors are with the Volkswagen AG, 38440 Wolfsburg, Germany {sascha.rosbach, vinit.james, simon.grossjohann, silviu.homoceanu, xing.li}@volkswagen.de

<sup>2</sup>The authors are with the Visual Inference Lab, Department of Computer Science, Technische Universität Darmstadt, 64289 Darmstadt, Germany stefan.roth@visinf.tu-darmstadt.de

## II. RELATED WORK

As of now, the decision-making system in automated vehicles is often decomposed into a hierarchical structure to reduce the complexity of the model-based system. This hierarchical decomposition, however, yields uncertainty due to insufficient knowledge of the subsequent level. In particular, these systems are prone to failure in unforeseen driving situations where simplified behavior models fail to match predefined templates [6]. General-purpose planning algorithms for automated driving aim to reduce planning-task decompositions to enable a scalable architecture that does not rely on behavior implementations. Starting with the work of McNaughton [9], attention has been drawn to parallel real-time planning that combines behavior planning and local motion planning. This approach enables sampling of a large set of actions that respect kinematic constraints. Thereby a sequence of sampled actions can represent complex maneuvers. This kind of general-purpose planner uses a linear reward function to map a complex feature space into rewards. Manual tuning of such a linear reward function is a tedious task performed by motion-planning experts. In our previous work, we automated the tuning process using IRL [5]. We found that inferred reward functions for situations defined a-priori surpass the performance of an expert-tuned reward function. However, the tuned linear reward functions do not generalize well over different situations as the objectives change continuously, e.g., the importance of keeping to the lane center in straight road segments while allowing deviations in curvy parts. In this work, we utilize the general-purpose planning approach and focus on situation-dependent reward function prediction.

A naive solution to the aforementioned problem is to use IRL with a-priori knowledge of situations and recover a set of situation-dependent reward functions. However, the complexity is limited to the number of identifiable scenarios that have to be mapped during inference. Prior work in IRL employs the latent maximum-entropy principle to condition reward functions on the situation [10]. Expectation Maximization (EM) is often used in combination with IRL to learn a prior over the situation [11]. Training can be performed on a mixture of different driving segments without a-priori knowledge about the situation. Hence, multiple reward functions are derived and a mixture is inferred during planning using Bayes' rule. Dirichlet process mixture models have been used to address some drawbacks of EM, specifically the reliance on a specified number of clusters [12], [13]. Recent work in IRL utilizes deep neural networks to extend the model capacity of mixture models. Wulfmeier et al. [7] proposed to learn spatial traversability by deep IRL using deep convolutional neural networks, which generates a direct mapping from raw sensory data to reward maps. This allows for encoding a situation into the reward function for a planning algorithm that operates on a grid-based representation. However, their IRL formulation relies on state-visitation frequency updates that do not scale to high-dimensional state-action spaces.

In contrast to our model-based IRL formulation, prior

work has also addressed high-dimensional state-spaces with model-free approaches. Finn et al. [14] proposed a sample-based deep IRL approach that learns the cost in the inner loop of a model-free policy optimization. This approach is suitable for problems with unknown system dynamics. Recent work also includes the idea of combining generative adversarial neural networks and imitation learning to clone the behavior of the demonstrated policy without finding the underlying reward functions [15].

In our work, we rely on a defined environment and vehicle transition model to explicitly incorporate traffic rules and provide safety [16]. We propose to couple the reliability of a model-based planning system with the generalization ability of deep inverse reinforcement learning. Our search-based planning algorithm operates in large and continuous state space and produces policies with multiple behaviors. This enables us to approximate the partition function required for the IRL gradient formulation with high accuracy [5]. Our IRL formulation uses a neural network to reason about static and kinematic features of the environment, thereby predicting situation-dependent driving styles.

## III. PRELIMINARIES

The interaction of the agent with the environment is often formulated as a Markov Decision Process (MDP), consisting of a 5-tuple  $\{\mathcal{S}, \mathcal{A}, T, R, \gamma\}$ , where  $\mathcal{S}$  denotes the set of states and  $\mathcal{A}$  describes the set of actions. The reward function  $R$  is discounted by  $\gamma$  and assigns a reward for every action  $a \in \mathcal{A}$  in state  $s \in \mathcal{S}$ . Our planning algorithm relies on an environment model  $M$  and an underlying vehicle transition function  $T$ . The model consists of static features of the environment that are derived from perception and localization, as well as kinematic features derived from the vehicle transition function. The planner generates a policy set  $\Pi$  by sampling actions  $a$  from distributions conditioned on vehicle dynamics for each state  $s$ . The features for each continuous action  $a$  are integrated in the policy generation process. The reward function  $R$  is given by a linear combination of  $K$  static and kinematic feature values  $f_i$  with weights  $\theta_i$  such that  $R(s, a) = \sum_{i \in K} -\theta_i f_i(s, a)$ . The value  $V^\pi$  of a policy  $\pi$  is the integral of discounted rewards during continuous transitions,  $V^\pi = \int_t \gamma_t R(s_t, a_t) dt$ . The PI feature  $f_i^\pi$  for a policy  $\pi$  is defined by  $f_i^\pi = \int_t \gamma_t f_i(s_t, a_t) dt$ . The demonstrations for our IRL formulation are based on odometry records  $\zeta$  of human driving. During the training in our simulation, we project the odometry record  $\zeta$  into the state-action space and are thereby able to formulate a demonstration from the policy set such that  $\pi^D \in \Pi$ .

### A. Maximum entropy IRL

The goal of IRL is to find the reward function weights  $\theta$  that enable the optimal policy  $\pi^*$  to be at least as good as the demonstrated policy  $\pi^D$  [17]. Thereby, the planner indirectly imitates the behavior of a demonstration [2]. In PI IRL, a probabilistic model is formulated that yields a probability distribution over policies,  $p(\pi|\theta)$  [18], [19]. The model is optimized such that the expected PI feature values

$\mathbb{E}_{p(\pi|\theta)}[\mathbf{f}^\pi]$  of the policy set  $\Pi$  match the empirical feature values  $\hat{\mathbf{f}}^{\Pi^D}$  of the demonstrations for each planning cycle of the MPC. Ziebart et al. [20] propose to maximize the entropy of the distribution to solve the ambiguity introduced by imperfect demonstrations, giving rise to the policy distribution

$$p(\pi|\theta) = \frac{1}{Z} \exp(-\theta^\top \mathbf{f}^\pi). \quad (1)$$

The calculation of the partition function  $Z = \sum_{\pi \in \Pi} \exp(-\theta^\top \mathbf{f}^\pi)$  is often intractable due to the exponential growth of the state-action space. Our planning algorithm approximates the partition function similar to Markov chain Monte Carlo methods. Maximizing the entropy of the distribution over policies subject to the feature constraints from demonstrated policies implies that the log-likelihood  $L(\theta)$  of the observed policies under the maximum entropy distribution is maximized. From this hypothesis, the log-likelihood gradient is obtained as

$$\nabla L(\theta) = \sum_{\pi \in \Pi} p(\pi|\theta) \mathbf{f}^\pi - \hat{\mathbf{f}}^{\Pi^D}. \quad (2)$$

Due to the dependency of the reward function on the driving situation, the probabilistic model  $p(\pi|\theta)$  that recovers a single reward function for the demonstrated trajectories does not scale.

#### IV. PATH INTEGRAL LATENT IRL

Instead of a single linear reward function, we consider that there are  $N$  different reward functions, each corresponding to situation-dependent behavior. As extension of our previous work, we incorporate latent variables in our PI IRL formulation [5], which allows us to infer a mixture of these reward functions during planning. In the following derivation, we consider a single demonstration  $\pi^D$  for every planning cycle. We adopt the formulation of Babes et al. [11] that utilizes Expectation Maximization (EM) in IRL, where  $\beta_c^{\pi^D}$  is the probability that a demonstration  $\pi^D$  belongs to a cluster  $c$ , and  $\psi(c)$  is the the estimated prior probability of a cluster  $c$ . The EM algorithm iteratively alternates between an estimation step (E-step) and a maximization step (M-step). Within the E-step, we compute the probability of a demonstration belonging to a cluster as

$$\beta_c^{t,\pi^D} = \frac{p(\pi^D|c, \theta^t, \psi^t) \psi^t(c)}{\sum_{c \in C} p(\pi^D|c, \theta^t, \psi^t) \psi^t(c)}, \quad (3)$$

exploiting the constraint  $\sum_{c \in C} \beta_c^{\pi^D} = 1$ . In Eq. 3,  $p(\pi^D|c, \theta, \psi)$  reduces to  $p(\pi^D|\theta_c)$ .

Within the M-step, we compute for each iteration  $t$  the prior probability  $\psi(c)$  as

$$\psi^{t+1}(c) = \frac{1}{B} \sum_{\pi^D \in \Pi^B} \beta_c^{t,\pi^D}, \quad (4)$$

where  $B$  is the number of demonstrations in a batch  $\Pi^B$ . Furthermore, we compute the reward functions  $\theta_c$  for each cluster as

$$\theta_c^{t+1} = \arg \max_{\theta} \sum_{\pi^D \in \Pi^B} \beta_c^{t,\pi^D} [\ln p(\pi^D|\theta_c)]. \quad (5)$$

We obtain the log-likelihood gradient for a cluster  $c$  as

$$\nabla_{\theta_c} L(\theta_c) = \beta_c^{\pi^D} \left( \sum_{\pi \in \Pi} p(\pi|\theta_c) \mathbf{f}^\pi - \mathbf{f}^{\pi^D} \right). \quad (6)$$

In contrast to the gradient formulation of IRL that finds a single reward function for the demonstrated policies, latent IRL finds a number of reward functions. The probability that a demonstration belongs to cluster  $\beta_c^{\pi}$  influences the update of the gradient. During training,  $\beta_c^{\pi}$  is inferred using Bayes' rule on the basis of the PI features of the demonstration. Due to the absence of features of the demonstration during online planning, we assume that the optimal solution on the basis of the last predicted reward function allows us to infer a situation-dependent reward function for the next planning cycle. In order to scale over a large number of situations, we next formulate an alternative approach based on deep learning.

#### V. PATH INTEGRAL DEEP IRL

We propose a deep learning approach for PI maximum entropy IRL that approximates a complex mapping between the situation and reward function. Our neural network predicts the reward function weights  $\theta_{k+1}$  based on the actions  $\mathbf{a}_k^\Pi$  and PI features  $\mathbf{f}_k^\Pi$  of the policy set  $\Pi_k$  at MPC cycle  $k$ , given by  $\theta_{k+1} \approx g(\Theta, \mathbf{f}_k, \mathbf{a}_k)$ .

The IRL problem can be formulated in the context of Bayesian inference as MAP estimation, which entails maximizing the joint posterior distribution of observing expert demonstrations  $\Pi^D$ . We calculate the maximum entropy probability based on the linear reward weights  $\theta$ , which are inferred by the network with parameters  $\Theta$  as

$$L(\theta) = L(g(\Theta, \mathbf{f}, \mathbf{a})) = \sum_{\pi^D \in \Pi^D} \ln p(\pi^D | g(\Theta, \mathbf{f}, \mathbf{a})). \quad (7)$$

Maximization leads to the optimal weights

$$\theta^* = \arg \max_{\theta} \sum_{\pi^D \in \Pi^D} \ln p(\pi^D | g(\Theta, \mathbf{f}, \mathbf{a})). \quad (8)$$

The gradient calculation from PI maximum entropy IRL from Eq. 2 lends itself naturally towards training deep neural networks, where the gradient of the log-likelihood  $L(\theta)$  can be calculated in terms of  $\Theta$  as

$$\begin{aligned} \frac{\partial L}{\partial \Theta} &= \frac{\partial L}{\partial \theta} \cdot \frac{\partial \theta}{\partial \Theta} \\ &= \left[ \sum_{\pi \in \Pi} p(\pi|\theta) \mathbf{f}^\pi - \hat{\mathbf{f}}^{\Pi^D} \right] \cdot \frac{\partial}{\partial \theta} g(\Theta, \mathbf{f}, \mathbf{a}). \end{aligned} \quad (9)$$

The gradient is separated into the maximum entropy gradient in terms of  $\theta$  and the gradient of  $\theta$  w.r.t. the network parameters  $\Theta$ , which can be directly obtained via backpropagation in the deep neural network.

##### A. Training methodology

Algo. 1 describes the training algorithm used for PI maximum entropy deep IRL. IRL training is often very time consuming since the MDP has to be solved with the current reward function in the inner loop of reward learning. In

---

**Algorithm 1:** Deep IRL Training

---

**Input:** odometry of human demonstration  $\zeta$ , model  $M$ , planning horizon  $H$ , reward discount factor  $\gamma$

**Output:** network weights  $\Theta^*$

```
1 function DeepIRLTraining( $\Pi^D, M, H, \gamma$ )
2    $\theta_0 \leftarrow$  init reward weights
3    $\Pi, \Pi^D \leftarrow$  Planning( $\theta_0, M, H, \gamma, \zeta$ )
4    $\Theta_0 \leftarrow$  init network weights
5   for  $e$  in Epochs do
6     for  $b$  in Batch do
7        $\theta_b \leftarrow$  network forward pass( $\Theta_b, f_b, a_b$ )
8        $\frac{\partial \mathcal{L}_b}{\partial \theta_b} = \sum_{\pi \in \Pi} p(\pi | \theta_b) f_b^\pi - \hat{f}_b^{\Pi^D}$ 
9        $\frac{\partial \mathcal{L}_b}{\partial \Theta_b} \leftarrow$  network backprop( $\Theta_b, f_b, a_b, \frac{\partial \mathcal{L}_b}{\partial \theta_b}$ )
10       $\Theta_{b+1} \leftarrow$  update weights( $\Theta_b, \frac{\partial \mathcal{L}_b}{\partial \Theta_b}$ )
```

---

our IRL formulation, we do not solve the MDP within the inner loop; instead we run our planning algorithm prior to training with a randomly initialized reward function  $\theta_0$  over training segments. Thereby, we generate a buffer of planning cycles consisting of a set of policy sets  $\Pi$  with corresponding PI features  $f^\Pi$  and actions  $a^\Pi$ . Due to the high-resolution sampling of actions, we ensure that there are policies that are geometrically close to human-recorded odometry and resemble human driving styles. Therefore the task of IRL is to find the unknown reward function that increases the likelihood of these trajectories to be considered as optimal policies. During the policy generation, we utilize a weighted Euclidean distance towards human driven trajectories  $\zeta$  to generate demonstrations  $\Pi^D$  for each policy set  $\Pi$ . Similar to the reward weights  $\theta$ , the network parameters  $\Theta$  are initialized with random values before starting the training. The training loop is run for a predefined number of epochs, ensuring that the convergence metric, the expected value difference (EVD), reaches the desired threshold value. For each epoch the training dataset is shuffled and divided into batches to perform mini-batch gradient decent.

### B. Neural network architecture

In this work a deep architecture is proposed, which maps the PI features  $f^\Pi$  and actions  $a^\Pi$  to linear reward weights  $\theta$ . The input to the network is a set of 15 PI features  $f^\Pi$  and 8 actions  $a^\Pi$  for every policy of a planning cycle. The continuous actions are reduced to steering angle and accelerations at discrete control points. The network architecture comprises of one-dimensional convolutions and average pooling to extract latent variables of policies describing the high-dimensional input space. The architecture uses one-dimensional convolutional layers to learn the causal relationship between sampled actions and resultant features of every policy. There is no inherent relationship among policies in the input space therefore no information is gained by performing higher dimensional convolutions. The architecture consists of convolutional building blocks, each comprising of two convolutional layers followed by an average pooling layer which performs dimension reduction. There are five such blocks in succession followed by fully-connected layers. A series of eight fully-connected layers at the end learn inter-policy relationships from the resultant latent variables of the

convolutional stack. A vector of linear reward weights with the same dimension as the no. of features, i.e. 15, is the output of the network. All activation functions in the network are chosen as Rectified Linear Units, except for the output layer where no activation function is used.

### C. Inference during MPC planning cycles

During real-time planning in an automated drive, the MPC re-plans in discrete time-steps  $k$ . After receiving the features and actions of the latest planning cycle, the neural network infers the reward weights. To enable smooth transitions of the reward functions, we utilize a predefined history size  $h$  to calculate the empirical mean of weights  $\hat{\theta}$ . The weights hence obtained are used to continuously re-parameterize the planning algorithm for the subsequent planning cycle.

## VI. EXPERIMENTS

In our experiments, we use our proposed PI deep IRL (PIDIRL), as well as PI latent IRL (PILIRL) and PI linear IRL (PIIRL) in real urban driving situations. We focus our experiments on situations that introduce conflicts in driving objectives so as to highlight the importance of situation-dependent reward functions.

### A. Data collection and simulation

Our experiments are conducted on a prototype vehicle, which uses a mediated perception architecture to produce feature maps of the environment. These feature maps are depicted in the ground-plane in Fig. 1. The MPC planning algorithm uses the feature maps to get state features and computes state-values as shown in as color-coded point clouds in Fig. 1. We gathered data on a set of streets in Hamburg, Germany located around Messe and Congress, see Fig. 3a. We concentrate on static infrastructure and objects while disabling intention prediction of dynamic objects to test our reward functions in different location-dependent situations on this track. Three drives over this course are segmented into train and test tracks, each of which is subdivided into four situations. Situation 1 resembles sharp curves, situation 2 stopping, starting and turns at traffic lights, situation 3 stopping and starting at a traffic lights, and situation 4 resembles lane following. The training of our algorithm is performed in a playback simulation of the recorded data as depicted in Fig. 1. After every planning cycle of the MPC, the position of the vehicle is reset to the odometry recording of the human demonstration

### B. Reward feature representation

We concentrate on a reward set consisting of  $K = 15$  manually engineered features, which are listed in the table of Fig. 3g. Infrastructural features are derived by a data fusion between objects and street network [21]. The kinematic characteristics of the policies are given by derivatives of the lateral and longitudinal actions as well as features related to behavior, e.g. lane change delay. In addition selection features are designed so as to evaluate policies based on more nuanced attributes such as maneuvers space, progress and policy end direction toward the road.

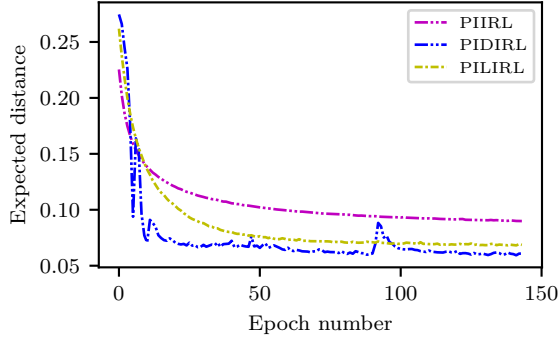


Fig. 2: Expected distance over epochs of training.

## VII. EVALUATION

We compare the results of the situation-dependent reward functions from PI deep IRL (PIDIRL) against PI latent IRL (PILIRL), and PI linear IRL (PIIRL) with and without a-priori knowledge about the situation. First, we compare the model training based on distance convergence toward the human demonstration. Second, we compare the driving style of our policies under learned reward functions against human driving in different driving situations. Third, we analyze the reward weights inferred on the different test segments.

### A. Training evaluation

We analyze the convergence of our model training based on a distance metric  $d(\zeta, \pi)$  towards the human driven demonstration  $\zeta$ . This distance metric is an integral of the Euclidean distance in longitudinal and lateral direction as well as the squared difference in the yaw angle over the policy as described in detail in [5]. Due to our goal of optimizing the likelihood of human driving behavior within the policy set  $\Pi$ , we measure our training convergence based on the expected distance (ED), given by  $\mathbb{E}[d(\zeta, \Pi)] = \sum_{\pi \in \Pi} p(\pi|\theta) d(\zeta, \pi)$ . We plot ED for each IRL approach over an equal amount of training epochs. In Fig. 2 the ED for PIDIRL, PILIRL and PIIRL are plotted for 143 epochs. The PIDIRL shows a large ED reduction 78% over the entire training dataset. The starting ED is high due to the randomly initialized network, but a high reduction of 74% is achieved in only 22 epochs. PILIRL also has a high initial ED due to higher variance in weight initialization in order to find multiple different clusters, it finally shows a reduction of 74%. Due to the limited model complexity, PIIRL shows lowest ED reduction of 60%, here trained over the entire training dataset without a-priori knowledge about the situations. The slight variance in the PIDIRL plot, particular in epoch 90, can be explained due to high learning rate and small batch size. From the statics above we can deduce that PIDIRL has the highest likelihood of selecting policies with humanlike driving behavior across different driving situations.

### B. Driving style evaluation

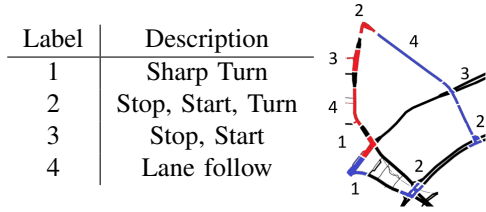
In this part of the evaluation, we compare the driving style of PIDIRL, PILIRL, and PIIRL models against manual

human driving over different test situations. We generate multiple PIIRL models with a-priori knowledge for each driving situations so as to compare the situation-dependent performance of the proposed PIDIRL. We plot the distance of the optimal policy obtained from these models to the odometry of the human demonstration for different driving situations. Our test segments contain 187 planning cycles for segment 1, 232 for segment 2, 217 for segment 3, and 224 for segment 4. In Fig. 3c-f, we fit a Gaussian distribution over the segment specific distance histogram with 200 bins of size 0.0025. From our experiments we deduce that our proposed PIDIRL method produces policies with 21% lower mean distance as compared to PILIRL in seg. 1 and in seg. 2. In seg. 3 the mean distances from both of these methods are comparable. Moreover, PIDIRL shows 13% lower mean distance as compared to the segment specific PIIRL for seg. 3 and comparable results for the other segments. In seg. 4 PIDIRL produces policies with higher mean distances as compared to PIIRL and PILIRL. This trend can be described by the dominant prior of PILIRL and overfitting of PIIRL for straight segments. This poses as a disadvantage for both of these methods as compared to our proposed PIDIRL which has better generalization capabilities. This generalization can be seen in the Fig. 3b, where PIDIRL has a similar mean distance over all segments. The reward function weights obtained from these models over the test situations are shown in Fig. 3h-k, where the weights are inversely proportional to the preference of that feature value in the policy. The reward weights are plotted in logarithmic scale. In these plots, it can be seen that PIDIRL produces nuanced reward functions for every planning cycle in the segment. The variance of the reward function produced by PIDIRL is proportional to the situation complexity. In simple situations, as shown in Fig. 3b at seg. 4, having a nuanced reward functions prediction mechanism like PIDIRL producing high variance as shown in Fig. 3k could lead to lower performance as compared to simpler reward functions namely PILIRL and PIIRL.

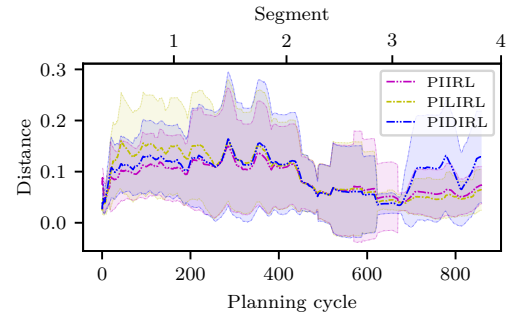
## VIII. CONCLUSION AND FUTURE WORK

We utilize path integral (PI) maximum entropy deep IRL to learn situation-dependent reward functions of a general-purpose planning algorithm. We propose a method to couple the reliability of a model-based planning system with the generalization capability of deep inverse reinforcement learning. Our experiments show that reward function predictions by our proposed neural network architecture are at par with multiple PI linear IRL reward function model that are trained with a-priori knowledge about the situation. In addition to this, we show that our deep IRL methodology has better generalization capabilities as compared to PI latent IRL which uses behavioral clustering of demonstrations. In future we plan to experiment with different network architectures in our proposed deep IRL paradigm so as to reduce the variance in reward functions prediction in simple driving situations.

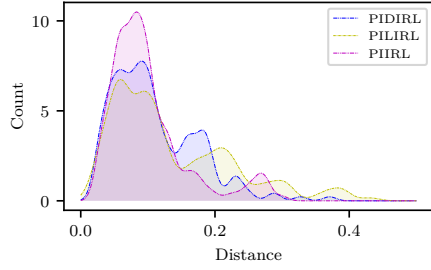




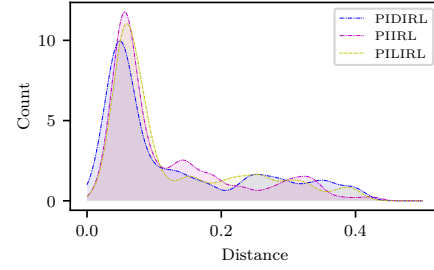
(a) The map depicts the training (blue) and test (red) segments of the selected Route in Hamburg City Center.



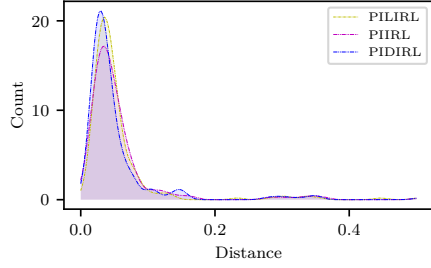
(b) Distance of the opt. policies on different driving segments.



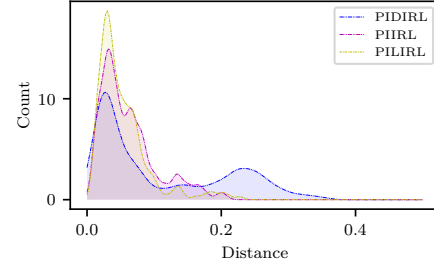
(c) Distribution of the opt. policy distances in seg. 1.



(d) Distribution of the opt. policy distances in seg. 2.



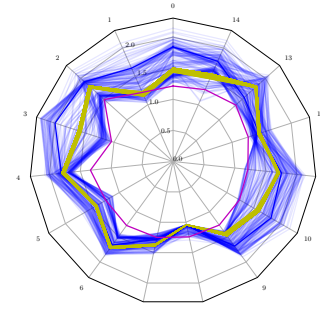
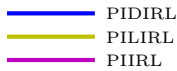
(e) Distribution of the opt. policy distances in seg. 3.



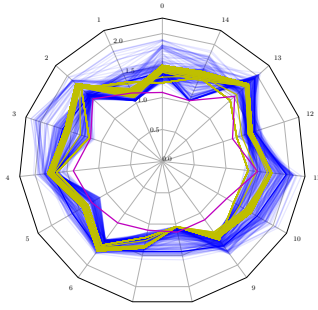
(f) Distribution of the opt. policy distances in seg. 4.

	Label	Feature
Kinematics	0	Long. Acc.
	1	Long. Jerk
	2	Long. Velocity
	3	Lat. Acc.
	4	Lat. Jerk
	9	Lat. Overshooting
	10	Lane Change Delay
Statics	5	Centerline
	6	Direction
	7	Proximity
	8	Curbs
Selection	11	State Class
	12	Manuver Space
	13	End Direction
	14	Min. Progress

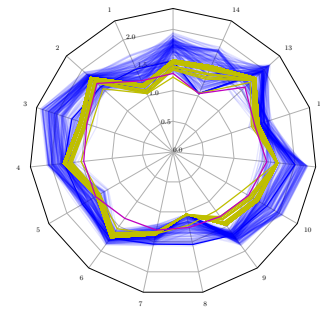
(g) The table lists the considered PI features. The parameters are plotted in logarithmic scale.



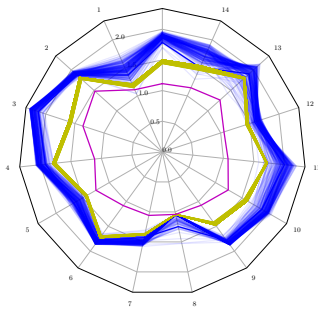
(h) Inf. reward weights on seg. 1.



(i) Inf. reward weights on seg. 2.



(j) Inf. reward weights on seg. 3.



(k) Inf. reward weights on seg. 4.

Fig. 3: (a) Tests conducted on a selected Route in Hamburg City center. (b-f) Graphs compares the performance of PIDIRL with PIILRL, and PIIRL over different driving segments. (g) Table of reward function features. (h-k) Graphs show the reward function predictions of our methods over driving segments.

## REFERENCES

- [1] S. Levine, "Reinforcement learning and control as probabilistic inference: Tutorial and review," *arXiv preprint arXiv:1805.00909*, 2018.
- [2] A. Y. Ng and S. J. Russell, "Algorithms for Inverse Reinforcement Learning," in *Proc. Int. Conf. Machine Learning (ICML)*, 2000, pp. 663–670.
- [3] M. Kuderer, S. Gulati, and W. Burgard, "Learning driving styles for autonomous vehicles from demonstration," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2015, pp. 2641–2646.
- [4] S. Heinrich, A. Zoufahl, and R. Rojas, "Real-time trajectory optimization under motion uncertainty using a GPU," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots and Syst. (IROS)*, 2015, pp. 3572–3577.
- [5] S. Rosbach, V. James, S. Großjohann, S. Homocceanu, and S. Roth, "Driving with style: Inverse reinforcement learning in general-purpose planning for automated driving," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots and Syst. (IROS)*, 2019, pp. 2658–2665.
- [6] S. Heinrich, "Planning Universal On-Road Driving Strategies for Automated Vehicles," Ph.D. dissertation, Freie Universität Berlin, 2018.
- [7] M. Wulfmeier, P. Ondruska, and I. Posner, "Maximum Entropy Deep Inverse Reinforcement Learning," *arXiv:1507.04888 [cs]*, p. 1507.04888, 2015.
- [8] M. Wulfmeier, D. Rao, D. Z. Wang, P. Ondruska, and I. Posner, "Large-scale cost function learning for path planning using deep inverse reinforcement learning," in *Int. J. Robotics Research*, vol. 36, no. 10, pp. 1073–1087, 2017.
- [9] M. McNaughton, "Parallel Algorithms for Real-time Motion Planning," Ph.D. dissertation, Carnegie Mellon University, 2011.
- [10] S. Wang, D. Schuurmans, and Y. Zhao, "The Latent Maximum Entropy Principle," in *Proc. IEEE Int. Symp. on Inf. Theory ISIT*, 2002, pp. 131–.
- [11] M. Babes, V. Marivate, K. Subramanian, and M. L. Littman, "Apprenticeship learning about multiple intentions," in *Proc. Int. Conf. Machine Learning (ICML)*, 2011, pp. 897–904.
- [12] M. Shimosaka, K. Nishi, J. Sato, and H. Kataoka, "Predicting driving behavior using inverse reinforcement learning with multiple reward functions towards environmental diversity," in *IEEE Intell. Vehicles Symp.*, 2015, pp. 567–572.
- [13] J. Choi and K.-E. Kim, "Nonparametric Bayesian inverse reinforcement learning for multiple reward functions," in *Adv. in Neural Inform. Process. Syst.*, 2012, pp. 305–313.
- [14] C. Finn, S. Levine, and P. Abbeel, "Guided Cost Learning: Deep Inverse Optimal Control via Policy Optimization," in *J. of Machine Learning Research*, pp. 49–58, 2016.
- [15] J. Ho and S. Ermon, "Generative adversarial imitation learning," in *Adv. in Neural Inform. Process. Syst.*, 2016, pp. 4565–4573.
- [16] C. M. Hruschka, D. Töpfer, and S. Zug, "Risk assessment for integral safety in automated driving," in *Int. Conf. on Intell. Autonomous Systems (ICoIAS)*, 2019, pp. 102–109.
- [17] S. Arora and P. Doshi, "A survey of inverse reinforcement learning: Challenges, methods and progress," in *arXiv Preprint arXiv:1806.06877*, 2018.
- [18] N. Aghasadeghi and T. Bretl, "Maximum entropy inverse reinforcement learning in continuous state spaces with path integrals," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots and Syst. (IROS)*, 2011, pp. 1561–1566.
- [19] E. Theodorou, J. Buchli, and S. Schaal, "A generalized path integral control approach to reinforcement learning," in *J. of Machine Learning Research*, vol. 11, no., pp. 3137–3181, 2010.
- [20] B. D. Ziebart, A. L. Maas, J. A. Bagnell, and A. K. Dey, "Maximum Entropy Inverse Reinforcement Learning," in *Proc. Nat. Conf. Artificial Intell. (AAAI)*, vol. 8, 2008, pp. 1433–1438.
- [21] K. Homeier and L. Wolf, "RoadGraph: High level sensor data fusion between objects and street network," in *Proc. IEEE Int. Conf. Intell. Transp. Syst. (ITSC)*, 2011, pp. 1380–1385.