

Learning Risk-Aware Costmaps via Inverse Reinforcement Learning for Off-Road Navigation

Samuel Triest¹, Mateo Guaman Castro¹, Parv Maheshwari²,
 Matthew Sivaprakasam¹, Wenshan Wang¹, and Sebastian Scherer¹

Abstract—The process of designing costmaps for off-road driving tasks is often a challenging and engineering-intensive task. Recent work in costmap design for off-road driving focuses on training deep neural networks to predict costmaps from sensory observations using corpora of expert driving data. However, such approaches are generally subject to overconfident mis-predictions and are rarely evaluated in-the-loop on physical hardware. We present an inverse reinforcement learning-based method of efficiently training deep cost functions that are uncertainty-aware. We do so by leveraging recent advances in highly parallel model-predictive control and robotic risk estimation. In addition to demonstrating improvement at reproducing expert trajectories, we also evaluate the efficacy of these methods in challenging off-road navigation scenarios. We observe that our method significantly outperforms a geometric baseline, resulting in 44% improvement in expert path reconstruction and 57% fewer interventions in practice. We also observe that varying the risk tolerance of the vehicle results in qualitatively different navigation behaviors, especially with respect to higher-risk scenarios such as slopes and tall grass.³

I. INTRODUCTION

Navigation in unstructured environments is a major challenge critical to many robotics applications such as off-road navigation, search and rescue, last-mile delivery and robotic exploration. As such, the problem has been widely studied in robotics [1–7]. A key component of navigation in unstructured environments is the capability to take in high-dimensional sensory data and convert it into intermediate representations for downstream planning and control. Typically, this representation takes the form of a costmap, where cells in the costmap describe some notion of the cell’s traversability, based on the observed features of that cell. However, determining the mapping between observations and cost is dependent on many factors and is often extremely challenging to do in practice.

Recent work in producing costmaps focuses on directly learning cell costs from map features and expert demonstrations using inverse reinforcement learning (IRL) [8–13] and deep neural networks. This is accomplished by training function approximators to predict a mapping from features to cost that best explains the behavior of an expert from large

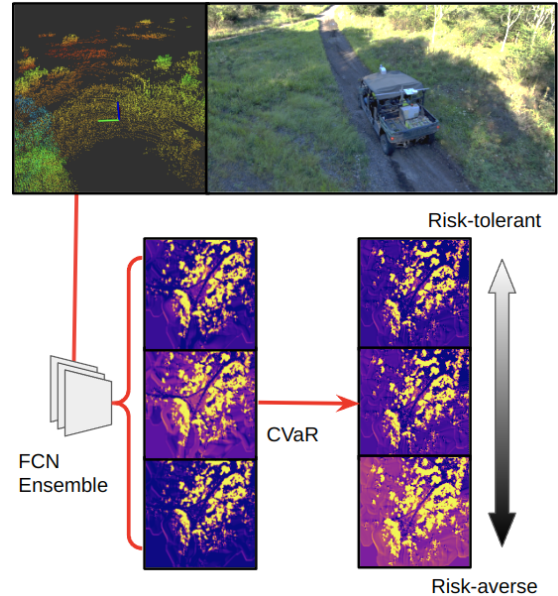


Fig. 1: We present an IRL-based costmap learning method that produces risk-aware costmaps from lidar data.

corpora of demonstrations. This deep, non-linear representation of the cost function in theory allows for more accurate costmap learning from fewer map features. However, neural networks are often susceptible to domain shift. While this problem is less significant with simple linear combinations of features [14,15], increasing the complexity of the function approximator to tens, or hundreds of thousands of parameters can lead to confident mis-predictions for rarely-visited terrain features. In unstructured environments, these mis-predictions can lead to highly undesirable outcomes, such as getting stuck, or damaging the robot. Additionally, standard practice for inverse RL relies on performing value iteration. This approach is both slow, and usually ignores non-holonomic vehicle constraints and dynamics.

In this paper, we describe a costmap learning procedure that extends the standard IRL framework in two key ways:

- 1) Similar to [13], we replace the slow value iteration step of traditional IRL methods with MPPI [16], allowing for faster training that respects vehicle dynamics.
- 2) We train an ensemble of cost functions to quantify uncertainty in cost using conditional value-at-risk.

In addition to showing improved results over an occupancy-based baseline, we also validate our costmap learning procedure via navigation experiments on a large all-terrain vehicle and demonstrate that our learned costmaps

* This work was supported by ARL awards #W911NF1820218 and #W911NF20S0005.

¹ Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, USA. {striest,mguamanc,msivapra,wenshanw,basti}@andrew.cmu.edu

² Department of Mathematics, Indian Institute of Technology Kharagpur. parvmaheshwari2002@iitkgp.ac.in

³More detailed algorithms and additional visualizations are provided in the appendix [Appendix](https://arxiv.org/abs/2305.12345) (appendix link: [tinyurl.com/mtkj63e8](https://arxiv.org/abs/2305.12345))

result in safer navigation than traditional, occupancy-based baselines [4]. Furthermore, we demonstrate that changing the risk-tolerance of the vehicle results in significant differences in navigation behavior over higher-uncertainty terrain such as slopes and tall grass.

II. RELATED WORK

Costmaps are a common representation in off-road navigation tasks due to their ease of use and ability to easily associate cost with spatial locations. As such, a large body of work exists on producing costmaps from sensory data.

Perhaps the most widely-used costmap generation method relies on computing cost from geometric properties of the terrain. The most straightforward form of this approach is to create an occupancy-based representation of the terrain and assign high cost to cells that are occupied with high-height terrain. Krüsi et al. [17], Fankhauser et al. [18] and Fan et al. [19] extend this representation with additional costs derived from geometric terrain features such as curvature and roughness. They all demonstrate that these features allow for improved navigation over rough terrain.

Recent work in off-road driving has focused on leveraging semantic segmentation approaches to assign traversability costs to different terrain types. Traditional work performs this segmentation in a first-person view (FPV) [10,20–22]. However, it remains challenging to convert this segmentation into a useful representation for navigation, as simple projection is not robust to occlusions or sensor mis-calibration. As such, recent work by Shaban et al. [23] instead performs semantic segmentation in a bird’s-eye view (BEV), using pointclouds from lidar. By first manually defining a mapping between semantic classes and cost, they are able to train a network to directly output costmaps. While they are able to demonstrate impressive navigation and prediction results, the mapping between semantic classes and cost is both heuristically defined and coarse, limiting the network’s ability to generalize between robots or easily adapt online.

Inverse RL for costmap generation for wheeled vehicles is also a well-explored problem [24–26]. Recent work in generating costmaps for off-road vehicles largely focus on introducing deep neural networks as cost function approximators. Wulfmeier et al. [9] trained a deep, fully convolutional network using inverse RL for urban driving scenarios and demonstrate improved performance in generating expert trajectories as compared to a hand-crafted baseline. They also observe an issue in that there are few gradient propagations in rarely-explored areas of the state space, causing their cost networks to perform more poorly on unseen state features. Lee et al. [13] focus on learning costmaps in a multi-agent highway driving scenario. In addition to replacing the value iteration step of IRL with MPPI [27], they also propose regularizing unvisited states to have zero cost to alleviate the costmap artifacts observed by Wulfmeier et al. They additionally add a time dimension to the state-space to account for dynamic agents in their environment. Zhu et al. [12] propose a modification to the standard IRL procedure that takes into account an approximation of the kinematics

of Ackermann-steered vehicles. They demonstrate that their algorithm results in improved ability to re-create expert behavior on unimproved roads. Zhang et al. [11] extend the basic IRL framework to include kinematic information. They concatenate information such as velocity and curvature into their intermediate feature representations in order to get costmaps that are able to reflect the current dynamics of their vehicle. Wigness et al. [10] combine semantic segmentation and IRL for a skid-steered robot in off-road scenarios. Unlike previous approaches which focus on extracting costmaps from lidar features, they instead produce costmaps from semantic classes from BEV-projected camera data. Using this feature space, they are then able to learn a costmap using a linear combination of these classes, as well as an obstacle layer. They also demonstrate the capability to quickly update navigation behaviors with a small number of additional demonstrations.

III. METHODOLOGY

A. Preliminaries

1) *MaxEnt IRL*: Maximum Entropy Inverse Reinforcement Learning (MaxEnt IRL) is a popular framework for extracting cost functions for a Markov Decision Process (MDP) from large corpora of human demonstrations [15]. MaxEnt IRL builds off of the results shown by Abbeel and Ng [14] that one can obtain a policy close to the expert policy by matching feature expectations. However, the baseline IRL problem is ill-posed, as there are an infinite number of cost functions that can explain a set of expert trajectories. Ziebart et al. [15] propose using the principle of maximum entropy [28] to address the ill-posed nature of unregularized IRL. They show that the gradient for maximizing the likelihood of the expert trajectories τ_E under the learned reward function with entropy regularization can be computed as Equation 1.

$$\nabla_{\theta} \sum_{\tau_E \in \mathcal{D}_E} L(\tau_E|\theta) = \tilde{f} - \sum_{s_i} D_{s_i}^L f_{s_i} \quad (1)$$

$$\nabla_f \sum_{\tau_E \in \mathcal{D}_E} L(\tau_E|\theta) = \sum_{s_i} [D_{s_i}^E - D_{s_i}^L] \quad (2)$$

Here, f_{s_i} are the features associated with state s_i , $D_{s_i}^L$ are the learner’s state visitations (i.e. the amount of time spent in state s_i), and \tilde{f} is the expert’s expected feature counts ($\tilde{f} = \sum_{s_i} D_{s_i}^E f_{s_i}$). Note that the computation of this gradient requires enumeration over the state-space of the MDP.

Performing gradient ascent using this objective has been shown to produce strong results for reward functions linear in state features (of the form $r = \theta^T f$). Wulfmeier et al. [8] extend this formulation by providing a gradient of the IRL objective with respect to state features (Equation 2). This gradient allows for the training of deep neural networks via backpropagation for IRL (MEDIRL).

2) *MPPI*: Traditionally, in order to compute state visitation frequencies $D_{s_i}^L$, value iteration is performed at each gradient step [8,11,15]. For navigation policies in particular, states are generally connected via some pre-defined edge

structure, such as a 4 or 8-connected grid [11], or an approximation of robot kinematics [12]. However, such approaches are relatively slow, and do not obey the true dynamics of the vehicle. We follow an approach similar to [13] in that we first compute solutions to the MDP in continuous space via a fast MPC algorithm (MPPI), then discretize the resulting solutions ($\tau = x_{1:T}$) into the resolution of the IRL state space. Given that MPPI can already be thought of as performing importance sampling of the optimal distribution from a proposal distribution, we can also use the existing update weights η from MPPI to weight the state visitation frequencies we compute from the MPPI proposal distribution. In order to compute the state distribution of the expert, we can simply set η to one. Note that we rely on a position extraction function $p(X) = [x, y]$ to get x-y positions from state. We find that this solution is faster to compute than prior implementations of IRL and respects the kinematic and dynamic constraints of our Ackermann-steered vehicle.

3) *Conditional Value-at-Risk*: While in theory the implicit regularization of MaxEnt IRL should result in a single weight vector for a given expert dataset, we observe that in practice, especially with neural networks, MaxEnt IRL can converge to meaningfully different solutions. This phenomenon is also noted by Wulfmeier et al. [9]. To combat potential uncertainty in feature space, we train an ensemble of function approximators using MEDIRL and combine their predictions using Conditional Value-at-Risk as a risk metric.

While originally used in econometrics applications, there has been interest from the robotics community in using CVaR as a risk metric for reasoning about distributions of cost [29–32] for its ability to capture more accurately the nuances of long-tailed and multimodal distributions. Intuitively, CVaR_ν can be thought of as the mean value of a distribution, when only considering the portion of the distribution exceeding a given quantile (a.k.a. Value at Risk (VaR)) $\nu \in [0, 1]$ (described formally in Equation 3) [33].

$$\text{CVaR}_\nu = \int_{f(x) > \text{VaR}_\nu}^{\infty} f(x)p(x)dx \quad (3)$$

Note that CVaR can also be used to capture the behavior of the lower tail of the distribution by taking values *below* a given quantile [32]. Via a small abuse of notation, we will refer to this mode of CVaR as CVaR_ν for $\nu \in [-1, 0]$ (Equation 4). This gives us a range of $\nu \in [-1, 1]$ that we can use to smoothly vary the risk-tolerance of the vehicle. Note that CVaR_0 corresponds to taking the mean of the ensemble (and being neutral to risk), similar to Ratliff et al. [25].

$$\text{CVaR}_{-\nu} = \int_{-\infty}^{f(x) < \text{VaR}_{1-\nu}} f(x)p(x)dx \quad (4)$$

B. Fast MEDIRL with Uncertainty Estimates

We present an efficient method for extracting costmaps from expert demonstrations with uncertainty by combining the above methods. An overview of our algorithm is presented in Figure 2. At a high level, our algorithm is similar to [8,13]. However, we train an ensemble of FCNs instead of a

single network. This allows us to use ensemble disagreement to estimate uncertainty, similar to prior work [34–37]. In contrast to prior work, we reason over costs. This allows us to condense our cost estimates with CVaR, as opposed to using some form of variance. At train time we construct an MDP by doing the following for each training example (τ^E, M_t) :

- 1) Extract a fixed time-window of expert trajectory $\tau_{t:t+H}^E$ (7.5s or 75 timesteps for our experiments), starting at the time the map features M_t were recorded.
- 2) Set the first expert state τ_0^E as the MDP start state.
- 3) Set the final expert state τ_T^E as the MDP goal state.
- 4) Obtain a costmap by randomly selecting a costmap predictor from the ensemble and running it on the lidar features at timestep 0.
- 5) Set state transition function to be vehicle dynamics.
- 6) Set cost function as a weighted combination of the costmap C_t and final-state distance-to-goal (Eq. 5).

$$J(\tau) = \sum_{s_i \in \tau} [C[p(x_i)]] + \kappa \|x_T - x_g\|_2 \quad (5)$$

We then solve this MDP via MPPI to obtain the learner’s state distribution under the predicted costmap. Using this state distribution, we are then able to update our network weights via the gradient provided in Equation 1.

An issue noted by Ratliff et al. [25] is that driving demonstrations may be significantly sub-optimal, especially with respect to homotopy classes around obstacles (e.g. an expert may drive left around a tree, when it is just as good or better to go right). We avoid this issue somewhat by goal-conditioning our MDP based on the final state of a short (7.5s) lookahead window. This also simplifies data collection in that human teleoperators are not required to set and achieve particular goal points during data collection.

At test time, we generate a costmap from each FCN and compute a cell-wise CVaR to condense the set of costmaps into a single risk-aware costmap. We can then solve the MDP via MPPI. Unlike train time, however, we can also warm-start the next optimization with the previous solution. In practice, this allows us to run fewer MPPI iterations at test time.

IV. HARDWARE IMPLEMENTATION

We tested our algorithms on a Yamaha Viking All-Terrain Vehicle (ATV) which was modified for autonomous driving by Mai et al. [38]. The navigation stack was run with a 12th gen Intel i7 cpu and NVIDIA 3080 laptop GPU. A NovAtel PROPAK-V3-RT2i GNSS unit provided global pose estimates as well as IMU data. Additionally, a Velodyne UltraPuck was used to collect pointclouds.

A. State Estimation

We use Super Odometry [39] to provide pose estimates and registered pointclouds, given unregistered lidar scans and IMU data from the NovAtel.

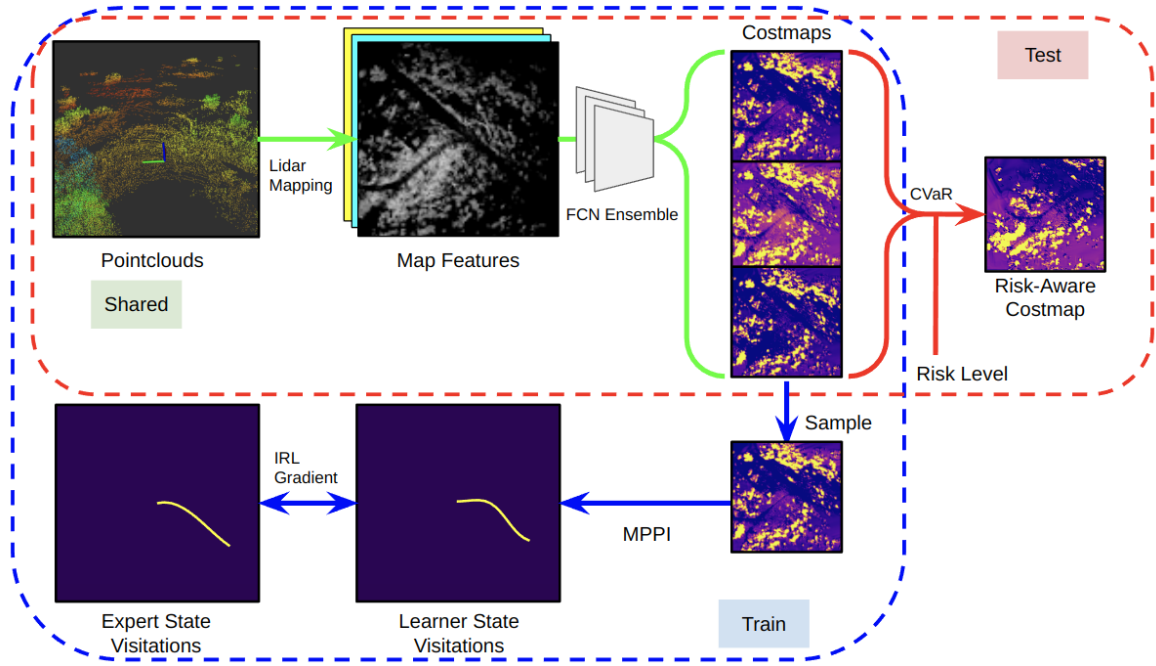


Fig. 2: An overview of our algorithm. We first process lidar scans into a tensor of map features. We then use an ensemble of FCNs to create a set of costmaps. In order to train these networks, we can sample from our set of costmaps, solve the resulting MDP with MPPI, and use the resulting state distribution to supervise our networks with inverse RL. At test time, the costmaps can be aggregated together using CVaR to create a costmap associated with a given risk level.

Feature	Calculation
Height Low	$\min_p [p_z], \forall p p_z < t_z + k_{overhang}$
Height Mean	$\frac{1}{ P } \sum_p [p_z], \forall p p_z < t_z + k_{overhang}$
Height High	$\max_p [p_z], \forall p p_z < t_z + k_{overhang}$
Height Max	$\max_p [p_z], \forall p$
Terrain (T)	See appendix
Slope	$0.5 (\frac{\partial}{\partial x} T + \frac{\partial}{\partial y} T)$
Diff	Height high - Terrain
SVD1	$\frac{\lambda_1 - \lambda_2}{\lambda_1}$
SVD2	$\frac{\lambda_2 - \lambda_3}{\lambda_2}$
SVD3	$\frac{\lambda_1}{\lambda_3}$
Roughness	$\frac{\lambda_1}{\lambda_3}$
Unknown	$\mathbf{1} [P = 0]$

TABLE I: List of grid map features, and their calculations given the points in a cell

B. Mapping and Feature Extraction

Given the registered pointclouds from Super Odometry, we then process the points into a grid map representation [40] to perform learning on. This grid map is represented as a dense three-dimensional tensor with two spatial dimensions corresponding to planar position, and the third dimension representing various terrain features of the cell. We represent terrain as an $80m \times 80m$ map at $0.5m$ resolution with 12 features per cell, resulting in a $160 \times 160 \times 12$ tensor. The map is centered at the robot position. We present a simplified description of our lidar mapping algorithm in the appendix, and a description of our lidar features in Table I.

C. Costmapping

Given the resulting $[W \times H \times D]$ tensor of map features, we run our ensemble of costmap networks to produce a $[B \times$

$W \times H]$ costmap (where B is the size of the ensemble). Our costmap network is a small, fully convolutional resnet [41], which was trained from scratch. We then perform a CVaR calculation for each cell to get a single risk-aware costmap of size $[W \times H]$, given a CVaR value ν .

D. Control

In order to perform navigation on top of these costmaps, we use a modified version of MPPI [27]. We observed that sampling noise from an Ornstein-Uhlenbeck process [42] results in better optimization than the original independent Gaussian noise for our particular application. Our implementation of MPPI minimizes a weighted sum of the costmap and final state distance to a specified goal point (Equation 5). We optimize trajectories through a kinematic bicycle model (KBM) with additional dynamics on throttle and steering. States and derivatives were clamped according to the actuation and safety limits of the vehicle. Similar model and optimization parameters were used for both the inner MPPI optimization of the MEDIRL algorithm and the MPC running on the vehicle. To meet the on-board compute and safety constraints of our vehicle, less computationally-heavy optimization parameters, and lower velocity limits were chosen for the on-board MPC.

E. Dataset

We ran our algorithm on a corpus of aggressive off-road driving data. Compared to our closest point of comparison, TartanDrive [43], this dataset contains significantly more off-trail scenarios such as pushing through tall grass and bushes, or driving down steep slopes with gravel. Most critically, our

Function Approximator	MHD
Occupancy	3.220 ± 0.027
Linear	1.847 ± 0.033
Linear Sigmoid	1.955 ± 0.046
Resnet	1.973 ± 0.022
Resnet Sigmoid	1.794 ± 0.033

TABLE II: Performance of various function approximators on reproducing expert behavior

dataset contains data from a lidar, allowing for features at a much farther range (40m vs. 10m forward in TartanDrive). A more detailed analysis of the driving data is presented in the appendix. In total, we collected roughly one hour of trajectories to train on, and about 15 minutes of separate trajectories to test on.

V. EXPERIMENTS, RESULTS AND ANALYSIS

A. Evaluation of IRL Methods

Similar to existing literature [10–12], we use modified Hausdorff distance [44] as our primary metric to measure the ability of our algorithm to generate costmaps that yield expert behavior (MHD). We compare our method to an occupancy-based baseline, which assigns high cost to all cells whose “diff” feature (Table I) exceeds the wheel radius. These obstacle cells were then C-space expanded to match the geometry of the ATV. For all experiments except the occupancy-based baseline, an ensemble of 16 function approximators was used, and learner trajectories were generated using the mean of the ensemble predictions (CVaR₀). Mean and standard deviation of results are computed over three seeds for each experiment. Results are presented in Table II.

Overall, the IRL-based methods outperformed the occupancy-based baseline by a significant margin. While all IRL-based methods performed similarly, we observe a small, but statistically significant reduction in MHD by using a Resnet. We also observe that ensembles of neural networks exhibit more variation in cost when changing CVaR values.

Additionally, we compare the iteration speeds and efficiency of prior IRL methods for offroad driving to our framework. We compare specifically to Zhang et al. [11] and Zhu et al. [12] (Table III). We are interested in evaluating the time to solve the inner RL loop, as well as the feasibility of these solutions under the motion constraints of our vehicle. Both our method and [11] were tested on an Intel Xeon CPU. Timing results for [11] are used from their paper as no source code was available. We compared to their GPU results. Using sampling from actual trajectory rollouts, we are able to achieve faster iteration speed per training sample. Additionally, we are also able to solve the inner RL problem more accurately, as we use the dynamics directly, as opposed to approximating the kinematics.

B. Large-Scale Navigation Experiments

Large-scale navigation experiments were conducted at a testing site near Pittsburgh, PA. The experiments consisted of navigation through a challenging 1.6km course that was defined by a series of GPS waypoints. Given the sensing horizon (40m in all directions) and focus on local planning,

Method	Time per Sample	Kinematics	Dynamics
[11]	2.0s	No	No
[12]	0.9s	Yes	No
Ours	0.3s	Yes	Yes

TABLE III: Comparison of properties of different IRL methods. In addition to having the fastest iteration time per sample, our method also respects both the kinematics and dynamics of the vehicle.

Method	Auto. Dist.	Auto. Speed	Interventions
Baseline	1465m	3.21m/s	7
IRL (CVaR -0.9)	1482m	3.25m/s	4
IRL (CVaR 0.0)	1498m	3.29m/s	3
IRL (CVaR 0.9)	1592m	3.14m/s	4

TABLE IV: Navigation metrics for each method. Distance and speed columns are not bolded as higher/lower values don’t necessarily mean better performance.

waypoints were spaced 50m apart. In all experiments, all components and parameters of our navigation stack were identical, except for the costmapping method. As we are interested in evaluating the effect of the costmap on overall navigation performance, we report the number of interventions from a human safety driver as our primary metric. We also report the total distance traveled and average speed, excluding distance traveled while intervening. The safety driver was instructed to intervene in the following two cases:

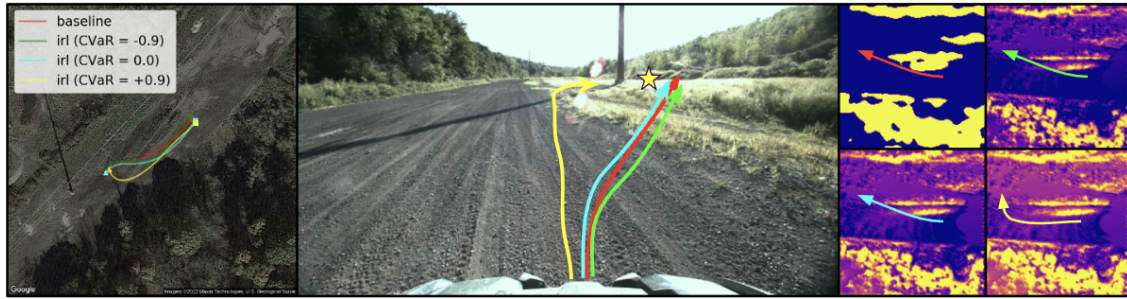
- 1) If the ATV was going to drive into an obstacle
- 2) If the ATV drove past a waypoint without going within a 4m radius of it

An intervention consisted of the safety driver teleoperating the vehicle past the point that resulted in the intervention such that a troublesome terrain feature would not result in multiple interventions. Results are provided in Table IV. Additional figures are provided in the appendix. Overall, it can be observed all IRL-based methods outperformed the occupancy-based baseline. All IRL-based methods resulted in a similar number of interventions, but exhibited some qualitatively different navigation behaviors. Figure 3 highlights several key differences in observed behaviors.

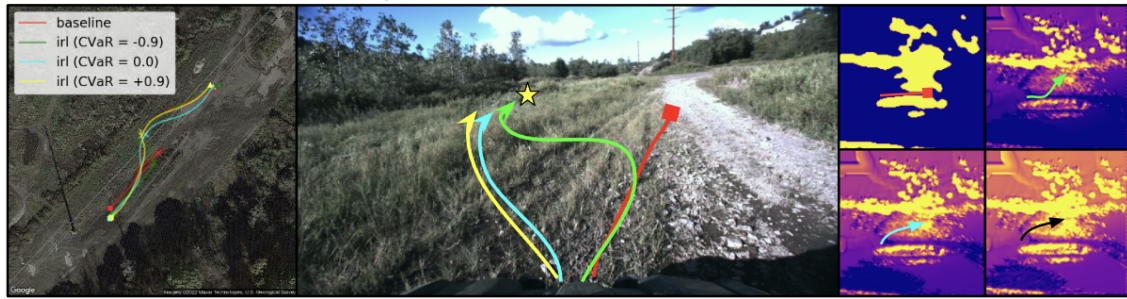
Overall, we are again able to observe significant improvement over the baseline when using IRL-based costmaps. All IRL-based methods had roughly the same number of interventions, but exhibited different navigation behaviors, especially with respect to tall grass. In general, more conservative settings of CVaR (> 0.8), caused the ATV to avoid patches of moderate-height grass when possible, while more aggressive settings of CVaR (< -0.8), caused the ATV to drive through taller patches of grass. As one would expect, average distance traveled increased with CVaR_v, as there were monotonically more high-cost regions to avoid. The occupancy-based method traveled the shortest amount of autonomous distance as a result of having more human interventions. Interestingly, we observed that the difference between the speeds and distance traveled for CVaR_{-0.9} and CVaR₀ were relatively small, while CVaR_{0.9} was considerably slower and traveled farther (this is potentially because



(a) Costmaps learned via IRL allow the robot to cut a corner over some low grass (about 15cm), while the baseline takes a longer route (note that obstacle inflation makes the grass patch appear as an obstacle to the baseline).



(b) Both the baseline and some IRL costmaps result in navigation over roughly 25cm grass, while IRL with a low risk tolerance (CVaR=0.9), takes a longer route around.



(c) The baseline method is unable to navigate to the waypoint in the tall grass, resulting in an intervention. IRL with a high risk tolerance (CVaR=-0.9) stays on the trail longer, then cuts through a taller grass patch. (black arrow used for clarity in costmap plot for CVaR=0.9)

Fig. 3: Several scenarios that during the navigation run that resulted in different behaviors. Left column: An enlarged BEV of the particular scenario, with trajectories from the individual runs superimposed (start=square, end=triangle, intervention=X, waypoint=yellow star). Middle column: FPV of the given scenario, with the paths from each trial annotated on. Right column: Visualization of the respective costmaps (blue=low cost, yellow=high cost) and trajectories for each. (top left: baseline, top right: IRL (CVaR -0.9), bottom left: IRL (CVaR 0.0), bottom right: IRL (CVaR 0.9).

of one fewer intervention for $CVaR_0$).

VI. CONCLUSION AND FUTURE WORK

We present an efficient method of training risk-aware costmaps with MEDIRL and demonstrate improved performance over occupancy-based baselines in challenging navigation scenarios. We also show that changing the risk-tolerance of the vehicle via CVaR results in different behaviors with respect to uncertain terrains such as tall grass.

We believe that there are three important potential extensions of this work. First, improving the set of features that MEDIRL can use should in theory lead to overall improvement. To that end, we are interested in adding visual features in addition to our lidar features. Methods such as that presented by Harley et al. [45] are immediately applicable

to our current feature representation, and would allow us to learn off of features such as RGB, semantic segmentation or even deep features. Second, we believe that significant navigation improvements can come from considering the velocity dimension. It follows intuition that different terrains can incur different costs depending on speed. As such, we believe that improvement in navigation can be achieved by predicting costmaps with an additional velocity dimension. Finally, while we demonstrate that different CVaR values result in different navigation behaviors, deciding the correct CVaR value remains an open problem. It seems that there are many factors that influence the level of risk an ATV should take, including speed, surroundings and operator preference. As such, we are also interested in designing controllers that are able to adapt their risk-tolerance online.

REFERENCES

- [1] S. R. Team, “Stanford racing team’s entry in the 2005 darpa grand challenge,” *Published on DARPA Grand Challenge*, 2005.
- [2] L. D. Jackel, E. Krotkov, M. Perschbacher *et al.*, “The darpa lagr program: Goals, challenges, methodology, and phase i results,” *Journal of Field robotics*, vol. 23, no. 11-12, pp. 945–973, 2006.
- [3] S. Scherer, V. Agrawal, G. Best *et al.*, “Resilient and modular subterranean exploration with a team of roving and flying robots,” *Field Robotics Journal*, pp. 678–734, May 2022.
- [4] “Arl autonomy stack.” 2022, accessed 13-September-2022. [Online]. Available: <https://www.arl.army.mil/business/collaborative-alliances/current-cras/sara-cra/sara-overview/>
- [5] P. Borges, T. Peynot, S. Liang *et al.*, “A survey on terrain traversability analysis for autonomous ground vehicles: Methods, sensors, and challenges,” *Field Robotics*, vol. 2, no. 1, pp. 1567–1627, 2022.
- [6] S. H. Young, “Robot autonomy in complex environments,” in *Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications III*, vol. 11746. SPIE, 2021, p. 1174602.
- [7] A. Agha, K. Otsu, B. Morrell *et al.*, “Nebula: Quest for robotic autonomy in challenging environments; team costar at the darpa subterranean challenge,” *arXiv preprint arXiv:2103.11470*, 2021.
- [8] M. Wulfmeier, P. Ondruska, and I. Posner, “Deep inverse reinforcement learning,” *CoRR*, abs/1507.04888, 2015.
- [9] M. Wulfmeier, D. Rao, D. Z. Wang *et al.*, “Large-scale cost function learning for path planning using deep inverse reinforcement learning,” *The International Journal of Robotics Research*, vol. 36, no. 10, pp. 1073–1087, 2017.
- [10] M. Wigness, J. G. Rogers, and L. E. Navarro-Serment, “Robot navigation from human demonstration: Learning control behaviors,” in *2018 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2018, pp. 1150–1157.
- [11] Y. Zhang, W. Wang, R. Bonatti *et al.*, “Integrating kinematics and environment context into deep inverse reinforcement learning for predicting off-road vehicle trajectories,” *arXiv preprint arXiv:1810.07225*, 2018.
- [12] Z. Zhu, N. Li, R. Sun *et al.*, “Off-road autonomous vehicles traversability analysis and trajectory planning based on deep inverse reinforcement learning,” in *2020 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2020, pp. 971–977.
- [13] K. Lee, D. Isele, E. A. Theodorou, and S. Bae, “Spatiotemporal costmap inference for mpc via deep inverse reinforcement learning,” *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 3194–3201, 2022.
- [14] P. Abbeel and A. Y. Ng, “Apprenticeship learning via inverse reinforcement learning,” in *Proceedings of the twenty-first international conference on Machine learning*, 2004, p. 1.
- [15] B. D. Ziebart, A. L. Maas, J. A. Bagnell *et al.*, “Maximum entropy inverse reinforcement learning,” in *Aaai*, vol. 8. Chicago, IL, USA, 2008, pp. 1433–1438.
- [16] G. Williams, A. Aldrich, and E. A. Theodorou, “Model predictive path integral control: From theory to parallel computation,” *Journal of Guidance, Control, and Dynamics*, vol. 40, no. 2, pp. 344–357, 2017.
- [17] P. Krüsi, P. Furgale, M. Bosse, and R. Siegwart, “Driving on point clouds: Motion planning, trajectory optimization, and terrain assessment in generic nonplanar environments,” *Journal of Field Robotics*, vol. 34, no. 5, pp. 940–984, 2017.
- [18] P. Fankhauser, M. Bjelonic, C. D. Bellicoso *et al.*, “Robust rough-terrain locomotion with a quadrupedal robot,” in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 5761–5768.
- [19] D. D. Fan, K. Otsu, Y. Kubo *et al.*, “Step: Stochastic traversability evaluation and planning for risk-aware off-road navigation,” *arXiv preprint arXiv:2103.02828*, 2021.
- [20] D. Maturana, P.-W. Chou, M. Uenoyama, and S. Scherer, “Real-time semantic mapping for autonomous off-road navigation,” in *Field and Service Robotics*. Springer, 2018, pp. 335–350.
- [21] P. Jiang, P. Osteen, M. Wigness, and S. Saripalli, “Rellis-3d dataset: Data, benchmarks and analysis,” 2020.
- [22] M. Wigness, S. Eum, J. G. Rogers *et al.*, “A rugd dataset for autonomous navigation and visual perception in unstructured outdoor environments,” in *International Conference on Intelligent Robots and Systems*, 2019.
- [23] A. Shaban, X. Meng, J. Lee *et al.*, “Semantic terrain classification for off-road autonomous driving,” in *Conference on Robot Learning*. PMLR, 2022, pp. 619–629.
- [24] N. D. Ratliff, J. A. Bagnell, and M. A. Zinkevich, “Maximum margin planning,” in *Proceedings of the 23rd international conference on Machine learning*, 2006, pp. 729–736.
- [25] N. D. Ratliff, D. Silver, and J. A. Bagnell, “Learning to search: Functional gradient techniques for imitation learning,” *Autonomous Robots*, vol. 27, no. 1, pp. 25–53, 2009.
- [26] J. A. Bagnell, D. Bradley, D. Silver *et al.*, “Learning for autonomous navigation,” *IEEE Robotics & Automation Magazine*, vol. 17, no. 2, pp. 74–84, 2010.
- [27] G. Williams, N. Wagener, B. Goldfain *et al.*, “Information theoretic mpc for model-based reinforcement learning,” in *2017 IEEE International Conference on Robotics and Automation*. IEEE, 2017, pp. 1714–1721.
- [28] E. T. Jaynes, “On the rationale of maximum-entropy methods,” *Proceedings of the IEEE*, vol. 70, no. 9, pp. 939–952, 1982.
- [29] A. Majumdar and M. Pavone, “How should a robot assess risk? towards an axiomatic theory of risk in robotics,” in *Robotics Research*. Springer, 2020, pp. 75–84.
- [30] A. Choudhry, B. Moon, J. Patrikar *et al.*, “Cvar-based flight energy risk assessment for multirotor uavs using a deep energy model,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 262–268.
- [31] D. D. Fan, A.-A. Agha-Mohammadi, and E. A. Theodorou, “Learning risk-aware costmaps for traversability in challenging environments,” *IEEE Robotics and Automation Letters*, vol. 7, no. 1, pp. 279–286, 2021.
- [32] X. Cai, M. Everett, J. Fink, and J. P. How, “Risk-aware off-road navigation via a learned speed distribution map,” *arXiv preprint arXiv:2203.13429*, 2022.
- [33] R. T. Rockafellar, S. Uryasev *et al.*, “Optimization of conditional value-at-risk,” *Journal of risk*, vol. 2, pp. 21–42, 2000.
- [34] A. G. Kendall, “Geometry and uncertainty in deep learning for computer vision,” Ph.D. dissertation, University of Cambridge, UK, 2019.
- [35] D. Pathak, D. Gandhi, and A. Gupta, “Self-supervised exploration via disagreement,” in *International conference on machine learning*. PMLR, 2019, pp. 5062–5071.
- [36] T. Yu, G. Thomas, L. Yu *et al.*, “Mopo: Model-based offline policy optimization,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 14 129–14 142, 2020.
- [37] S. J. Wang, S. Triest, W. Wang *et al.*, “Rough terrain navigation using divergence constrained model-based reinforcement learning,” in *5th Annual Conference on Robot Learning*, 2021.
- [38] J. Mai, “System design, modelling, and control for an off-road autonomous ground vehicle,” Master’s thesis, Carnegie Mellon University, Pittsburgh, PA, July 2020.
- [39] S. Zhao, H. Zhang, P. Wang *et al.*, “Super odometry: Imu-centric lidar-visual-inertial estimator for challenging environments,” in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 8729–8736.
- [40] P. Fankhauser and M. Hutter, “A universal grid map library: Implementation and use case for rough terrain navigation,” in *Robot Operating System (ROS)*. Springer, 2016, pp. 99–120.
- [41] K. He, X. Zhang, S. Ren, and J. Sun, “Identity mappings in deep residual networks,” in *European conference on computer vision*. Springer, 2016, pp. 630–645.
- [42] G. E. Uhlenbeck and L. S. Ornstein, “On the theory of the brownian motion,” *Physical review*, vol. 36, no. 5, p. 823, 1930.
- [43] S. Triest, M. Sivaprakasam, S. J. Wang *et al.*, “Tartandrive: A large-scale dataset for learning off-road dynamics models,” *arXiv preprint arXiv:2205.01791*, 2022.
- [44] M.-P. Dubuisson and A. K. Jain, “A modified hausdorff distance for object matching,” in *Proceedings of 12th international conference on pattern recognition*, vol. 1. IEEE, 1994, pp. 566–568.
- [45] A. W. Harley, Z. Fang, J. Li *et al.*, “A simple baseline for bev perception without lidar,” *arXiv preprint arXiv:2206.07959*, 2022.