



Inverse reinforcement learning for autonomous navigation via differentiable semantic mapping and planning

Tianyu Wang¹ · Vikas Dhiman² · Nikolay Atanasov¹

Received: 1 November 2022 / Accepted: 9 June 2023 / Published online: 6 July 2023
© The Author(s) 2023

Abstract

This paper focuses on inverse reinforcement learning for autonomous navigation using distance and semantic category observations. The objective is to infer a cost function that explains demonstrated behavior while relying only on the expert's observations and state-control trajectory. We develop a map encoder, that infers semantic category probabilities from the observation sequence, and a cost encoder, defined as a deep neural network over the semantic features. Since the expert cost is not directly observable, the model parameters can only be optimized by differentiating the error between demonstrated controls and a control policy computed from the cost estimate. We propose a new model of expert behavior that enables error minimization using a closed-form subgradient computed only over a subset of promising states via a motion planning algorithm. Our approach allows generalizing the learned behavior to new environments with new spatial configurations of the semantic categories. We analyze the different components of our model in a minigrid environment. We also demonstrate that our approach learns to follow traffic rules in the autonomous driving CARLA simulator by relying on semantic observations of buildings, sidewalks, and road lanes.

Keywords Inverse reinforcement learning · Semantic mapping · Autonomous navigation

1 Introduction

Autonomous systems operating in unstructured, partially observed, and changing real-world environments need an understanding of semantic meaning to evaluate the safety, utility, and efficiency of their performance. For example, while a bipedal robot may navigate along sidewalks, an autonomous car needs to follow the road lanes and traffic signs. Designing a cost function that encodes such rules by hand is impractical for complex tasks. It is, however, often possible to obtain demonstrations of desirable behavior that

indirectly capture the role of semantic context in the task execution. Semantic labels provide rich information about the relationship between object entities and their utility for task execution. In this work, we consider an inverse reinforcement learning (IRL) problem in which observations containing semantic information about the environment are available.

Consider imitating a driver navigating in an unknown environment as a motivating scenario (see Fig. 1). The car is equipped with sensors that can reveal information about the semantic categories of surrounding objects and areas. An expert driver can reason about a course of action based on this contextual information. For example, staying on the road relates to making progress, while hitting the sidewalk or a tree should be avoided. One key challenge in IRL is to infer a cost function when such expert reasoning is not explicit. If reasoning about semantic entities can be learned from the expert demonstrations, the cost model may generalize to new environment configurations. To this end, we propose an IRL algorithm that learns a cost function from semantic features of the environment. Simultaneously recognizing the environment semantics and encoding costs over them is a very challenging task. While other works learn a black-box neural network parametrization to map observations

✉ Tianyu Wang
tiw161@eng.ucsd.edu

Vikas Dhiman
vikas.dhiman@maine.edu

Nikolay Atanasov
natanasov@eng.ucsd.edu

¹ Electrical and Computer Engineering, University of California, San Diego, 9500 Gilman Dr, La Jolla, CA 92093, USA

² Electrical and Computer Engineering, The University of Maine, 168 College Ave, Orono, ME 04469, USA

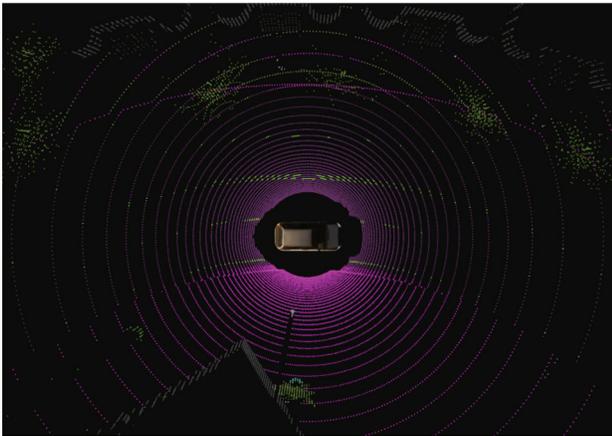


Fig. 1 Autonomous vehicle in an urban street, simulated via the CARLA simulator (Dosovitskiy et al., 2017). The vehicle is equipped with a LiDAR scanner, four RGB cameras, and a segmentation algorithm, providing a semantically labeled point cloud. An expert driver demonstrates lane keeping (green) and avoidance of sidewalks (pink) and buildings (gray). This paper considers inferring the expert's cost function and generating behavior that can imitate the expert's response to semantic observations in new operational conditions (Color figure online)

directly to costs (Wulfmeier et al., 2016; Song, 2019), we take advantage of semantic segmentation and occupancy mapping before inferring the cost function. A metric-semantic map is constructed from causal partial semantic observations of the environment to provide features for cost function learning. Contrary to most IRL algorithms, which are based on the maximum entropy expert model (Ziebart et al., 2008; Wulfmeier et al., 2016), we propose a new expert model allowing bounded rational deviations from optimal behavior (Baker et al., 2007). Instead of dynamic programming over the entire state space, our formulation allows efficient deterministic search over a subset of promising states. A key advantage of our approach is that this deterministic planning process can be differentiated in closed-form with respect to the parameters of the learnable cost function.

This work makes the following contributions:

1. We propose a cost function representation composed of a *map encoder*, capturing semantic class probabilities from online, first-person, distance and semantic observations and a *cost encoder*, defined as a deep neural network over the semantic features.
2. We propose a new expert model which enables cost parameter optimization with a closed-form subgradient of the cost-to-go, computed only over a subset of promising states.
3. We evaluate our model in autonomous navigation experiments in a 2D minigrid environment (Chevalier-Boisvert et al., 2018) with multiple semantic categories (e.g. wall, lawn, lava) as well as an autonomous driving task that

respects traffic rules in the CARLA simulator (Dosovitskiy et al., 2017).

2 Related work

2.1 Imitation learning

Imitation learning (IL) has a long history in reinforcement learning and robotics (Ross et al., 2011; Atkeson & Schaal, 1997; Argall et al., 2009; Pastor et al., 2009; Zhu et al., 2018; Rajeswaran et al., 2018; Pan et al., 2020). The goal is to learn a mapping from observations to a control policy to mimic expert demonstrations. Behavioral cloning (Ross et al., 2011) is a supervised learning approach that directly maximizes the likelihood of the expert demonstrated behavior. However, it typically suffers from distribution mismatch between training and testing and does not consider long-horizon planning. Another view of IL is through inverse reinforcement learning where the learner recovers a cost function under which the expert is optimal (Neu & Szepesvári, 2007; Ng & Russell, 2000; Abbeel & Ng, 2004). Recently, Ghasemipour et al. (2020) and Ke et al. (2020) independently developed a unifying probabilistic perspective for common IL algorithms using various f-divergence metrics between the learned and expert policies as minimization objectives. For example, behavioral cloning minimizes the Kullback–Leibler (KL) divergence between the learner and expert policy distribution while adversarial training methods, such as AIRL (Fu et al., 2018) and GAIL (Ho & Ermon, 2016) minimize the KL divergence and Jensen Shannon divergence, respectively, between state-control distributions under the learned and expert policies.

2.2 Inverse reinforcement learning

Learning a cost function from demonstration requires a control policy that is differentiable with respect to the cost parameters. Computing policy derivatives has been addressed by several successful IRL approaches (Neu & Szepesvári, 2007; Ratliff et al., 2006; Ziebart et al., 2008). Early works assume that the cost is linear in the feature vector and aim at matching the feature expectations of the learned and expert policies. Ratliff et al. (2006) compute subgradients of planning algorithms to guarantee that the expected reward of an expert policy is better than any other policy by a margin. Value iteration networks (VIN) by Tamar et al. (2016) show that the value iteration algorithm can be approximated by a series of convolution and maxpooling layers, allowing automatic differentiation to learn the cost function end-to-end. Ziebart et al. (2008) develop a dynamic programming algorithm to maximize the likelihood of observed expert data and learn a policy with maximum entropy (MaxEnt). Many works (Levine et al., 2011; Wulfmeier et al., 2016; Song, 2019)

extend MaxEnt to learn a nonlinear cost function using Gaussian Processes or deep neural networks. Finn et al. (2016) use a sampling-based approximation of the MaxEnt partition function to learn the cost function under unknown dynamics for high-dimensional continuous systems. However, the cost in most existing work is learned offline using full observation sequences from the expert demonstrations. A major contribution of our work is to develop cost representations and planning algorithms that rely only on causal partial observations. In the case where demonstrations are suboptimal with respect to the true cost function, a learned cost function can be recovered with preference-based comparisons (Brown et al., 2020; Jeon et al., 2020), self-supervision (Chen et al., 2021) or human corrections and improvements (Bajcsy et al., 2017; Jain et al., 2015). In this work, we assume that only the demonstrations are provided and we cannot assess the demonstrator’s suboptimality with respect to the unknown true cost.

2.3 Mapping and planning

There has been significant progress in semantic segmentation techniques, including deep neural networks for RGB image segmentation (Papandreou et al., 2015; Badrinarayanan et al., 2017; Chen et al., 2018) and point cloud labeling via spherical depth projection (Wu et al., 2018; Dohan et al., 2015; Milioto et al., 2019; Cortinhal et al., 2020). Maps that store semantic information can be generated from segmented images (Sengupta et al., 2012; Lu et al., 2019). Gan et al. (2020) and Sun et al. (2018) generalize binary occupancy mapping (Hornung et al., 2013) to multi-class semantic mapping in 3D. In this work, we parameterize the navigation cost of an autonomous vehicle as a nonlinear function of such semantic map features to explain expert demonstrations.

Achieving safe and robust navigation is directly coupled with the quality of the environment representation and the cost function specifying desirable behaviors. Traditional approaches combine geometric mapping of occupancy probability (Hornung et al., 2013) or distance to the nearest obstacle (Oleynikova et al., 2017) with hand-specified planning cost functions. Recent advances in deep reinforcement learning demonstrated that control inputs may be predicted directly from sensory observations (Levine et al., 2016). However, special model designs (Khan et al., 2018) that serve as a latent map are needed in navigation tasks where simple reactive policies are not feasible. Gupta et al. (2017) decompose visual navigation into two separate stages explicitly: mapping the environment from first-person RGB images in local coordinates and planning through the constructed map with VIN (Tamar et al., 2016). Our model constructs a global map instead and, yet, remains scalable with the size of the environment due to our sparse tensor implementation.

This paper is a revised and extended version of our previous conference publications (Wang et al., 2020a, b). In our previous work (Wang et al., 2020a), we proposed differentiable mapping and planning stages to learn the expert cost function. The cost function is parameterized as a neural network over binary occupancy probabilities, updated from local distance observations. An A* motion planning algorithm computes the policy at the current state and backpropagates the gradient in closed-form to optimize the cost parameterization. We proposed an extension of the occupancy map from binary to multi-class in Wang et al. (2020b), which allows the cost function to capture semantic features from the environment. This paper unifies our results in a common differentiable multi-class mapping and planning architecture and presents an in-depth analysis of the various model components via experiments in a minigrid environment and the CARLA simulator. This work also introduces a new sparse tensor implementation of the multi-class occupancy mapping stage to enable its use in large environments.

3 Problem formulation

3.1 Environment and agent models

Consider an agent aiming to reach a goal in an a priori unknown environment with different terrain types. Figure 2 shows a grid-world illustration of this setting. Let $\mathbf{x}_t \in \mathcal{X}$ denote the agent state (e.g., pose, twist, etc.) at discrete time t . In this work, we will consider $\mathbf{x}_t \in SE(2)$ composed of 2D position and orientation. Let $\mathbf{x}_g \in \mathcal{X}$ be the goal state. The agent state evolves according to known deterministic dynamics, $\mathbf{x}_{t+1} = f(\mathbf{x}_t, \mathbf{u}_t)$, with control input $\mathbf{u}_t \in \mathcal{U}$. The control space \mathcal{U} is assumed finite. Let $\mathcal{K} = \{0, 1, 2, \dots, K\}$ be a set of class labels, where 0 denotes “free” space and $k \in \mathcal{K} \setminus \{0\}$ denotes a particular semantic class such as road, sidewalk, or car. Let $m^* : \mathcal{X} \rightarrow \mathcal{K}$ be a function specifying the true semantic occupancy of the environment by labeling states with semantic classes. We implicitly assume that m^* assigns labels to agent positions rather than to other state variables. We do not introduce an output function, mapping an agent state to its position, to simplify the notation. Let \mathcal{M} be the space of possible environment realizations m^* . Let $c^* : \mathcal{X} \times \mathcal{U} \times \mathcal{M} \rightarrow \mathbb{R}_{\geq 0}$ be a cost function specifying desirable agent behavior in a given environment, e.g., according to an expert user or an optimal design. We assume that the agent does not have access to either the true semantic map m^* or the true cost function c^* . However, the agent is able to obtain point-cloud observations $\mathbf{P}_t = \{(\mathbf{p}_l, \mathbf{y}_l)\}_l \in \mathcal{P}$ at each step t , where \mathbf{p}_l is the measurement location. In the following sections, we consider $\mathbf{p}_l \in \mathbb{R}^2$ for MiniGrid experiments in Sect. 7 and $\mathbf{p}_l \in \mathbb{R}^3$ for CARLA experiments in Sect. 8.

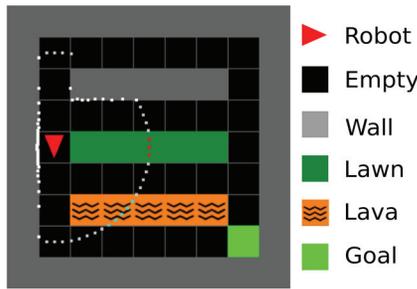


Fig. 2 A 9×9 grid environment with cells from four semantic classes: empty, wall, lawn, lava. An autonomous agent (red triangle, facing down) starts from the top left corner and is heading towards the goal in the bottom right. The agent prefers traversing the lawn but dislikes lava. LiDAR points detect the semantic labels of the corresponding tiles (gray on empty, white on wall, purple on lawn and cyan on lava) (Color figure online)

The vector of weights $y_l = [y_l^1, \dots, y_l^K]^\top$, where $y_l^k \in \mathbb{R}$, indicates the likelihood that semantic class $k \in \mathcal{K} \setminus \{0\}$ was observed. For example, $y_l \in \mathbb{R}^K$ can be obtained from the softmax output of a semantic segmentation algorithm (Papan-dreou et al., 2015; Badrinarayanan et al., 2017; Chen et al., 2018) that predicts the semantic class of the corresponding measurement location p_l in an RGBD image. The observed point cloud P_t depends on the agent state x_t and the environment realization m^* .

3.2 Expert model

We assume the expert chooses a control according to a Boltzmann-rational policy (Ramachandran & Amir, 2007; Neu & Szepesvári, 2007) given the *true* cost c^* and the *true* environment m^* ,

We assume that an expert user or algorithm demonstrates desirable agent behavior in the form of a training set $\mathcal{D} := \{(x_{t,n}, u_{t,n}^*, P_{t,n}, x_{g,n})\}_{t=1, n=1}^{T_n, N}$. The training set consists of N demonstrated executions with different lengths T_n for $n \in \{1, \dots, N\}$. Each demonstration trajectory contains the agent states $x_{t,n}$, expert controls $u_{t,n}^*$, and sensor observations $P_{t,n}$ encountered during navigation to a goal state $x_{g,n}$.

The design of an IRL algorithm depends on a model of the stochastic control policy $\pi^*(u | x; c^*, m^*)$ used by the expert to generate the training data \mathcal{D} , given the *true* cost c^* and environment m^* . The state of the art relies on the MaxEnt model (Ziebart et al., 2008), which assumes that the expert minimizes the weighted sum of the stage cost $c^*(x, u; m^*)$ and the negative policy entropy over the agent trajectory.

We propose a new model of expert behavior to explain rational deviation from optimality. We assume that the expert is aware of the optimal value function:

$$Q^*(x_t, u_t; c^*, m^*)$$

$$\begin{aligned} &:= \min_{T, u_{t+1:T-1}} \sum_{k=t}^{T-1} c^*(x_k, u_k; m^*) \\ &\text{s.t. } x_{k+1} = f(x_k, u_k), x_T = x_g. \end{aligned} \quad (1)$$

but does not always choose strictly rational actions. Instead, the expert behavior is modeled as a Boltzmann policy over the optimal value function:

$$\begin{aligned} &\pi^*(u_t | x_t; c^*, m^*) \\ &= \frac{\exp(-\frac{1}{\alpha} Q^*(x_t, u_t; c^*, m^*))}{\sum_{u \in \mathcal{U}} \exp(-\frac{1}{\alpha} Q^*(x_t, u; c^*, m^*))} \end{aligned} \quad (2)$$

where α is a temperature parameter. The Boltzmann policy stipulates an exponential preference of controls that incur low long-term costs. We will show in Sect. 5 that this expert model allows very efficient policy search as well as computation of the policy gradient with respect to the stage cost, which is needed for inverse cost learning. In contrast, the MaxEnt policy requires either value iteration over the full state space (Ziebart et al., 2008) or sampling-based estimation of a partition function (Finn et al., 2016). Appendix A provides a comparison between our model and the MaxEnt formulation.

3.3 Problem statement

Given the training set \mathcal{D} , our goal is to:

- learn a cost function estimate $c_t : \mathcal{X} \times \mathcal{U} \times \mathcal{P}^t \times \Theta \rightarrow \mathbb{R}_{\geq 0}$ that depends on an observation sequence $P_{1:t}$ from the true latent environment and is parameterized by $\theta \in \Theta$,
- design a stochastic policy π_t from c_t such that the agent behavior under π_t matches the demonstrations in \mathcal{D} .

The optimal value function corresponding to a stage cost estimate c_t is:

$$\begin{aligned} &Q_t(x_t, u_t; P_{1:t}, \theta) \\ &:= \min_{T, u_{t+1:T-1}} \sum_{k=t}^{T-1} c_t(x_k, u_k; P_{1:t}, \theta) \\ &\text{s.t. } x_{k+1} = f(x_k, u_k), x_T = x_g. \end{aligned} \quad (3)$$

Following the expert model proposed in Sect. 3.2, we define a Boltzmann policy corresponding to Q_t :

$$\begin{aligned} &\pi_t(u_t | x_t; P_{1:t}, \theta) \\ &\propto \exp\left(-\frac{1}{\alpha} Q_t(x_t, u_t; P_{1:t}, \theta)\right) \end{aligned} \quad (4)$$

and aim to optimize the stage cost parameters θ to match the demonstrations in \mathcal{D} .

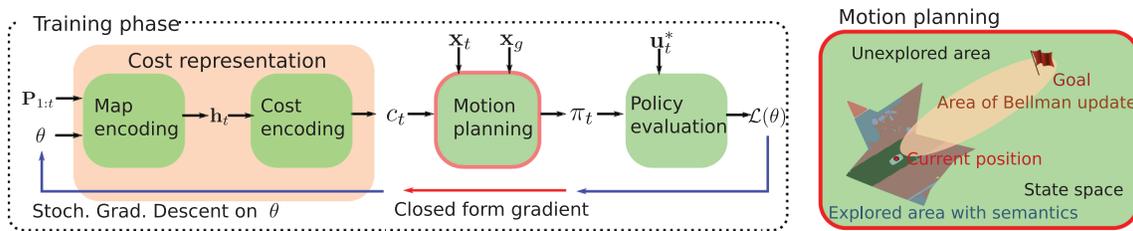


Fig. 3 Architecture for cost function learning from demonstrations with semantic observations. Our main contribution is a cost representation, combining a probabilistic *semantic map encoder*, with recurrent dependence on semantic observations $\mathbf{P}_{1:t}$, and a *cost encoder*, defined

Problem 1 Given demonstrations \mathcal{D} , optimize the cost function parameters θ so that log-likelihood of the demonstrated controls $\mathbf{u}_{t,n}^*$ is maximized by policy functions $\pi_{t,n}$ obtained according to (4):

$$\min_{\theta} \mathcal{L}(\theta) := - \sum_{n=1}^N \sum_{t=1}^{T_n} \log \pi_{t,n}(\mathbf{u}_{t,n}^* | \mathbf{x}_{t,n}; \mathbf{P}_{1:t,n}, \theta). \quad (5)$$

The problem setup is illustrated in Fig. 3. An important consequence of our expert model is that the computation of the optimal value function corresponding to a given stage cost estimate is a standard deterministic shortest path (DSP) problem (Bertsekas, 1995). However, the challenge is to make the value function computation differentiable with respect to the cost parameters θ in order to propagate the loss in (5) back through the DSP problem to update θ . Once the parameters are optimized, the associated agent behavior can be generalized to navigation tasks in new partially observable environments by evaluating the cost c_t based on the observations $\mathbf{P}_{1:t}$ iteratively and re-computing the associated policy π_t .

4 Cost function representation

We propose a cost function representation with two components: a *semantic occupancy map encoder* with parameters Ψ and a *cost encoder* with parameters ϕ . The model is differentiable by design, allowing its parameters to be optimized by the subsequent planning algorithm described in Sect. 5.

4.1 Semantic occupancy map encoder

We develop a semantic occupancy map that stores the likelihood of the different semantic categories in \mathcal{K} in different areas of the map. We discretize the state space \mathcal{X} into J cells and let $\mathbf{m} = [m^1, \dots, m^J]^T \in \mathcal{K}^J$ be an a priori unknown vector of true semantic labels over the cells. Given the agent states $\mathbf{x}_{1:t}$ and observations $\mathbf{P}_{1:t}$ over time, our

model maintains the semantic occupancy posterior $\mathbb{P}(\mathbf{m} = \mathbf{k} | \mathbf{x}_{1:t}, \mathbf{P}_{1:t})$, where $\mathbf{k} = [k^1, \dots, k^J]^T \in \mathcal{K}^J$. The representation complexity may be reduced significantly if one assumes independence among the map cells m^j : $\mathbb{P}(\mathbf{m} = \mathbf{k} | \mathbf{x}_{1:t}, \mathbf{P}_{1:t}) = \prod_{j=1}^J \mathbb{P}(m^j = k^j | \mathbf{x}_{1:t}, \mathbf{P}_{1:t})$.

We generalize the binary occupancy grid mapping algorithm (Thrun et al., 2005; Hornung et al., 2013) to obtain incremental Bayesian updates for the multi-class probability at each cell m^j . In detail, at time $t - 1$, we maintain a vector $\mathbf{h}_{t-1,j}$ of class log-odds at each cell and update them given the observation \mathbf{P}_t obtained from state \mathbf{x}_t at time t .

Definition 1 The vector of class log-odds associated with cell m^j at time t is $\mathbf{h}_{t,j} = [h_{t,j}^0, \dots, h_{t,j}^K]^T$ with elements:

$$h_{t,j}^k := \log \frac{\mathbb{P}(m^j = k | \mathbf{x}_{1:t}, \mathbf{P}_{1:t})}{\mathbb{P}(m^j = 0 | \mathbf{x}_{1:t}, \mathbf{P}_{1:t})} \text{ for } k \in \mathcal{K}. \quad (6)$$

Note that by definition, $h_{t,j}^0 = 0$. Applying Bayes rule to (6) leads to a recursive Bayesian update for the log-odds vector:

$$\begin{aligned} h_{t,j}^k &= h_{t-1,j}^k + \log \frac{p(\mathbf{P}_t | m^j = k, \mathbf{x}_t)}{p(\mathbf{P}_t | m^j = 0, \mathbf{x}_t)} \\ &= h_{t-1,j}^k + \sum_{(p_l, y_l) \in \mathbf{P}_t} \left(\log \frac{\mathbb{P}(m^j = k | \mathbf{x}_t, (p_l, y_l))}{\mathbb{P}(m^j = 0 | \mathbf{x}_t, (p_l, y_l))} - h_{0,j}^k \right), \end{aligned} \quad (7)$$

where $p(\mathbf{P}_t | m^j = k, \mathbf{x}_t)$ is the likelihood of observing \mathbf{P}_t from agent state \mathbf{x}_t when cell m^j has semantic label k . Here, we assume that the observations $(p_l, y_l) \in \mathbf{P}_t$ at time t , given the cell m^j and state \mathbf{x}_t , are independent among each other and of the previous observations $\mathbf{P}_{1:t-1}$. The semantic class posterior can be recovered from the log-odds vector $\mathbf{h}_{t,j}$ via a softmax function $\mathbb{P}(m^j = k | \mathbf{x}_{1:t}, \mathbf{P}_{1:t}) = \sigma^k(\mathbf{h}_{t,j})$, where $\sigma : \mathbb{R}^{K+1} \rightarrow \mathbb{R}^{K+1}$ satisfies:

$$\begin{aligned} \sigma(\mathbf{z}) &= [\sigma^0(\mathbf{z}), \dots, \sigma^K(\mathbf{z})]^T, \\ \sigma^k(\mathbf{z}) &= \frac{\exp(z^k)}{\sum_{k' \in \mathcal{K}} \exp(z^{k'})}, \end{aligned}$$

$$\log \frac{\sigma^k(\mathbf{z})}{\sigma^{k'}(\mathbf{z})} = z^k - z^{k'}. \tag{8}$$

To complete the Bayesian update in (7), we propose a parametric inverse observation model, $\mathbb{P}(m^j = k \mid \mathbf{x}_t, (\mathbf{p}_l, \mathbf{y}_l))$, relating the class likelihood of map cell m^j to a labeled point $(\mathbf{p}_l, \mathbf{y}_l)$ obtained from state \mathbf{x}_t .

Definition 2 Consider a labeled point $(\mathbf{p}_l, \mathbf{y}_l)$ observed from state \mathbf{x}_t . Let $\mathcal{J}_{t,l} \subset \{1, \dots, J\}$ be the set of map cells intersected by the sensor ray from \mathbf{x}_t toward \mathbf{p}_l . Let m^j be an arbitrary map cell and $d(\mathbf{x}, m^j)$ be the distance between \mathbf{x} and the center of mass of m^j . Define the inverse observation model of the class label of cell m^j as:

$$\mathbb{P}(m^j = k \mid \mathbf{x}_t, (\mathbf{p}_l, \mathbf{y}_l)) = \begin{cases} \sigma^k(\Psi_l \bar{\mathbf{y}}_l \delta p_{t,l,j}), & \delta p_{t,l,j} \leq \epsilon, j \in \mathcal{J}_{t,l} \\ \sigma^k(\mathbf{h}_{0,j}), & \text{otherwise,} \end{cases} \tag{9}$$

where $\Psi_l \in \mathbb{R}^{(K+1) \times (K+1)}$ is a learnable parameter matrix, $\delta p_{t,l,j} := d(\mathbf{x}_t, m^j) - \|\mathbf{p}_l - \mathbf{x}_t\|_2$, $\epsilon > 0$ is a hyperparameter (e.g., set to half the size of a cell), and $\bar{\mathbf{y}}_l := [0, \mathbf{y}_l^\top]^\top$ is augmented with a trivial observation for the “free” class.

Intuitively, the inverse observation model specifies that cells intersected by the sensor ray are updated according to their distance to the ray endpoint and the detected semantic class probability, while the class likelihoods of other cells remain unchanged and equal to the prior. For example, if m^j is intersected, the likelihood of the class label is determined by a softmax squashing of a linear transformation of the measurement vector \mathbf{y}_l with parameters Ψ_l , scaled by the distance $\delta p_{t,l,j}$. Otherwise, Definition 2 specifies an uninformative class likelihood in terms of the prior log-odds vector $\mathbf{h}_{0,j}$ of cell m^j (e.g., $\mathbf{h}_{0,j} = \mathbf{0}$ specifies a uniform prior over the semantic classes).

Definition 3 The log-odds vector of the inverse observation model associated with cell m^j and point observation $(\mathbf{p}_l, \mathbf{y}_l)$ from state \mathbf{x}_t is $\mathbf{g}_j(\mathbf{x}_t, (\mathbf{p}_l, \mathbf{y}_l))$ with elements:

$$g_j^k(\mathbf{x}_t, (\mathbf{p}_l, \mathbf{y}_l)) = \log \frac{\mathbb{P}(m^j = k \mid \mathbf{x}_t, (\mathbf{p}_l, \mathbf{y}_l))}{\mathbb{P}(m^j = 0 \mid \mathbf{x}_t, (\mathbf{p}_l, \mathbf{y}_l))}. \tag{10}$$

The log-odds vector of the inverse observation model, \mathbf{g}_j , specifies the increment for the Bayesian update of the cell log-odds $\mathbf{h}_{t,j}$ in (7). Using the softmax properties in (8) and Definition 2, we can express \mathbf{g}_j as:

$$\mathbf{g}_j(\mathbf{x}_t, (\mathbf{p}_l, \mathbf{y}_l)) = \begin{cases} \Psi_l \bar{\mathbf{y}}_l \delta p_{t,l,j}, & \delta p_{t,l,j} \leq \epsilon, j \in \mathcal{J}_{t,l} \\ \mathbf{h}_{0,j}, & \text{otherwise.} \end{cases} \tag{11}$$

Note that the inverse observation layer model definition in (9) resembles a single neural network layer. One can also specify

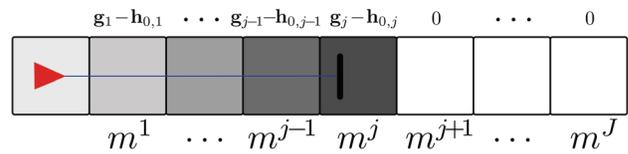


Fig. 4 Illustration of the log-odds update in (13) for a single point observation. The sensor ray (blue) hits an obstacle (black) in cell m^j . The log-odds increment $\mathbf{g}_j - \mathbf{h}_{0,j}$ on each cell is shown in grayscale (Color figure online)

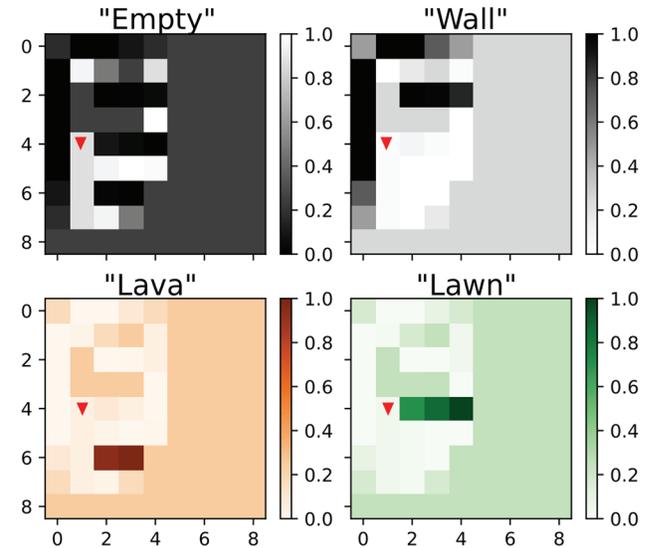


Fig. 5 The semantic occupancy probability of each class for the example in Fig. 2. Using the map encoder described in Sect. 4.1, the semantic categories (wall, lawn, lava, etc.) can be identified correctly after training

a more expressive multi-layer neural network that maps the observation \mathbf{y}_l and the distance differential $\delta p_{t,l,j}$ along the l -th ray to the log-odds vector:

$$\mathbf{g}_j(\mathbf{x}_t, (\mathbf{p}_l, \mathbf{y}_l); \Psi_l) = \begin{cases} \text{NN}(\bar{\mathbf{y}}_l, \delta p_{t,l,j}; \Psi_l) & \delta p_{t,l,j} \leq \epsilon, j \in \mathcal{J}_{t,l} \\ \mathbf{h}_{0,j} & \text{otherwise.} \end{cases} \tag{12}$$

Proposition 1 Given a labeled point cloud $\mathbf{P}_t = \{(\mathbf{p}_l, \mathbf{y}_l)\}_l$ obtained from state \mathbf{x}_t at time t , the Bayesian update of the log-odds vector of any map cell m^j is:

$$\mathbf{h}_{t,j} = \mathbf{h}_{t-1,j} + \sum_{(\mathbf{p}_l, \mathbf{y}_l) \in \mathbf{P}_t} [\mathbf{g}_j(\mathbf{x}_t, (\mathbf{p}_l, \mathbf{y}_l)) - \mathbf{h}_{0,j}]. \tag{13}$$

Figure 4 illustrates the increment of the log-odds vector $\mathbf{h}_{t,j}$ for a single point $(\mathbf{p}_l, \mathbf{y}_l)$. The log-odds are increased more at m^j than other cells far away from the observed point. When ϵ in (12) is set to half the cell size, values of the cells beyond the observed point are unchanged. Figure 5 shows

the semantic class probability prediction for the example in Fig. 2 using the inverse observation model in Definition 2 and the log-odds update in (13).

4.2 Cost encoder

We also develop a cost encoder that uses the semantic occupancy log odds \mathbf{h}_t to define a cost function estimate $c_t(\mathbf{x}, \mathbf{u})$ at a given state-control pair (\mathbf{x}, \mathbf{u}) . A convolutional neural network (CNN) (Goodfellow et al., 2016) with parameters ϕ can extract cost features from the multi-class occupancy map: $c_t = \text{CNN}(\mathbf{h}_t; \phi)$. We adopt a fully convolutional network (FCN) architecture (Badrinarayanan et al., 2017) to parameterize the cost function over the semantic class probabilities. The model is a multi-scale architecture that performs down-samples and up-samples to extract feature maps at different layers. Features from multiple scales ensure that the cost function is aware of both local and global context from the semantic map posterior. FCNs are also translation equivariant (Cohen & Welling, 2016), ensuring that map regions of the same semantic class infer the same cost, irrespective of the specific locations of those regions. Our model architecture (illustrated in Fig. 7) consists of a series of convolutional layers with 32 channels, batch normalization (Ioffe & Szegedy, 2015) and ReLU layers, followed by a max-pooling layer with 2×2 window with stride 2. The feature maps go through another series of convolutional layers with 64 channels, batch normalization, ReLU and max-pooling layers before they are upsampled by reusing the max-pooling indices. The feature maps then go through two series of upsampling, convolution, batch normalization and ReLU layers to produce the final cost function c_t . We add a small positive constant to the ReLU output to ensure that $c_t > 0$ and there are no negative cycles or cost-free paths during planning.

In summary, the semantic map encoder (parameterized by $\{\Psi_l\}_l$) takes the agent state history $\mathbf{x}_{1:t}$ and point cloud observation history $\mathbf{P}_{1:t}$ as inputs to encode a semantic map probability as discussed in Sect. 4.1. The FCN cost encoder (parameterized by ϕ) in turn defines a cost function from the extracted semantic features. The learnable parameters of the cost function, $c_t(\mathbf{x}, \mathbf{u}; \mathbf{P}_{1:t}, \theta)$, are $\theta = \{\{\Psi_l\}_l, \phi\}$.

5 Cost learning via differentiable planning

We focus on optimizing the parameters θ of the cost representation $c_t(\mathbf{x}, \mathbf{u}; \mathbf{P}_{1:t}, \theta)$ developed in Sect. 4. Since the true cost c^* is not directly observable, we need to differentiate the loss function $\mathcal{L}(\theta)$ in (5), which, in turn, requires differentiating through the DSP problem in (3) with respect to the cost function estimate c_t .

Previous works rely on dynamic programming to solve the DSP problem in (3). For example, the VIN model (Tamar

et al., 2016) approximates T iterations of the value iteration algorithm by a neural network with T convolutional and minpooling layers. This allows VIN to be differentiable with respect to the stage cost c_t but it scales poorly with the size of the problem due to the full Bellman backups (convolutions and minpooling) over the state and control space. We observe that it is not necessary to determine the optimal cost-to-go $Q_t(\mathbf{x}, \mathbf{u})$ at every state $\mathbf{x} \in \mathcal{X}$ and control $\mathbf{u} \in \mathcal{U}$. Instead of dynamic programming, a motion planning algorithm, such as a variant of A* (Likhachev et al., 2004) or RRT (LaValle, 1998; Karaman & Frazzoli, 2011), may be used to solve problem (3) efficiently and determine the optimal cost-to-go $Q_t(\mathbf{x}, \mathbf{u})$ only over a subset of promising states. The subgradient method of Shor (2012); Ratliff et al. (2006) may then be employed to obtain the subgradient of $Q_t(\mathbf{x}_t, \mathbf{u}_t)$ with respect to c_t along the optimal path.

5.1 Deterministic shortest path

Given a cost estimate c_t , we use the A* algorithm (Algorithm 1) to solve the DSP problem in (3) and obtain the optimal cost-to-go Q_t . The algorithm starts the search from the goal state \mathbf{x}_g and proceeds backwards towards the current state \mathbf{x}_t . It maintains an *OPEN* set of states, which may potentially lie along a shortest path, and a *CLOSED* list of states, whose optimal value $\min_{\mathbf{u}} Q_t(\mathbf{x}, \mathbf{u})$ has been determined exactly. At each iteration, the algorithm pops a state \mathbf{x} from *OPEN* with the smallest $g(\mathbf{x}) + \epsilon h(\mathbf{x}_t, \mathbf{x})$ value, where $g(\mathbf{x})$ is an estimate of the cost-to-go from \mathbf{x} to \mathbf{x}_g and $h(\mathbf{x}_t, \mathbf{x})$ is a heuristic function that does not overestimate the true cost from \mathbf{x}_t to \mathbf{x} and satisfies the triangle inequality. We find all predecessor states \mathbf{x}' and their corresponding control \mathbf{u}' that lead to \mathbf{x} under the known dynamics model $\mathbf{x} = f(\mathbf{x}', \mathbf{u}')$ and update their g values if there is a lower cost trajectory from \mathbf{x}' to \mathbf{x}_g through \mathbf{x} . The algorithm terminates when all neighbors of the current state \mathbf{x}_t are in the *CLOSED* set. The following relations are satisfied at any time throughout the search:

$$Q_t(\mathbf{x}, \mathbf{u}) = c_t(\mathbf{x}, \mathbf{u}) + g(f(\mathbf{x}, \mathbf{u})), \forall f(\mathbf{x}, \mathbf{u}) \in \text{CLOSED},$$

$$Q_t(\mathbf{x}, \mathbf{u}) \leq c_t(\mathbf{x}, \mathbf{u}) + g(f(\mathbf{x}, \mathbf{u})), \forall f(\mathbf{x}, \mathbf{u}) \notin \text{CLOSED}.$$

The algorithm terminates only after all neighbors $f(\mathbf{x}_t, \mathbf{u})$ of the current state \mathbf{x}_t are in *CLOSED* to guarantee that the optimal cost-to-go $Q_t(\mathbf{x}_t, \mathbf{u})$ at \mathbf{x}_t is exact. A simple choice of heuristic that guarantees the above relations is $h(\mathbf{x}, \mathbf{x}') = 0$, which reduces A* to Dijkstra's algorithm. Alternatively, the cost encoder output may be designed to ensure that $c_t(\mathbf{x}, \mathbf{u}) \geq 1$, which allows using Manhattan distance, $h(\mathbf{x}, \mathbf{x}') = \|\mathbf{x} - \mathbf{x}'\|_1$, as the heuristic.

Algorithm 1: A* motion planning

```

1 Function Plan ( $x_t, x_g, c_t, h, \epsilon$ ):
2    $OPEN \leftarrow \{x_t\}, CLOSED \leftarrow \{\}$ 
3    $g(x) \leftarrow \infty, \forall x \in \mathcal{X}, g(x_t) \leftarrow 0$ 
4   while  $\exists u \in \mathcal{U}$  s.t.  $f(x_t, u) \notin CLOSED$  do
5     Remove  $x$  from  $OPEN$  with smallest  $g(x) + \epsilon h(x_t, x)$ 
     and insert in  $CLOSED$ 
6     for  $(x', u') \in \text{Predecessors}(x)$  do
7       if  $x' \notin CLOSED$  and  $g(x') > g(x) + c_t(x', u')$ 
       then
8          $g(x') \leftarrow g(x) + c_t(x', u')$ 
9          $CHILD(x') \leftarrow x$ 
10        if  $x' \in OPEN$  then
11          Update priority of  $x'$  with  $g(x') + \epsilon h(x_t, x')$ 
12        else
13           $OPEN \leftarrow OPEN \cup \{x'\}$ 
14    return  $g(f(x_t, u)) \forall u \in \mathcal{U}$ 
15 Function Predecessors ( $x$ ):
16 return  $\{(x', u') \in \mathcal{X} \times \mathcal{U} \mid x = f(x', u')\}$ 

```

Finally, a Boltzmann policy $\pi_t(u \mid x)$ can be defined using the g values returned by A* for any $x \in \mathcal{X}$:

$$\pi_t(u \mid x) \propto \exp\left(-\frac{1}{\alpha} (c_t(x, u) + g(f(x, u)))\right). \tag{14}$$

The policy discourages controls that lead to states outside of $CLOSED$ because $c_t(x, u) + g(f(x, u))$ overestimates $Q_t(x, u)$. For any unvisited states, the policy is uniform since g values are initialized to infinity. In practice, we only need to query the policy at the current state x_t , which is always in $CLOSED$, for the loss function $\mathcal{L}(\theta)$ in (5) during training and policy inference during testing.

5.2 Backpropagation through planning

Having solved the DSP problem in (3) for a fixed cost function c_t , we now discuss how to optimize the cost parameters θ such that the planned policy in (14) minimizes the loss in (5). Our goal is to compute the gradient $\frac{d\mathcal{L}(\theta)}{d\theta}$, using the chain rule, in terms of $\frac{\partial \mathcal{L}(\theta)}{\partial Q_t(x_t, u_t)}$, $\frac{\partial Q_t(x_t, u_t)}{\partial c_t(x, u)}$, and $\frac{\partial c_t(x, u)}{\partial \theta}$. The first gradient term can be obtained analytically from (5) and (4), as we show later, while the third one can be obtained via backpropagation (automatic differentiation) through the neural network cost model $c_t(x, u; P_{1:T}, \theta)$ developed in Sect. 4. We focus on computing the second gradient term.

We rewrite $Q_t(x_t, u_t)$ in a form that makes its subgradient with respect to $c_t(x, u)$ obvious. Let $\mathcal{T}(x_t, u_t)$ be the set of trajectories, $\tau = x_t, u_t, x_{t+1}, u_{t+1}, \dots, x_{T-1}, u_{T-1}, x_T$, of length $T - t + 1$ that start at x_t, u_t , satisfy transitions $x_{i+1} = f(x_i, u_i)$, and terminate at $x_T = x_g$. Let $\tau^* \in \mathcal{T}(x_t, u_t)$ be an optimal trajectory corresponding to the optimal cost-to-go $Q_t(x_t, u_t)$. Define a state-control vis-

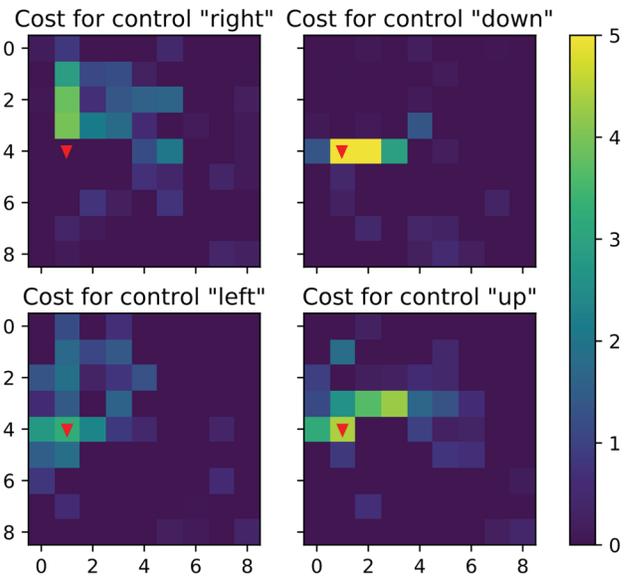


Fig. 6 Learned cost function for the example in Fig. 2. The cost of control “right” is the smallest at the agent’s location after training. The agent correctly predicts that it should move right and step on the lawn

itation function which counts the number of times transition (x, u) appears in τ :

$$\mu_\tau(x, u) := \sum_{k=t}^{T-1} \mathbb{1}_{(x_k, u_k)=(x, u)}. \tag{15}$$

The optimal cost-to-go $Q_t(x_t, u_t)$ can be viewed as a minimum over trajectories $\mathcal{T}(x_t, u_t)$ of the inner product between the cost function c_t and the visitation function μ_τ :

$$Q_t(x_t, u_t) = \min_{\tau \in \mathcal{T}(x_t, u_t)} \sum_{x \in \mathcal{X}, u \in \mathcal{U}} c_t(x, u) \mu_\tau(x, u), \tag{16}$$

where \mathcal{X} can be assumed finite because both T and \mathcal{U} are finite. We use the subgradient method (Shor, 2012; Ratliff et al., 2006) to compute a subgradient of $Q_t(x_t, u_t)$ with respect to c_t .

Lemma 1 Let $f(x, y)$ be differentiable and convex in x . Then, $\nabla_x f(x, y^*)$, where $y^* := \arg \min_y f(x, y)$, is a subgradient of the piecewise-differentiable convex function $g(x) := \min_y f(x, y)$.

Applying Lemma 1 to (16) leads to the following subgradient of the optimal cost-to-go function:

$$\frac{\partial Q_t(x_t, u_t)}{\partial c_t(x, u)} = \mu_{\tau^*}(x, u) \tag{17}$$

which can be obtained along the optimal trajectory τ^* by tracing the $CHILD$ relations returned by Algorithm 1. Figure 8 shows an illustration of this subgradient computation

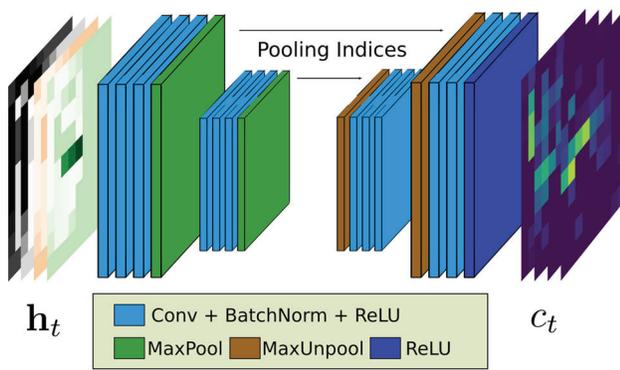


Fig. 7 A fully convolutional encoder-decoder neural network similar to that in Badrinarayanan et al. (2017) is used as the cost encoder to learn features from semantic map h_t to cost function c_t

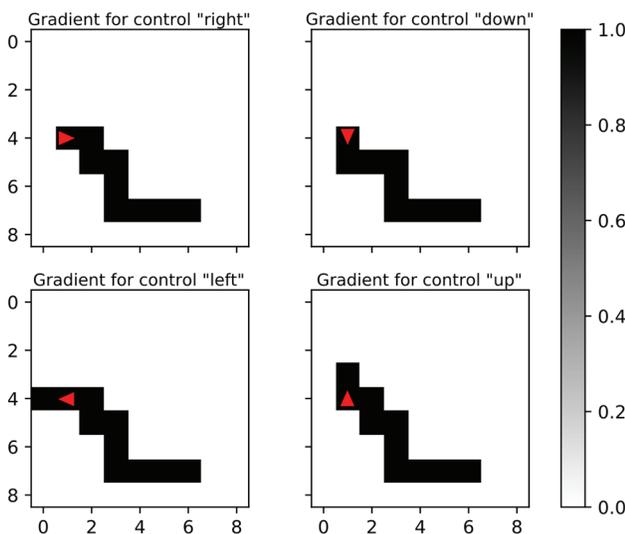


Fig. 8 Subgradient of the optimal cost-to-go $Q_t(x_t, u_t)$ for each control u_t with respect to the cost $c_t(x, u)$ in Fig. 6

with respect to the cost estimate in Fig. 6 for the example in Fig. 2. The result in (17) and the chain rule allow us to obtain a complete subgradient of $\mathcal{L}(\theta)$.

Proposition 2 A subgradient of the loss function $\mathcal{L}(\theta)$ in (5) with respect to θ can be obtained as:

$$\begin{aligned} \frac{\partial \mathcal{L}(\theta)}{\partial \theta} &= - \sum_{n=1}^N \sum_{t=1}^{T_n} \frac{d \log \pi_{t,n}(u_{t,n}^* | x_{t,n})}{d\theta} \\ &= - \sum_{n=1}^N \sum_{t=1}^{T_n} \sum_{u_{t,n} \in \mathcal{U}} \frac{d \log \pi_{t,n}(u_{t,n}^* | x_{t,n})}{d \log \pi_{t,n}(u_{t,n} | x_{t,n})} \frac{d Q_{t,n}(x_{t,n}, u_{t,n})}{d\theta} \\ &= - \sum_{n=1}^N \sum_{t=1}^{T_n} \sum_{u_{t,n} \in \mathcal{U}} \frac{1}{\alpha} \left(\mathbb{1}_{\{u_{t,n}=u_{t,n}^*\}} - \pi_{t,n}(u_{t,n} | x_{t,n}) \right) \\ &\quad \times \sum_{(x,u) \in \tau^*} \frac{\partial Q_{t,n}(x_{t,n}, u_{t,n})}{\partial c_t(x, u)} \frac{\partial c_t(x, u)}{\partial \theta} \end{aligned} \tag{18}$$

5.3 Algorithms

The computation graph implied by Proposition 2 is illustrated in Fig. 3. The graph consists of a cost representation layer and a differentiable planning layer, allowing end-to-end minimization of $\mathcal{L}(\theta)$ via stochastic subgradient descent. The training algorithm for solving Problem 1 is shown in Algorithm 2. The testing algorithm that enables generalizing the learned semantic mapping and planning behavior to new sensory data in new environments is shown in Algorithm 3.

Algorithm 2: Train cost parameters θ

Input: Dataset $\mathcal{D} = \{(x_{t,n}, u_{t,n}^*, P_{t,n}, x_{g,n})\}_{t=1, n=1}^{T_n, N}$

- 1 **while** θ not converged **do**
- 2 $\mathcal{L}(\theta) \leftarrow 0$
- 3 **for** $n = 1, \dots, N$ and $t = 1, \dots, T_n$ **do**
- 4 Update $c_{t,n}$ using $x_{t,n}$ and $P_{t,n}$ as in Sec. 4
- 5 Get $Q_{t,n}(x_{t,n}, u)$ from Alg. 1 with cost $c_{t,n}$
- 6 Get $\pi_{t,n}(u | x_{t,n})$ in (4) from $Q_{t,n}(x_{t,n}, u)$
- 7 $\mathcal{L}(\theta) \leftarrow \mathcal{L}(\theta) - \log \pi_{t,n}(u_{t,n}^* | x_{t,n})$
- 8 Update $\theta \leftarrow \theta - \eta \nabla \mathcal{L}(\theta)$ via Prop. 2

Output: Trained cost function parameters θ

Algorithm 3: Test control policy π_t

Input: Start state x_s , goal state x_g , cost parameters θ

- 1 Current state $x_t \leftarrow x_s$
- 2 **while** $x_t \neq x_g$ and navigation not failed **do**
- 3 Make an observation P_t
- 4 Update c_t using x_t and P_t as in Sec. 4
- 5 Get $Q_t(x_t, u)$ from Alg. 1 with cost c_t
- 6 Get $\pi_t(u | x_t)$ in (4) from $Q_t(x_t, u)$
- 7 $x_t \leftarrow f(x_t, u_t)$ with $u_t := \arg \max_u \pi_t(u | x_t)$

Output: Navigation succeeds or fails at x_t

6 Sparse tensor implementation

In this section, we propose a sparse tensor implementation of the map and cost variables introduced in Sect. 4. The region explored during a single navigation trajectory is usually a small subset of the full environment due to the agent’s limited sensing range. The map and cost variables $h_t, g_t, c_t(x, u)$ thus contains many 0 elements corresponding to “free” space or unexplored regions and only a small subset of the states in $c_t(x, u)$ are queried during planning and parameter optimization in Sect. 5. Representing these variables as dense matrices is computationally and memory inefficient. Instead, we propose an implementation of the map encoder and cost encoder that exploits the sparse structure of these matrices. Choy et al. (2019) developed the Minkowski Engine, an automatic differentiation neural network library for sparse tensors. This

library is tailored for our case as we require automatic differentiation for operations among the variables \mathbf{h}_t , \mathbf{g}_t , c_t in order to learn the cost parameters θ .

During training, we pre-compute the variable $\delta p_{t,l,j}$ over all points \mathbf{p}_l from a point cloud \mathbf{P}_t and all grid cells m^j . This results in a matrix $\mathbf{R}_t \in \mathbb{R}^{K \times J}$ where the entry corresponding to cell m^j stores the vector $\mathbf{y}_l \delta p_{t,l,j}$.¹ The matrix \mathbf{R}_t is then converted to COOrdinate list (COO) format (Tew, 2016), specifying the nonzero indices $\mathbf{C}_t \in \mathbb{R}^{N_{nz} \times 1}$ and their feature values $\mathbf{F}_t \in \mathbb{R}^{N_{nz} \times K}$, where $N_{nz} \ll J$ if \mathbf{R}_t is sparse. To construct \mathbf{C}_t and \mathbf{F}_t , we append non-zero features $\mathbf{y}_l \delta p_{t,l,j}$ to \mathbf{F}_t and their coordinates j in \mathbf{R}_t to \mathbf{C}_t . The inverse observation model log-odds \mathbf{g}_t can be computed from \mathbf{C}_t and \mathbf{F}_t via (11) and represented in COO format as well. Hence, a sparse representation of the semantic occupancy log-odds \mathbf{h}_t can be obtained by accumulating \mathbf{g}_t over time via (13).

We use the sparse tensor operations (e.g., convolution, batch normalization, pooling, etc.) provided by the Minkowski Engine in place of their dense tensor counterparts in the cost encoder defined in Sect. 4.2. For example, the convolution kernel does not slide sequentially over each entry in a dense tensor but is defined only over the indices in \mathbf{C}_t , skipping computations at the 0 elements. To ensure that the sparse tensors are compatible in the backpropagation step of the cost parameter learning (Sect. 5.2), the analytic subgradient in (18) should also be provided in sparse COO format. We implement a custom operation in which the forward function computes the cost-to-go $Q_t(x_t, \mathbf{u}_t)$ from $c_t(x, \mathbf{u})$ via Algorithm 1 and the backward function multiplies the sparse matrix $\frac{\partial Q_t(x_t, \mathbf{u}_t)}{\partial c_t(x, \mathbf{u})}$ with the previous gradient in the computation graph, $\frac{\partial \mathcal{L}(\theta)}{\partial Q_t(x_t, \mathbf{u}_t)}$, to get $\frac{\partial \mathcal{L}(\theta)}{\partial c_t(x, \mathbf{u})}$. The output gradient $\frac{\partial \mathcal{L}(\theta)}{\partial c_t(x, \mathbf{u})}$ is used as input to the downstream operations defined in Sects. 4.2 and 4.1 to update the cost parameters θ .

7 MiniGrid experiment

We first demonstrate our inverse reinforcement learning approach in a synthetic minigrid environment (Chevalier-Boisvert et al., 2018). We consider a simplified setting to help visualize and understand the differentiable semantic mapping and planning components.² A more realistic autonomous driving setting is demonstrated in Sect. 8.

¹ In our experiments, we found that storing only \mathbf{y}_l at the cell m^j where \mathbf{p}_l lies, instead of along the sensor ray, does not degrade performance.

² Our code for the minigrid experiments is open-sourced at <https://github.com/tianyudwang/sirl>.

7.1 Experiment setup

7.1.1 Environment

Grid environments of sizes 16×16 and 64×64 are generated by sampling a random number of random length rectangles with semantic labels from $\mathcal{K} := \{\text{empty}, \text{wall}, \text{lava}, \text{lawn}\}$. One such environment is shown in Fig. 9. The agent motion is modeled over a 4-connected grid such that a control \mathbf{u}_t from $\mathcal{U} := \{\text{up}, \text{down}, \text{left}, \text{right}\}$ causes a transition from \mathbf{x}_t to one of the four neighboring tiles \mathbf{x}_{t+1} . A *wall* tile is not traversable and a transition to it does not change the agent's position.

7.1.2 Sensor

At each step t , the agent receives 72 labeled points $\mathbf{P}_t = \{\mathbf{p}_l, \mathbf{y}_l\}_l$, obtained from ray-tracing a 360° field of view at angular resolution of 5° with maximum range of 3 grid cells and returning the grid location \mathbf{p}_l of the hit point and its semantic class encoded in a one-hot vector \mathbf{y}_l . See Fig. 2 for an illustration. The sensing range is smaller than the environment size, making the environment only partially observable at any given time.

7.1.3 Demonstrations

Expert demonstrations are obtained by running a shortest path algorithm on the true map \mathbf{m}^* , where the cost of arriving at an *empty*, *wall*, *lava*, or *lawn* tile is 1, 100, 10, 0.5, respectively. We generate 10,000, 1000, and 1000 random map configurations for training, validation, and testing, respectively. Start and goal locations are randomly assigned and maps without a feasible path are discarded. To avoid overfitting, we use the model parameters that perform best in validation for testing.

7.2 Models

7.2.1 DeepMaxEnt

We use the DeepMaxEnt IRL algorithm of Wulfmeier et al. (2016) as a baseline. DeepMaxEnt is an extension of the MaxEnt IRL algorithm (Ziebart et al., 2008), which uses a deep neural network to learn a cost function directly from LiDAR observations. In contrast to our model, DeepMaxEnt does not have an explicit map representation. The cost representation is a multi-scale FCN (Wulfmeier et al., 2016) adapted to the 16×16 and 64×64 domains. Value iteration over the cost matrix is approximated by a finite number of Bellman backup iterations, equal to the number of map cells. The original experiments in Wulfmeier et al. (2016) use the mean and variance of the height of 3D LiDAR

Table 1 Validation and test results for the 16×16 and 64×64 minigrid environments

Model	16×16				64×64			
	<i>NLL</i>	<i>Acc</i> (%)	<i>TSR</i> (%)	<i>MHD</i>	<i>NLL</i>	<i>Acc</i> (%)	<i>TSR</i> (%)	<i>MHD</i>
DeepMaxEnt	0.333	87.7	85.5	0.783	0.160	92.5	86.3	2.305
Ours	0.247	91.9	93.0	0.208	0.153	95.2	95.6	1.097

We report the negative log-likelihood (*NLL*) and prediction accuracy (*Acc*) of the validation set expert controls and the trajectory success rate (*TSR*) and modified Hausdorff distance (*MHD*) between the agent and the expert trajectories on the test set. See Sect. 7.3 for precise definitions of the metrics

points in each cell, as well as a binary indicator of cell visibility, as input features to the FCN neural network. Since our synthetic experiments are in 2D, the point count in each grid cell is used instead of the height mean and variance. This is a fair adaptation since Wulfmeier et al. (2016) argued that obstacles generally represent areas of larger height variance which corresponds to more points within obstacles cells for our observations. We compare against the original DeepMaxEnt model in Sect. 8.

7.2.2 Ours

Our model takes as inputs the semantic point cloud \mathbf{P}_t and the agent position \mathbf{x}_t at each time step and updates the semantic map probability via Sect. 4.1. The cost encoder goes through two scales of convolution and down(up)-sampling as introduced in Sect. 4.2. The models are trained using the Adam optimizer (Kingma & Ba, 2014) in Pytorch (Paszke et al., 2019). The neural network model training and online inference during testing are performed on an Intel i7-7700K CPU and an NVIDIA GeForce GTX 1080Ti GPU.

7.3 Evaluation metrics

The following metrics are used for evaluation: negative log-likelihood (*NLL*) and control accuracy (*Acc*) for the validation set and trajectory success rate (*TSR*) and modified Hausdorff distance (*MHD*) for the test set. Given learned cost parameters θ^* and a validation set $\mathcal{D} = \{(\mathbf{x}_{t,n}, \mathbf{u}_{t,n}^*, \mathbf{P}_{t,n}, \mathbf{x}_{g,n})\}_{t=1, n=1}^{T_n, N}$, policies $\pi_{t,n}(\cdot | \mathbf{x}_{t,n}; \mathbf{P}_{1:t,n}, \theta^*)$ are computed online via Algorithm 1 at each demonstrated state $\mathbf{x}_{t,n}$ and are evaluated according to:

$$NLL(\theta^*, \mathcal{D}) = -\frac{1}{\sum_{n=1}^N T_n} \sum_{n=1, t=1}^{N, T_n} \log \pi_{t,n}(\mathbf{u}_{t,n}^* | \mathbf{x}_{t,n}; \mathbf{P}_{1:t,n}, \theta^*) \quad (19)$$

$$Acc(\theta^*, \mathcal{D}) = \frac{1}{\sum_{n=1}^N T_n} \sum_{n=1, t=1}^{N, T_n} \mathbb{1} \left\{ \mathbf{u}_{t,n}^* = \arg \max_{\mathbf{u}} \pi_{t,n}(\mathbf{u} | \mathbf{x}_{t,n}; \mathbf{P}_{1:t,n}, \theta^*) \right\} \quad (20)$$

In the test set, the agent is initialized at the starting pose and iteratively applies control inputs $\mathbf{u}_{t,n} = \arg \max_{\mathbf{u}} \pi_{t,n}(\mathbf{u} | \mathbf{x}_{t,n}; \mathbf{P}_{1:t,n}, \theta^*)$ as described in Algorithm 3. The agent trajectories can deviate from expert trajectories and the agent has to recover from states which were not encountered by the expert. To find whether the agent eventually reaches the goal, we report the success rate *TSR*, where success is defined as reaching the goal state $\mathbf{x}_{g,n}$ within twice the number of steps of the expert trajectory. In addition, *MHD* compares how far the agent trajectories τ_n^A deviate from the expert trajectories τ_n^E :

$$MHD \left(\left\{ \tau_n^A \right\}, \left\{ \tau_n^E \right\} \right) = \frac{1}{N} \sum_{n=1}^N \max \left\{ \frac{1}{T^A} \sum_{t=1}^{T^A} d(\tau_{t,n}^A, \tau_n^E), \frac{1}{T^E} \sum_{t=1}^{T^E} d(\tau_{t,n}^E, \tau_n^A) \right\}, \quad (21)$$

where $d(\tau_{t,n}^A, \tau_n^E)$ is the minimum Euclidean distance from the state $\tau_{t,n}^A$ at time t in τ_n^A to any state in τ_n^E .

7.4 Results

The results are shown in Table 1. Our model outperforms DeepMaxEnt in every metric. Specifically, low *NLL* on the validation set indicates that map encoder and cost encoder in our model are capable of learning a cost function that matches the expert demonstrations. During testing in unseen environment configurations, our model also achieves a higher score in successfully reaching the goal. In addition, the difference in the agent trajectory and the expert trajectory is smaller, as measured by the *MHD* metric.

The outputs of our model components, i.e., map encoder, cost encoder and subgradient computation, are visualized in Fig. 9. The map encoder integrates past observations and holds a correct estimate of the semantic probability of each cell. The subgradients in the last column enable us to propagate the negative log-likelihood of the expert controls back to the cost model parameters. The cost visualizations indicate that the learned cost function correctly assigns higher costs to *wall* and *lava* cells (in brighter scale) and lower costs to *lawn* cells (in darker scale).

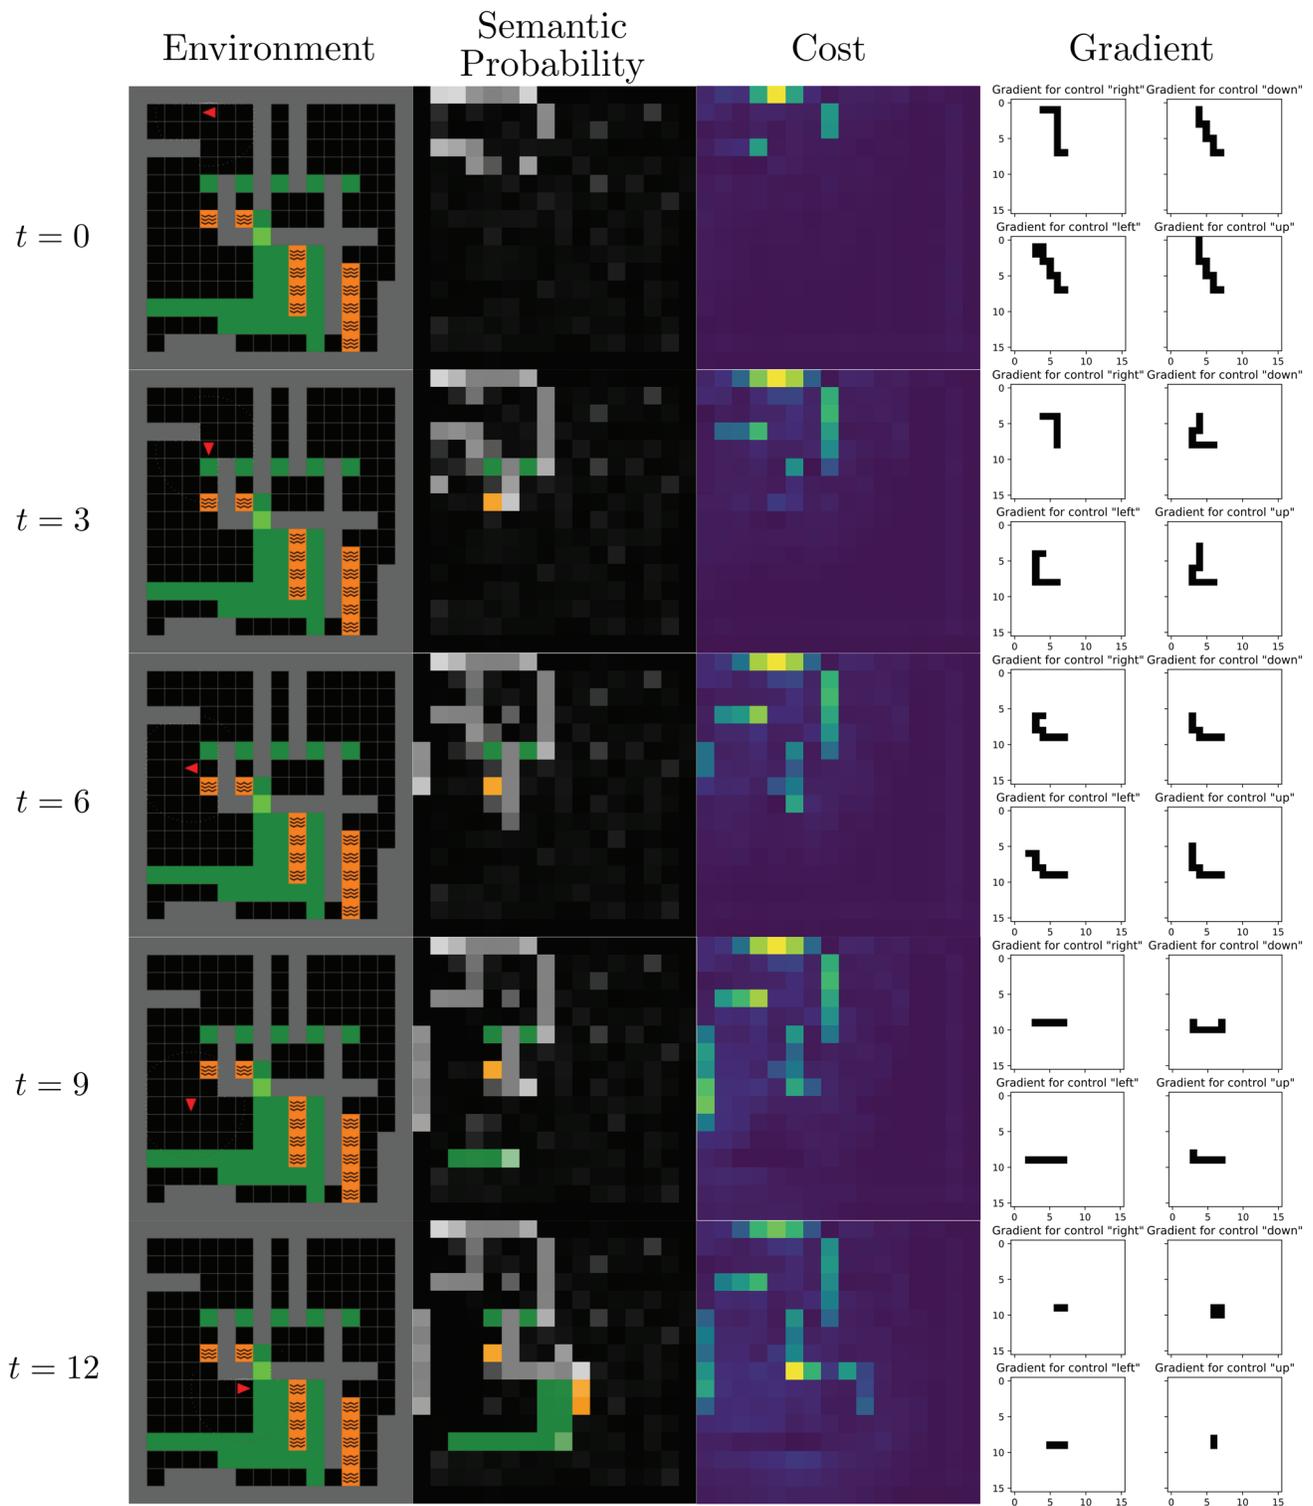


Fig. 9 Examples of probabilistic multi-class occupancy estimation, cost encoder output, and subgradient computation. The first column shows the agent in the true environment at different time steps. The second column shows the semantic occupancy estimates of the different cells. The third column shows the predicted cost of arriving at each cell.

Note that the learned cost function correctly assigns higher costs (in brighter scale) to *wall* and *lava* cells and lower costs (in darker scale) to *lawn* cells. The last column shows subgradients obtained via 17 during backpropagation to update the cost parameters (Color figure online)

Table 2 Average inference speed comparison between our model and DeepMaxEnt for predicting one control in testing

Grid size	16 × 16	64 × 64
DeepMaxEnt	5.8 ms	19.7 ms
Ours	2.7 ms	3.1 ms

7.5 Inference speed

The problem setting in this paper requires the agent to replan at each step when a new observation P_t arrives and updates the cost function c_t . Our planning algorithm is computationally efficient because it searches only through a subset of promising states to obtain the optimal cost-to-go $Q_t(x_t, u_t)$. On the other hand, the value iteration in DeepMaxEnt has to perform Bellman backups on the entire state space even though most of the environment is not visited and the cost in these unexplored regions is inaccurate. Table 2 shows the average inference speed to predict a new control u_t at each step during testing.

8 CARLA experiment

Building on the insights developed in the 2D minigrid environment in Sect. 7, we design an experiment in a realistic autonomous driving simulation.

8.1 Experiment setting

8.1.1 Environment

We evaluate our approach using the CARLA simulator (0.9.9) (Dosovitskiy et al., 2017), which provides high-fidelity autonomous vehicle simulation in urban environments. Demonstration data is collected from maps {Town01, Town02, Town03}, while Town04 is used for validation and Town05 for testing. Town05 includes different street layouts (e.g., intersections, buildings and freeways) and is larger than the training and validation maps.

8.1.2 Sensors

The vehicle is equipped with a LiDAR sensor that has 20 meters maximum range and 360° horizontal field of view. The vertical field of view ranges from 0° (facing forward) to −40° (facing down) with 5° resolution. A total of 56000 LiDAR rays are generated per scan P_t and point measurements are returned only if a ray hits an obstacle (see Fig. 10). The vehicle is also equipped with 4 semantic segmentation cameras that detect 13 different classes, including *road*, *road line*, *sidewalk*, *vegetation*, *car*, *building*, etc. The cameras

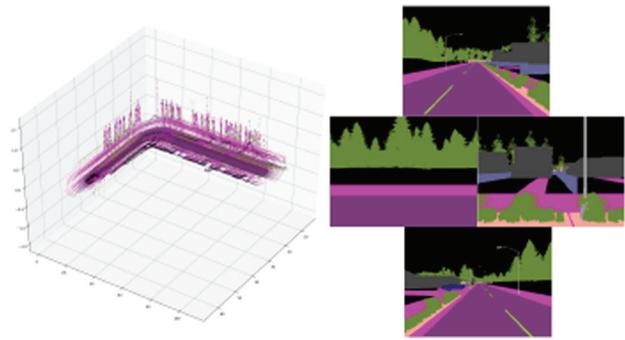


Fig. 10 Example of 3D LiDAR points and semantic segmentation camera facing four directions. The LiDAR points are annotated with semantic class labels

face front, left, right, and rear, each capturing a 90° horizontal field of view (see Fig. 10). The semantic label of each LiDAR point is retrieved by projecting the point in the camera’s frame and querying the pixel value in the segmented image.

8.1.3 Demonstrations

In each map, we collect 100 expert trajectories by running an autonomous navigation agent provided by the CARLA Python API. On the graph of all available waypoints, the expert samples two waypoints as start and goal and searches the shortest path as a list of waypoints. The expert uses a PID controller to generate a smooth and continuous trajectory to connect the waypoints on the shortest path. The expert respects traffic rules, such as staying on the road, and keeping in the current lane. The ground plane is discretized into a 256×256 grid of 0.5 meter resolution. Expert trajectories that do not fit in the given grid size are discarded. For planning purposes, the agent motion is modeled over a 4-connected grid with control space $\mathcal{U} := \{up, down, left, right\}$. A planned sequence of such controls is followed using the CARLA PID controller. Simulation features not related to the experiment are disabled, including spawning other vehicles and pedestrians, changing traffic signals and weather conditions, etc. Designing an agent that understands more complicated environment settings with other moving objects and changing traffic lights will be considered in future research.

8.2 Models and metrics

8.2.1 DeepMaxEnt

We use the DeepMaxEnt IRL algorithm Wulfmeier et al. (2016) with a multi-scale FCN cost encoder as a baseline again. Unlike the previous 2D experiment in Sect. 7, we use the input format from the original paper. Specifically,

observed 3D point clouds are mapped into a 2D grid with three channels: the mean and variance of the height of the points as well as the cell visibility of each cell. This model does not utilize the point cloud semantic labels.

8.2.2 DeepMaxEnt + Semantics

The input features are augmented with additional channels that contain the number of points in a cell of each particular semantic class. This model uses the additional semantic information but does not explicitly map the environment over time.

8.2.3 Ours

We ignore the height information in the 3D point clouds $P_{1:t}$ and maintain a 2D semantic map. The cost encoder is a two scale convolution and down(up)-sampling neural network, described in Sect. 4.2. Additionally, our model is implemented using sparse tensors, described in Sect. 6, to take advantage of the sparsity in the map h_t and cost c_t . The models are implemented using the Minkowski Engine (Choy et al., 2019) and the PyTorch library (Paszke et al., 2019) and are trained with the Adam optimizer (Kingma & Ba, 2014). The neural network training and the online inference during testing are performed on an Intel i7-7700K CPU and an NVIDIA GeForce GTX 1080Ti GPU.

8.2.4 Metrics

The metrics, *NLL*, *Acc*, *TSR*, and *MHD*, introduced in Sect. 7.3, are used for evaluation.

8.3 Results

Table 3 shows the performance of our model in comparison to DeepMaxEnt and DeepMaxEnt + Semantics. Our model learns to generate policies closest to the expert demonstrations in the validation map *Town04* by scoring best in *NLL* and *Acc* metrics. During testing in map *Town05*, the models predict controls at each step online to generate the agent trajectory. Ours achieves the highest success rate of reaching the goal without hitting sidewalks and other obstacles. Among the successful trajectories, Ours is also closest to the expert by achieving the minimum *MHD*. The results demonstrate that the map encoder captures both geometric and semantic information, allowing accurate cost estimation and generation of trajectories that match the expert behavior. Figure 11 shows an example of a generated trajectory during testing in the previously unseen *Town05* environment (also see Extension 1). The map encoder predicts correct semantic class labels for each cell and the cost encoder assigns higher costs to sidewalks than the road. We notice that the

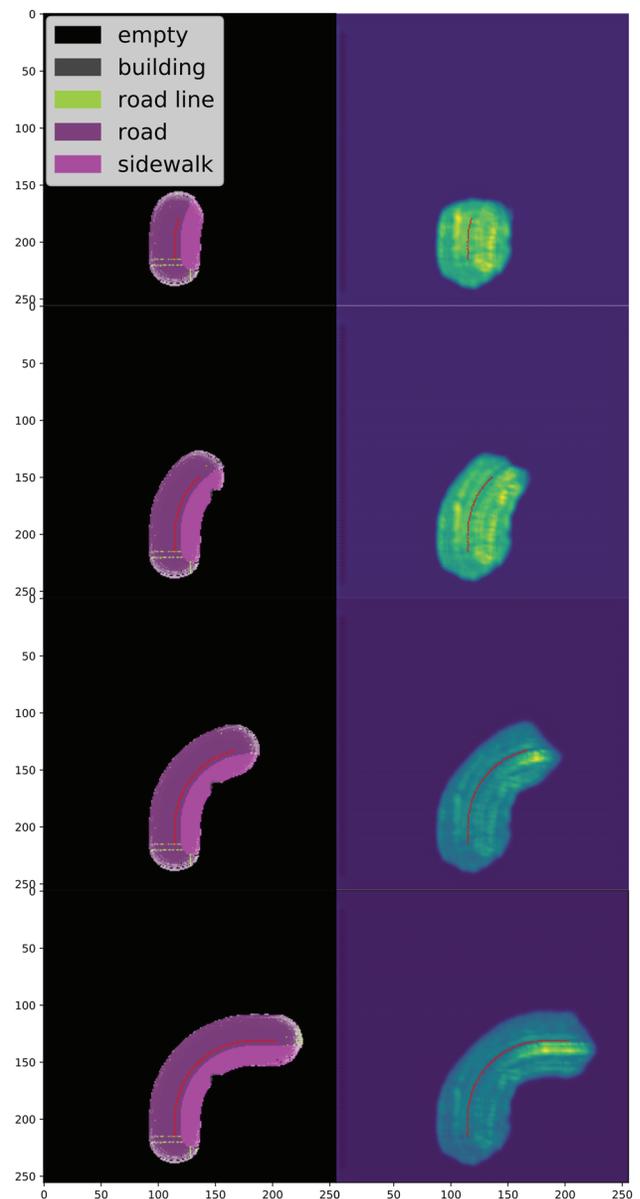


Fig. 11 Examples of semantic occupancy estimation and cost encoding during different steps in a test trajectory marked in red (also see Extension 1). The left column shows the most probable semantic class of the map encoder and the right column shows the cost to arrive at each state. Our model correctly distinguishes the road from other categories (e.g., sidewalk, building, etc) and assigns lower cost to road than sidewalks (Color figure online)

addition of semantic information actually degrades the performance of DeepMaxEnt. We conjecture that the increase in the number of input channels, due to the addition of the number of LiDAR points per category, makes the convolutional neural network layers prone to overfit on the training set but generalize poorly on the validation and test sets. Additional examples of agent trajectories and semantic mapping predictions are given in Online Resource 1. We also report runtime analysis for test-time model inference in Table 4.

Table 3 Test results from the CARLA *Town05* environment, including the negative log-likelihood (*NLL*) and prediction accuracy (*Acc.*) of the validation set expert controls and the trajectory success rate (*TSR*) and

Model	<i>NLL</i>	<i>Acc</i> (%)	<i>TSR</i> (%)	<i>MHD</i>
DeepMaxEnt	0.673	85.3	89	4.331
DeepMaxEnt + Semantics	0.742	82.6	87	4.752
Ours	0.406	94.2	93	2.538

Table 4 Runtime analysis of our model during testing in the CARLA simulator

Simulation	Data preprocessing	Model inference
39.3 ± 1.5 ms	14.9 ± 3.2 ms	6.4 ± 2.5 ms

We report per-step runtime averaged over 100 test trajectories

Each time step is divided into (1) simulator update, where the agent is set at new states and image and lidar observations are generated, (2) data preprocessing, where semantic labels are retrieved for point clouds and data are moved to GPU, and (3) model inference.

8.4 Evaluation with dynamic obstacles

In this section, we study the effects of dynamic obstacles in the scene on our model’s map and cost encoders. We create three different scenarios where the agent vehicle has to overtake a lower speed non-player character (NPC vehicle). In Scenario 1, the NPC is spawned in the left lane to the agent and 20 ms ahead. The agent is expected to stay in its own lane when overtaking the NPC. In Scenario 2, the NPC is spawned in the same lane as the agent and 20 ms ahead. The agent has to move to its left lane to overtake the NPC. Scenario 3 is a mixture of the first two where the NPC could be either in the same lane or in the left lane to the agent. A visualization of the first two scenarios is shown in Fig. 12. Training and evaluation are conducted in the *Town05* map since it contains multi-lane streets while other maps contain mostly single-lane streets. We sample 100 trajectories for testing within the top-left quadrant of the map and 200 trajectories for training from other quadrants as illustrated in Fig. 13. We train our model in all three scenarios and test it in the same scenario where it is trained. Each trajectory is discretized on a 128×128 grid of 1 meter resolution.

To effectively capture the most current information of the dynamic NPC vehicle, we multiply the grid log-odds $h_{t,j}$ with a decay rate $\gamma \in \{1.0, 0.9, 0.8, 0.7\}$, i.e.,

$$h_{t,j}^k = \gamma h_{t-1,j}^k + \log \frac{p(\mathbf{P}_t | m^j = k, \mathbf{x}_t)}{p(\mathbf{P}_t | m^j = 0, \mathbf{x}_t)}. \quad (22)$$

modified Hausdorff distance (*MHD*) between the agent and the expert trajectories on the test set

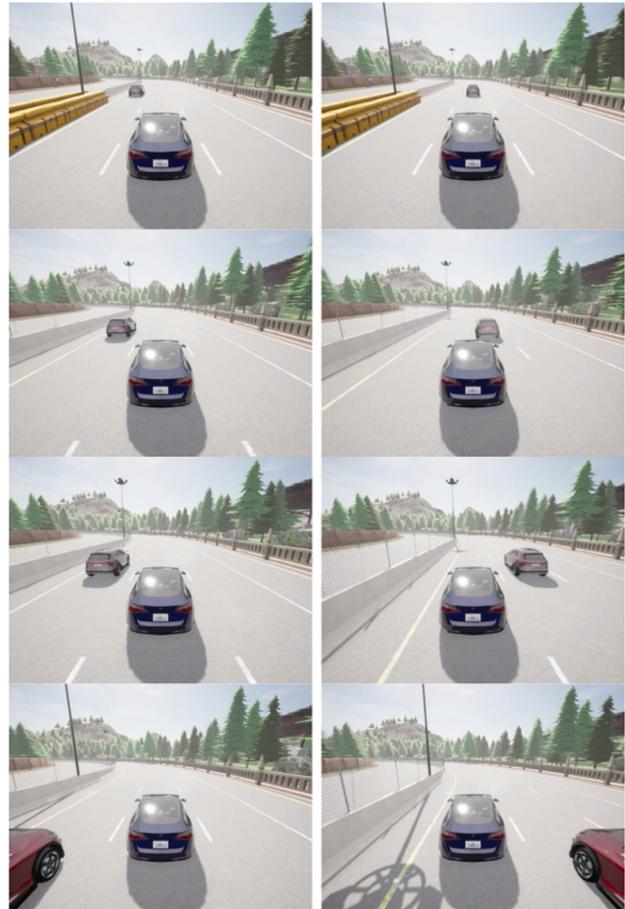


Fig. 12 Two scenarios with dynamic obstacles. Left column (scenario 1): the agent vehicle (blue) keeps in its own lane when overtaking the NPC vehicle (red) in the left lane. Right column (scenario 2): when the NPC and agent vehicles are spawned in the same lane, the agent switches to the left lane to overtake (Color figure online)

The map encoder is the same as in previous experiments when $\gamma = 1.0$, while when $\gamma < 1$ past observation information is slowly removed. Note that we use the same decay rate across all semantic classes since we do not assume prior knowledge of which classes are dynamic or static. Alternatively, it is possible to use a different decay rate for each class, $h_{t-1,j}^k \rightarrow \gamma^k h_{t-1,j}^k$, or set γ^k as a learnable parameter to be optimized with the overall objective in (5). The semantic probabilities

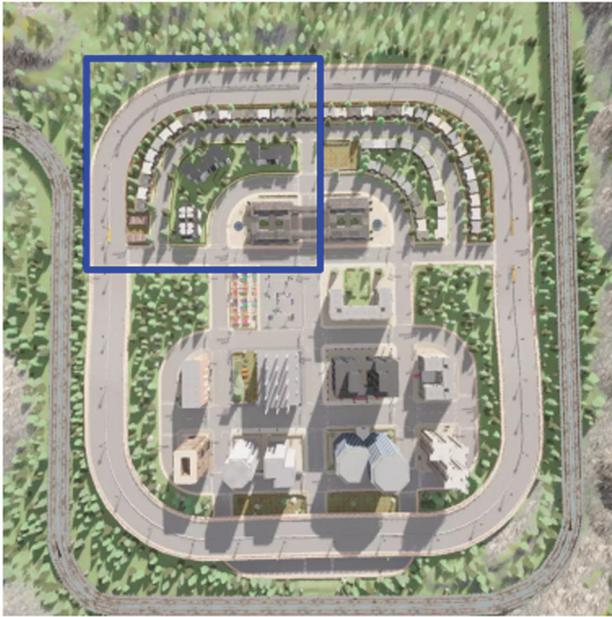


Fig. 13 Bird’s-eye view of the *Town05* map. The top-left quadrant is reserved for testing while training trajectories are sampled from other regions

expected as lane changing is a harder task when a moving NPC vehicle is present. The performance in the mixed scenario is on par with that in Scenario 2, suggesting that our policy class in (4) may not capture a multi-modal distribution in the demonstrated behaviors effectively. Within each scenario, we find that the model generally works better when the decay rate γ is close to 1.0. We suspect that since both vehicles are moving forward in the same direction, it does not hurt to map the NPC’s past locations. However, when γ is small, forgetting the NPC’s history makes its semantic probability smaller (as shown in Fig. 14), and thus the agent has a higher chance of colliding into the NPC vehicle. Finally, we find that MHD is consistent across all settings which shows that the agent trajectories are close to the expert’s, when they are successful.

8.5 Evaluation with noisy semantic observations

In this section, we study how noisy semantic observations can affect downstream cost prediction and policy inference. First, we consider noise added to the contours of each segmentation region. We replace each pixel value in the original 600×800 semantic segmentation image with a random pixel within its local 5×5 pixel window. This makes the segment boundaries blurry while the interior of each semantic region is unchanged (see Fig. 15b). With this noise model, only 0.2% of the lidar points are labeled incorrectly. We considered two additional noise models in which 2% and 20% of all pixels are randomly changed to an incorrect label chosen among the remaining semantic labels (see Fig. 15c, d).

We find that these noise models have very little influence on the policy inference. To understand this, we study how much the map encoder output changes when using noisy semantic images. We calculate the total variation distance between the semantic map probabilities obtained from the original and perturbed semantic images. Specifically, let $\mathbb{P}_a := \mathbb{P}(m^j = k \mid \mathbf{x}_{1:T}, \mathbf{P}_{1:T})$ be the semantic posterior probability of a trajectory using the original semantic segmentation images and, correspondingly, let $\mathbb{P}_b, \mathbb{P}_c, \mathbb{P}_d$ denote the posteriors using perturbed images from Fig. 15. The total variation distance between two discrete probability measures is

$$TV(\mathbb{P}_a, \mathbb{P}_b) := \frac{1}{2} \sum_k |\mathbb{P}_a(m^j = k) - \mathbb{P}_b(m^j = k)| \quad (23)$$

Table 6 and Fig. 16 show the maximum and a histogram, respectively, of the total variation across all grid cells. These results show that our map encoder is robust to noise even when 20% of the labels in the semantic images are wrong.

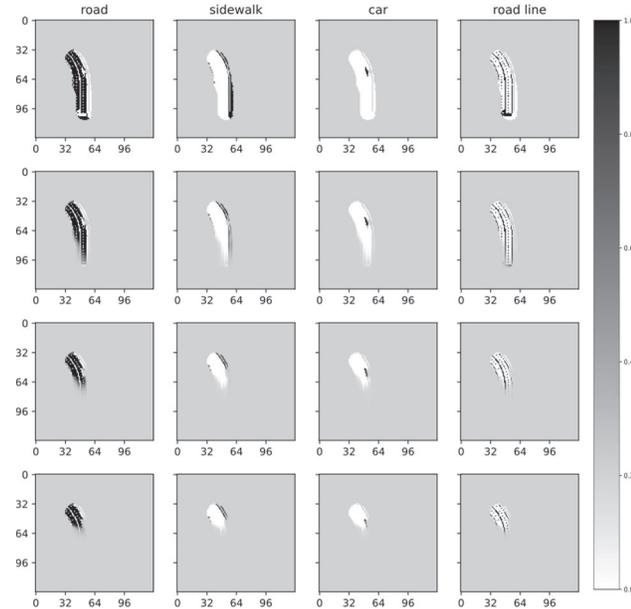


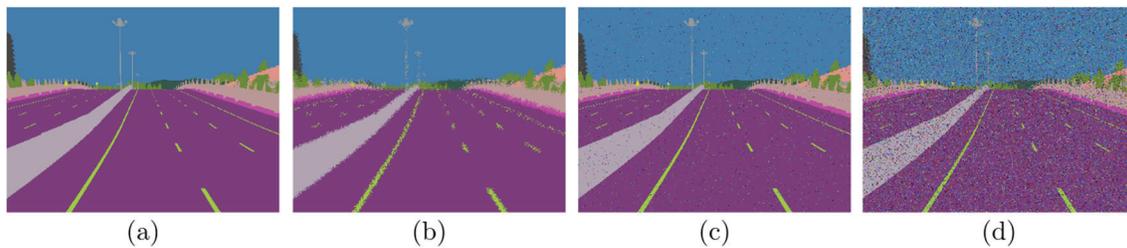
Fig. 14 Semantic probability of each class with different decay rate $\gamma \in \{1, 0.9, 0.8, 0.7\}$

and cost encoder output of the same trajectory with different decay rates is shown in Fig. 14.

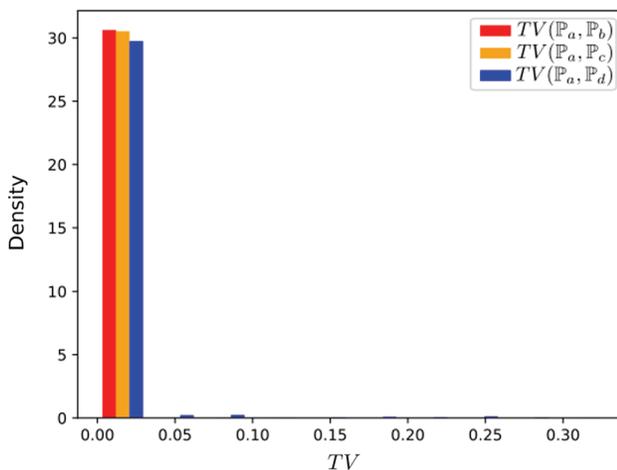
We report the results of our model with different decay rates in each scenario in Table 5. In addition to the *TSR* and *MHD* metrics, we report the collision rate (*CR*) between the agent and the NPC vehicles in the test trajectories. We find that *CR* is higher in Scenario 2 than in Scenario 1, which is

Table 5 Test results with dynamic obstacles from CARLA *Town05* map, including trajectory success rate (*TSR*), collision rate (*CR*), and modified Hausdorff distance (*MHD*) between the agent and the expert trajectories on the test set

Decay rate γ	Scenario 1 (no lane change)			Scenario 2 (lane change required)			t	Scenario 3 (mixed scenario)		
	<i>TSR</i> (%)	<i>CR</i> (%)	<i>MHD</i>	<i>TSR</i> (%)	<i>CR</i> (%)	<i>MHD</i>		<i>TSR</i> (%)	<i>CR</i> (%)	<i>MHD</i>
1.0	92	2	2.878	84	12	2.528		84	8	2.708
0.9	92	3	2.795	88	11	2.446		84	9	2.690
0.8	90	2	2.721	78	15	2.512		80	11	2.912
0.7	86	7	3.224	73	24	2.885		78	16	2.948

**Fig. 15** Noisy semantic segmentation observations: **a** original image, **b** each pixel is replaced with a random pixel within its local 5×5 pixel window **c** 2% of all pixels are randomly changed, **d** 20% of all pixels are randomly changed**Table 6** Maximum total variation distance between the original and perturbed semantic probabilities across all grid cells

$\max_{m_j} TV(\mathbb{P}_a, \mathbb{P}_b)$	$\max_{m_j} TV(\mathbb{P}_a, \mathbb{P}_c)$	$\max_{m_j} TV(\mathbb{P}_a, \mathbb{P}_d)$
0.002	0.185	0.511

**Fig. 16** Histogram of the total variation distance between the semantic probabilities from the original and the perturbed semantic images. Even with 20% incorrectly labeled pixels in the semantic segmentation images, most semantic probabilities are unaffected

9 Conclusion

This paper introduced an inverse reinforcement learning approach for inferring navigation costs from demonstrations with semantic category observations. Our cost model consists

of a probabilistic multi-class occupancy map and a deep fully convolutional cost encoder defined over the class likelihoods. The cost function parameters are optimized by computing the optimal cost-to-go of a deterministic shortest path problem, defining a Boltzmann control policy over the cost-to-go, and backpropagating the log-likelihood of the expert controls with a closed-form subgradient. Experiments in simulated minigrid environments and the CARLA autonomous driving simulator show that our approach outperforms methods that do not encode semantic information probabilistically over time. Our work offers a promising solution for learning complex behaviors from visual observations that generalize to new environments.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s10514-023-10118-4>.

Author Contributions Conceptualization: TW, NA; Methodology: TW, NA; Software, investigation, and analysis: TW, VD; Writing, review and editing: TW, VD, NA; Supervision: NA.

Funding We gratefully acknowledge support from NSF CRII IIS-1755568, ONR SAI N00014-18-1-2828 and NSF No. 2218063.

Declarations

Conflict of interest The authors have no relevant financial or non-financial interests to disclose.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as

long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendix A: Comparison between Boltzmann and maximum entropy policies

This appendix compares the MaxEnt expert model of Ziebart et al. (2008) to the expert model proposed in Sect. 3.2. The MaxEnt model has been widely studied in the context of reinforcement learning and inverse reinforcement learning (Haarnoja et al., 2017; Finn et al., 2016; Levine, 2018). On the other hand, while a Boltzmann policy is a well-known method for exploration in reinforcement learning, it has not been used to model expert or learner behavior in inverse reinforcement learning.

The work of Haarnoja et al. (2017) shows that both a Boltzmann policy and the MaxEnt policy are special cases of an energy-based policy:

$$\pi(\mathbf{u}_t | \mathbf{x}_t) \propto \exp(-E(\mathbf{x}_t, \mathbf{u}_t)) \quad (\text{A1})$$

with appropriate choices of the energy function E . We study the two policies in the discounted infinite-horizon setting, as this is the most widely used setting for the MaxEnt model. Extensions to first-exit and finite-horizon formulations are possible. Consider a Markov decision process with finite state space \mathcal{X} , finite control space \mathcal{U} , transition model $p(\mathbf{x}' | \mathbf{x}, \mathbf{u})$, stage cost $c(\mathbf{x}, \mathbf{u})$, and discount factor $\gamma \in (0, 1)$.

Proposition 3 (Haarnoja et al. (2017, Thm. 1)) *Define the maximum entropy Q-value as:*

$$Q_{ME}(\mathbf{x}_t, \mathbf{u}_t) := c(\mathbf{x}_t, \mathbf{u}_t) + \min_{\pi} \mathbb{E}_{\pi, p} \left[\sum_{k=t+1}^{\infty} \gamma^{k-t} (c(\mathbf{x}_k, \mathbf{u}_k) - \alpha \mathcal{H}(\pi(\cdot | \mathbf{x}_k))) \right], \quad (\text{A2})$$

where $\mathcal{H}(\pi(\cdot | \mathbf{x})) = -\sum_{\mathbf{u} \in \mathcal{U}} \pi(\mathbf{u} | \mathbf{x}) \log \pi(\mathbf{u} | \mathbf{x})$ is the Shannon entropy of $\pi(\cdot | \mathbf{x})$. Then, the maximum entropy policy satisfies:

$$\pi_{ME}(\mathbf{u}_t | \mathbf{x}_t) \propto \exp\left(-\frac{1}{\alpha} Q_{ME}(\mathbf{x}_t, \mathbf{u}_t)\right). \quad (\text{A3})$$

Similarly, define the usual Q-value as:

$$Q_{BM}(\mathbf{x}_t, \mathbf{u}_t) := c(\mathbf{x}_t, \mathbf{u}_t) + \min_{\pi} \mathbb{E}_{\pi, p} \left[\sum_{k=t+1}^{\infty} \gamma^{k-t} c(\mathbf{x}_k, \mathbf{u}_k) \right] \quad (\text{A4})$$

and the Boltzmann policy associated with it as:

$$\pi_{BM}(\mathbf{u}_t | \mathbf{x}_t) \propto \exp\left(-\frac{1}{\alpha} Q_{BM}(\mathbf{x}_t, \mathbf{u}_t)\right). \quad (\text{A5})$$

The value functions Q_{ME} and Q_{BM} can be seen as the fixed points of the following Bellman contraction operators:

$$\begin{aligned} \mathcal{T}_{ME}[Q](\mathbf{x}_t, \mathbf{u}_t) &:= c(\mathbf{x}_t, \mathbf{u}_t) \\ &- \gamma \alpha \mathbb{E}_p \left[\log \sum_{\mathbf{u}_{t+1} \in \mathcal{U}} \exp\left(-\frac{1}{\alpha} Q(\mathbf{x}_{t+1}, \mathbf{u}_{t+1})\right) \right] \end{aligned} \quad (\text{A6})$$

$$\mathcal{T}_{BM}[Q](\mathbf{x}_t, \mathbf{u}_t) := c(\mathbf{x}_t, \mathbf{u}_t) + \gamma \mathbb{E}_p \left[\min_{\mathbf{u}_{t+1} \in \mathcal{U}} Q(\mathbf{x}_{t+1}, \mathbf{u}_{t+1}) \right]. \quad (\text{A7})$$

In the latter, the Q values are bootstrapped with a “hard” min operator, while in the former they are bootstrapped with a “soft” min operator given by the log-sum-exponential operation. The form of the Bellman equations resembles the online SARSA update and offline Q-learning update in reinforcement learning. Consider temporal difference control with transitions $(\mathbf{x}, \mathbf{u}, c, \mathbf{x}', \mathbf{u}')$ using SARSA backups:

$$Q(\mathbf{x}, \mathbf{u}) \leftarrow Q(\mathbf{x}, \mathbf{u}) + \eta [c(\mathbf{x}, \mathbf{u}) + \gamma Q(\mathbf{x}', \mathbf{u}') - Q(\mathbf{x}, \mathbf{u})]$$

and Q-learning backups:

$$Q(\mathbf{x}, \mathbf{u}) \leftarrow Q(\mathbf{x}, \mathbf{u}) + \eta [c(\mathbf{x}, \mathbf{u}) + \gamma \min_{\mathbf{u}'} Q(\mathbf{x}', \mathbf{u}') - Q(\mathbf{x}, \mathbf{u})]$$

where η is a step-size parameter. If we additionally assume that the controls are sampled from the energy-based policy in (A1) defined by Q , the SARSA algorithm specifies the MaxEnt policy, while the Q-learning algorithm specifies the Boltzmann policy.

We show a visualization of the MaxEnt and Boltzmann policies, π_{ME} , π_{BM} , as well as their corresponding value functions Q_{ME} , Q_{BM} , in the infinite horizon setting with discount $\gamma = 0.95$ and $\alpha = 1$. The 4-connected grid environment in Fig. 17 has obstacles only along the outside border. The true cost is 0 to arrive at the goal (which is an absorbing state), 1 to any state (except the goal) inside the grid, infinity to any obstacle outside the border. Note that Q_{ME} and Q_{BM} are very different in absolute value. In fact, Q_{ME} is negative for all states due to the additional entropy term in (A2).

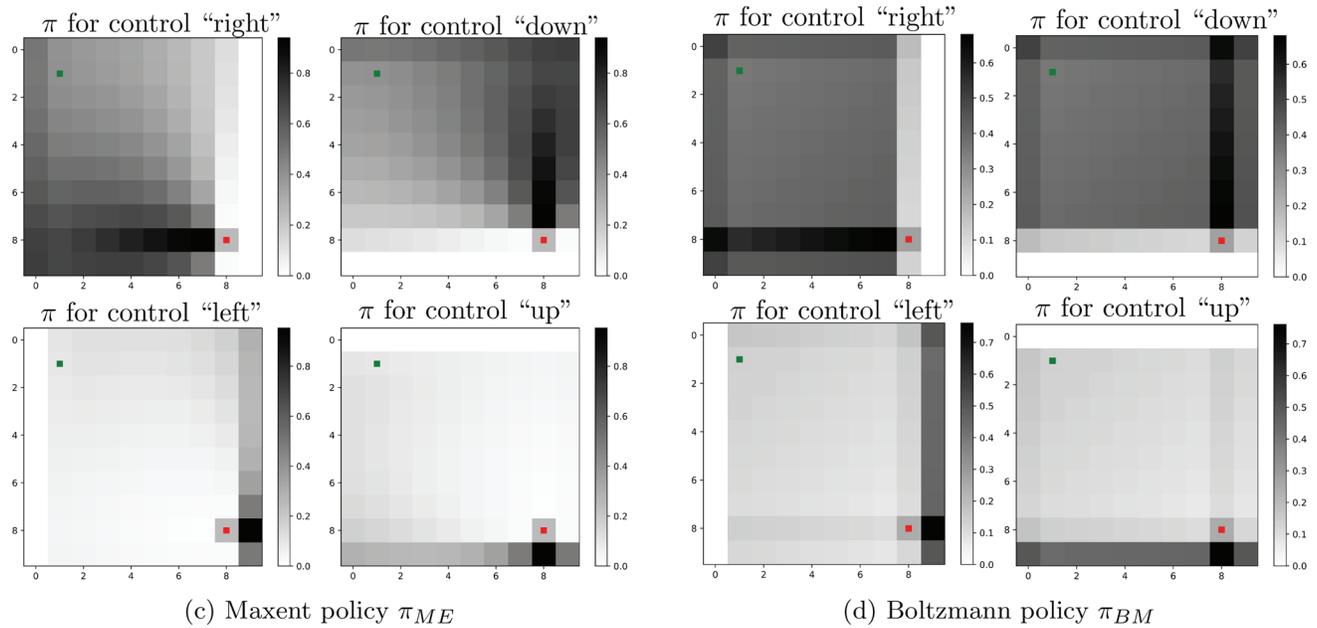
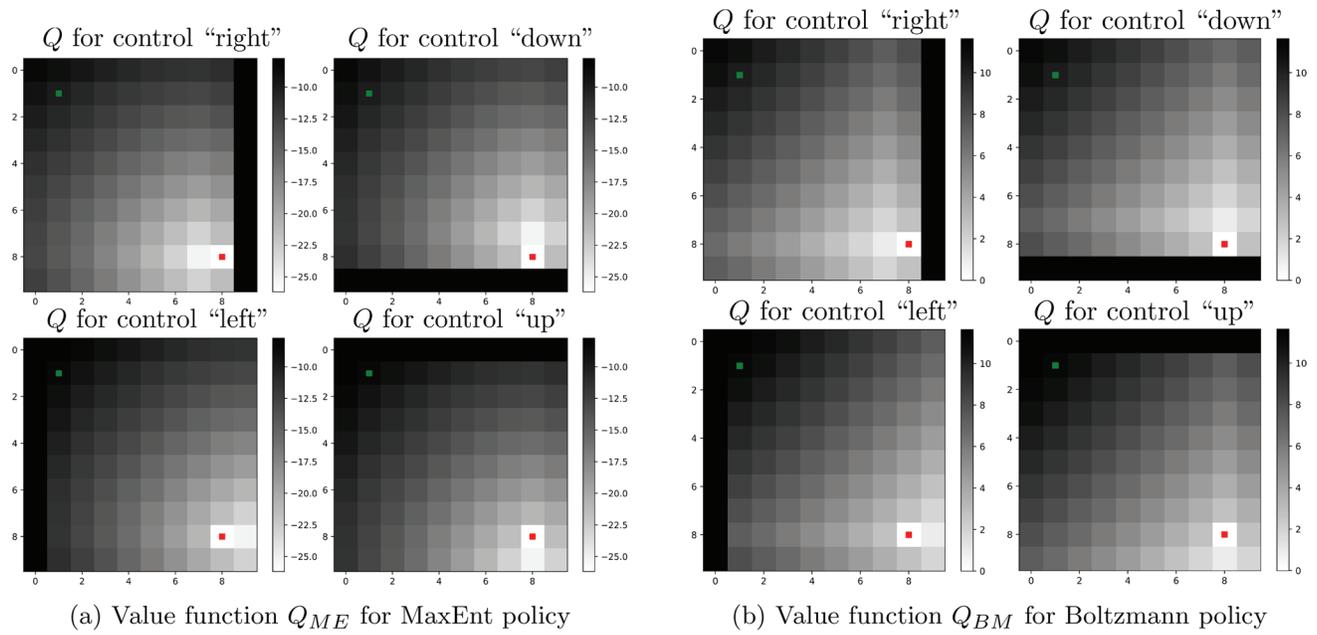


Fig. 17 Value functions corresponding to the MaxEnt and Boltzmann policies in infinite horizon setting with discount $\gamma = 0.95$. The environment only has obstacles around the outer boundary. The start state is marked in green and the goal in red. The controls are $\{right, down, left, up\}$ at each state with constant true cost of 0 to arrive at the goal (which is an absorbing state), 1 to any state except the goal

in the grid and infinity to any obstacle outside the border. Darker color indicates higher cost-to-go values to reach the goal (top two rows) or higher probability of choosing a control (bottom two rows). Although the absolute values of Q_{ME} and Q_{BM} are different, both have similar relative value differences across the controls, providing well-performing policies π_{ME} and π_{BM} (Color figure online)

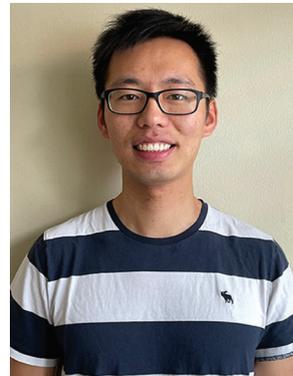
However, the relative value differences across the controls are similar and, thus, both policies π_{ME} and π_{BM} generate desirable paths from start to goal.

References

- Abbeel, P., & Ng, A. Y. (2004). Apprenticeship learning via inverse reinforcement learning. In *International conference on machine learning* (p. 1).
- Argall, B. D., Chernova, S., Veloso, M., & Browning, B. (2009). A survey of robot learning from demonstration. *Robotics and Autonomous Systems*, 57(5), 469–483.
- Atkeson, C. G., & Schaal, S. (1997). Robot learning from demonstration. In *International conference on machine learning* (Vol. 97, pp. 12–20).
- Badrinarayanan, V., Kendall, A., & Cipolla, R. (2017). SegNet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(12), 2481–2495.
- Bajcsy, A., Losey, D. P., O'malley, M. K., & Dragan, A. D. (2017). Learning robot objectives from physical human interaction. In *Conference on robot learning*.
- Baker, C. L., Tenenbaum, J. B., & Saxe, R. R. (2007). Goal inference as inverse planning. In *Annual meeting of the cognitive science society* (Vol. 29).
- Bertsekas, D. (1995). *Dynamic programming and optimal control*. Athena Scientific.
- Brown, D. S., Goo, W., & Niekum, S. (2020). Better-than-demonstrator imitation learning via automatically-ranked demonstrations. In *Conference on robot learning* (pp. 330–359).
- Chen, L., Paleja, R., & Gombolay, M. (2021). Learning from suboptimal demonstration via self-supervised reward regression. In *Conference on robot learning*.
- Chen, L.-C., Zhu, Y., Papandreou, G., Schroff, F., & Adam, H. (2018). Encoder-decoder with atrous separable convolution for semantic image segmentation. In *European conference on computer vision* (pp. 801–818).
- Chevalier-Boisvert, M., Willems, L., & Pal, S. (2018). Minimalistic grid-world environment for OpenAI Gym. <https://github.com/maximecb/gym-minigrid>. GitHub.
- Choy, C., Gwak, J., & Savarese, S. (2019). 4D spatio-temporal convnets: Minkowski convolutional neural networks. In *IEEE conference on computer vision and pattern recognition* (pp. 3075–3084).
- Cohen, T., & Welling, M. (2016). Group equivariant convolutional networks. In *International conference on machine learning* (pp. 2990–2999).
- Cortinhal, T., Tzelepis, G., & Aksoy, E. E. (2020). SalsaNext: Fast, uncertainty-aware semantic segmentation of lidar point clouds for autonomous driving. arXiv preprint [arXiv:2003.03653](https://arxiv.org/abs/2003.03653).
- Dohan, D., Matejek, B., & Funkhouser, T. (2015). Learning hierarchical semantic segmentations of lidar data. In *International conference on 3D vision* (pp. 273–281).
- Dosovitskiy, A., Ros, G., Codevilla, F., Lopez, A., & Koltun, V. (2017). CARLA: An open urban driving simulator. In *Proceedings of the 1st annual conference on robot learning* (pp. 1–16).
- Finn, C., Christiano, P., Abbeel, P., & Levine, S. (2016). A connection between generative adversarial networks, inverse reinforcement learning, and energy-based models. arXiv preprint [arXiv:1611.03852](https://arxiv.org/abs/1611.03852).
- Finn, C., Levine, S., & Abbeel, P. (2016). Guided cost learning: Deep inverse optimal control via policy optimization. In: *International conference on machine learning* (pp. 49–58).
- Fu, J., Luo, K., & Levine, S. (2018). Learning robust rewards with adversarial inverse reinforcement learning. In *International conference on learning representations*.
- Gan, L., Zhang, R., Grizzle, J. W., Eustice, R. M., & Ghaffari, M. (2020). Bayesian spatial kernel smoothing for scalable dense semantic mapping. *IEEE Robotics and Automation Letters*, 5(2), 790–797.
- Ghasemipour, S. K. S., Zemel, R., & Gu, S. (2020). A divergence minimization perspective on imitation learning methods. In *Conference on robot learning* (pp. 1259–1277).
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.
- Gupta, S., Davidson, J., Levine, S., Sukthankar, R., & Malik, J. (2017). Cognitive mapping and planning for visual navigation. In *Computer vision and pattern recognition (CVPR)*.
- Haarnoja, T., Tang, H., Abbeel, P., & Levine, S. (2017). Reinforcement learning with deep energy-based policies. In *International conference on machine learning* (pp. 1352–1361).
- Ho, J., & Ermon, S. (2016). Generative adversarial imitation learning. In *Advances in neural information processing systems* (pp. 4565–4573).
- Hornung, A., Wurm, K. M., Bennewitz, M., Stachniss, C., & Burgard, W. (2013). OctoMap: An efficient probabilistic 3D mapping framework based on octrees. *Autonomous Robots*, 34(3), 189–206.
- Ioffe, S., & Szegedy, C. (2015). Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning* (Vol. 37, pp. 448–456).
- Jain, A., Sharma, S., Joachims, T., & Saxena, A. (2015). Learning preferences for manipulation tasks from online coactive feedback. *The International Journal of Robotics Research*, 34, 1296–1313.
- Jeon, H. J., Milli, S., & Dragan, A. (2020). Rewardrational (implicit) choice: A unifying formalism for reward learning. *Advances in Neural Information Processing Systems*, 33, 4415–4426.
- Karaman, S., & Frazzoli, E. (2011). Sampling-based algorithms for optimal motion planning. *The International Journal of Robotics Research*, 30(7), 846–894.
- Ke, L., Choudhury, S., Barnes, M., Sun, W., Lee, G., & Srinivasa, S. (2020). Imitation learning as f-divergence minimization. In *International workshop on the algorithmic foundations of robotics*.
- Khan, A., Zhang, C., Atanasov, N., Karydis, K., Kumar, V., & Lee, D. D. (2018). Memory augmented control networks. In *International conference on learning representations*.
- Kingma, D. P., & Ba, J. (2014). ADAM: A method for stochastic optimization. In *International conference on learning representations*.
- LaValle, S. (1998). *Rapidly-exploring random trees: A new tool for path planning (TR 98-11)*. Department of Computer Science, Iowa State University.
- Levine, S. (2018). Reinforcement learning and control as probabilistic inference: Tutorial and review. arXiv preprint [arXiv:1805.00909](https://arxiv.org/abs/1805.00909).
- Levine, S., Popovic, Z., & Koltun, V. (2011). Nonlinear inverse reinforcement learning with Gaussian processes. In *Advances in neural information processing systems* (pp. 19–27).
- Levine, S., Finn, C., Darrell, T., & Abbeel, P. (2016). End-to-end training of deep visuomotor policies. *The Journal of Machine Learning Research*, 17(1), 1334–1373.
- Likhachev, M., Gordon, G., & Thrun, S. (2004). ARA*: Anytime A* with provable bounds on sub-optimality. In *Advances in neural information processing systems* (pp. 767–774).
- Lu, C., van de Molengraft, M. J. G., & Dubbelman, G. (2019). Monocular semantic occupancy grid mapping with convolutional variational encoder-decoder networks. *IEEE Robotics and Automation Letters*, 4(2), 445–452.
- Milioto, A., Vizzo, I., Behley, J., & Stachniss, C. (2019). RangeNet++: Fast and accurate lidar semantic segmentation. In *IEEE/RSJ international conference on intelligent robots and systems (IROS)* (pp. 4213–4220).

- Neu, G., & Szepesvári, C. (2007). Apprenticeship learning using inverse reinforcement learning and gradient methods. In *Conference on uncertainty in artificial intelligence* (pp. 295–302).
- Ng, A. Y., & Russell, S. (2000). Algorithms for inverse reinforcement learning. In *International conference on machine learning* (pp. 663–670).
- Oleynikova, H., Taylor, Z., Fehr, M., Siegart, R., & Nieto, J. (2017). Voxblox: Incremental 3D Euclidean signed distance fields for onboard MAV planning. In *IEEE/RSJ international conference on intelligent robots and systems (IROS)* (pp. 1366–1373).
- Pan, Y., Cheng, C.-A., Saigol, K., Lee, X., Yan, X., Theodorou, E. A., & Boots, B. (2020). Imitation learning for agile autonomous driving. *The International Journal of Robotics Research*, 39(2–3), 286–302.
- Papandreou, G., Chen, L.-C., Murphy, K. P., & Yuille, A. L. (2015). Weakly- and semi-supervised learning of a deep convolutional network for semantic image segmentation. In *IEEE international conference on computer vision* (pp. 1742–1750).
- Pastor, P., Hoffmann, H., Asfour, T., & Schaal, S. (2009). Learning and generalization of motor skills by learning from demonstration. In *IEEE international conference on robotics and automation* (pp. 763–768).
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., & Desmaison, A. (2019). Pytorch: An imperative style, high-performance deep learning library. In *Advances in neural information processing systems* (pp. 8026–8037).
- Rajeswaran, A., Kumar, V., Gupta, A., Vezzani, G., Schulman, J., Todorov, E., & Levine, S. (2018). Learning complex dexterous manipulation with deep reinforcement learning and demonstrations. In *Proceedings of robotics: Science and systems (RSS)*.
- Ramachandran, D., & Amir, E. (2007). Bayesian inverse reinforcement learning. In *International joint conferences on artificial intelligence organization* (Vol. 7, pp. 2586–2591).
- Ratliff, N. D., Bagnell, J. A., & Zinkevich, M. A. (2006). Maximum margin planning. In *International conference on machine learning* (pp. 729–736).
- Ross, S., Gordon, G., & Bagnell, D. (2011). A reduction of imitation learning and structured prediction to no-regret online learning. In *International conference on artificial intelligence and statistics* (pp. 627–635).
- Sengupta, S., Sturges, P., Ladický, L., & Torr, P. H. (2012). Automatic dense visual semantic mapping from street-level imagery. In *IEEE/RSJ international conference on intelligent robots and systems* (pp. 857–862).
- Shor, N. Z. (2012). *Minimization methods for nondifferentiable functions* (Vol. 3). Springer.
- Song, Y. (2019). Inverse reinforcement learning for autonomous ground navigation using aerial and satellite observation data (unpublished master's thesis). Carnegie Mellon University.
- Sun, L., Yan, Z., Zaganidis, A., Zhao, C., & Duckett, T. (2018). Recurrent-OctoMap: Learning state-based map refinement for long-term semantic mapping with 3D-Lidar data. *IEEE Robotics and Automation Letters*, 3(4), 3749–3756.
- Tamar, A., Wu, Y., Thomas, G., Levine, S., & Abbeel, P. (2016). Value iteration networks. In *Advances in neural information processing systems* (pp. 2154–2162).
- Tew, P. A. (2016). An investigation of sparse tensor formats for tensor libraries (unpublished doctoral dissertation). Massachusetts Institute of Technology.
- Thrun, S., Burgard, W., & Fox, D. (2005). *Probabilistic robotics*. MIT Press.
- Wang, T., Dhiman, V., & Atanasov, N. (2020a). Learning navigation costs from demonstration in partially observable environments. In *IEEE international conference on robotics and automation*.
- Wang, T., Dhiman, V., & Atanasov, N. (2020b). Learning navigation costs from demonstrations with semantic observations. In *Conference on learning for dynamics and control*.
- Wu, B., Wan, A., Yue, X., & Keutzer, K. (2018). SqueezeSeg: Convolutional neural nets with recurrent CRF for real-time road-object segmentation from 3D lidar point cloud. In *International conference on robotics and automation* (pp. 1887–1893).
- Wulfmeier, M., Wang, D. Z., & Posner, I. (2016). Watch this: Scalable cost-function learning for path planning in urban environments. In *IEEE/RSJ international conference on intelligent robots and systems (IROS)* (pp. 2089–2095).
- Zhu, Y., Wang, Z., Merel, J., Rusu, A., Erez, T., Cabi, S., Tunyasuvunakool, S., Kramar, J., Hadsell, R., de Freitas, N., & Heess, N. (2018). Reinforcement and imitation learning for diverse visuomotor skills. In *Robotics: Science and systems*.
- Ziebart, B. D., Maas, A., Bagnell, J., & Dey, A. K. (2008). Maximum entropy inverse reinforcement learning. In *AAAI conference on artificial intelligence* (pp. 1433–1438).

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Tianyu Wang is a PhD student in Electrical and Computer Engineering at the University of California San Diego, CA, USA. He received a B.S. degree in Physics from Haverford College, PA in 2016 and an M.S. degree in Electrical and Computer Engineering from the University of California San Diego, La Jolla, CA in 2018. His research interests include reinforcement learning and imitation learning with applications to robotics and autonomous driving.



Vikas Dhiman is an Assistant Professor in the ECE department at the University of Maine. His works lie in the localization, mapping and control algorithms for applications in robotics. He was a Postdoctoral Researcher at the University of California, San Diego (2019-21). He graduated in Elec. Engg (2008) from Indian Institute of Technology, Roorkee, earned his MS (2014) from University at Buffalo, and received his PhD (2019) from the University of Michigan, Ann Arbor.



Nikolay Atanasov is an Assistant Professor of Electrical and Computer Engineering at the University of California San Diego, La Jolla, CA, USA. He obtained a B.S. degree in Electrical Engineering from Trinity College, Hartford, CT, USA in 2008, and M.S. and Ph.D. degrees in Electrical and Systems Engineering from University of Pennsylvania, Philadelphia, PA, USA in 2012 and 2015, respectively. Dr. Atanasov's research focuses on robotics, control theory, and machine learning

with emphasis on active perception problems for autonomous mobile robots. He works on probabilistic models that unify geometric and semantic information in simultaneous localization and mapping (SLAM) and on optimal control and reinforcement learning algorithms for minimizing probabilistic model uncertainty. Dr. Atanasov's work has been recognized by the Joseph and Rosaline Wolf award for the best Ph.D. dissertation in Electrical and Systems Engineering at the University of Pennsylvania in 2015, the Best Conference Paper Award at the IEEE International Conference on Robotics and Automation (ICRA) in 2017, the NSF CAREER Award in 2021, and the IEEE RAS Early Academic Career Award in Robotics and Automation in 2023.