Learning Occupancy Priors of Human Motion From Semantic Maps of Urban Environments

Andrey Rudenko[®], Luigi Palmieri[®], Johannes Doellinger[®], Achim J. Lilienthal[®], and Kai O. Arras

Abstract—Understanding and anticipating human activity is an important capability for intelligent systems in mobile robotics, autonomous driving, and video surveillance. While learning from demonstrations with on-site collected trajectory data is a powerful approach to discover recurrent motion patterns, generalization to new environments, where sufficient motion data are not readily available, remains a challenge. In many cases, however, semantic information about the environment is a highly informative cue for the prediction of pedestrian motion or the estimation of collision risks. In this work, we infer occupancy priors of human motion using only semantic environment information as input. To this end, we apply and discuss a traditional Inverse Optimal Control approach, and propose a novel approach based on Convolutional Neural Networks (CNN) to predict future occupancy maps. Our CNN method produces flexible context-aware occupancy estimations for semantically uniform map regions and generalizes well already with small amounts of training data. Evaluated on synthetic and real-world data, it shows superior results compared to several baselines, marking a qualitative step-up in semantic environment assessment.

Index Terms—Deep learning for visual perception, human detection and tracking, human motion analysis, human motion prediction, semantic scene understanding.

I. INTRODUCTION

U NDERSTANDING and predicting human motion is an increasingly popular subject of research with the goal of improving the safety and efficiency of autonomous systems in spaces shared with people. Application areas include mobile service robots, intelligent vehicles, collaborative production assistants, video surveillance, or urban city planning. Human motion in these scenarios is influenced by many factors, including other agents in the scene and the environment itself, which can be represented by a topometric map and semantic information [1].

Manuscript received July 11, 2020; accepted February 4, 2021. Date of publication February 24, 2021; date of current version March 22, 2021. This letter was recommended for publication by Associate Editor M. Hanheide and Editor T. Asfour upon evaluation of the reviewers' comments. This work was supported by the European Union's Horizon 2020 Research and Innovation Programme under Grant Agreement 732737 (ILIAD). (*Corresponding author: Andrey Rudenko.*)

Andrey Rudenko is with the Bosch Corporate Research, Renningen, Germany, and also with the Mobile Robotics and Olfaction Lab, Örebro University, Orebro 70182, Sweden (e-mail: andrey.rudenko@de.bosch.com).

Luigi Palmieri and Kai O. Arras are with the Bosch Corporate Research, Renningen 71272, Germany (e-mail: palmieri@informatik.uni-freiburg.de; kaioliver.arras@de.bosch.com).

Johannes Doellinger is with the Bosch Center for Artificial Intelligence, Renningen 71272, Germany (e-mail: johannes.doellinger@de.bosch.com).

Achim J. Lilienthal is with the Mobile Robotics and Olfaction Lab, Örebro University, Orebro 70182, Sweden (e-mail: achim.lilienthal@oru.se).

Digital Object Identifier 10.1109/LRA.2021.3062010



Fig. 1. Predicting occupancy priors in semantically-rich urban environments. **Top left:** an urban scene from the Stanford Drone dataset. **Top right:** semantic map of the environment. **Bottom left:** CNN-predicted occupancy distribution priors of walking people in the environment, encoded with a heatmap: warmer colors correspond to states with higher probability of observing pedestrians. **Bottom right:** ground truth occupancy distribution.

While indoor human navigation is often motivated by avoiding collisions with static and dynamic obstacles, surface semantics have a strong impact in outdoor (e.g. urban) environments. For instance, pedestrians walk most of the time on sidewalks, sometimes on streets, unpaved areas and greenspaces, and very rarely over obstacles. Modeling the influence of semantics is a challenging task, typically approached with data-driven methods using human trajectory data in a given environment [2]–[4] without knowing the goal of the target agent. Powerful in scenes known beforehand, such approaches may suffer from poor generalization to never-seen or changing environments where no data is available.

In this letter, we research the possibility of inferring occupancy priors of walking people in previously unseen places with limited input, namely using only the semantic map of the area. A prior occupancy distribution is intuitively interpretable and beneficial for a large variety of applications, such as improved goal estimation [5], [6] (for instance in Fig. 1, not all walking directions are likely to be the goal of a person) and crossing intention recognition in autonomous driving tasks [7], [8], where possible "illegal crosswalks" could be easily detected. The usage of semantic map-based occupancy priors may further improve the accuracy of map-based motion prediction approaches [9],

2377-3766 © 2021 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.



Fig. 2. Structure of the *semapp* network. Input semantic tensor of size |<height>x<width>x<features>| is downsampled twice before passing through a bottleneck and getting upsampled to the original resolution again. Black arrows indicate skip connections. The kernel sizes and strides are denoted as |<kernel width>x<kernel height>s<stride>|.

which often assume constant priors for each semantic class. Such occupancy estimation can guide a cleaning robot towards more heavily used areas, or a service robot in search of people to assist.

Traditionally, Inverse Reinforcement Learning (IRL) has been used to learn semantic preferences of walking people in urban and semantically-rich environments [10]-[12]. It is indeed possible to use the learned preferences to simulate trajectories and infer the prior occupancy distribution in a new environment. In this work we review the IRL methodology, applied to the occupancy prior distribution inference, and discuss its limitations. In order to address them, we introduce a novel extension to the recent method by Doellinger et al. [13], which uses a Convolutional Neural Network (CNN) to predict average occupancy maps indoors, with semantic map input for the urban scenes. We train our method on scenes from the Stanford Drone Dataset [14], as well as on simulated environments. In comparison to several baselines, our CNN method predicts much more accurate prior occupancy priors in terms of KL-divergence to the ground truth distributions. Crucially, it makes a qualitative improvement of estimating flexible occupancy priors for semantically-uniform areas by considering local context and interconnections between different semantic regions.

In summary, we make the following contributions:

- We analyze and discuss state-of-the-art methodology for inferring occupancy prior distribution in semantically-rich urban environments.
- Addressing the limitations of the prior art, we propose a novel method based on Convolutional Neural Networks (displayed in Fig. 1 and 2).
- We execute a thorough comparison of the discussed methods with several baselines, and show qualitative and quantitative improvement of the KL-divergence scores when using our CNN method.

The letter is structured as follows: in Section II we review the related work, in Section III we detail the proposed solutions and in Section IV we describe the training and evaluation. Results are presented in Section V, and a discussion in Section VI concludes the letter.

II. RELATED WORK

The ability to understand a human environment and its affordances is useful in a number of tasks where intelligent autonomous systems need to reason on observed events, anticipate future events, evaluate risks and act in a dynamic world. Examples include person and group tracking [15]–[17], in particular over a camera network with non-overlapping fields of view, human-aware motion planning [18]-[20], motion behavior learning [21], human motion prediction [22], [23], human-robot interaction [24], video surveillance [25] or collision risk assessment [26]. Apart from basic geometric properties of the workspace, its semantics have a large impact on human motion in these tasks. Modeling this impact is challenging, therefore a popular approach is learning the motion patterns directly from data without explicitly specifying semantic features [2]–[4], [27]. However, many of those methods either need additional training input in new environments or experience transfer issues. To the best of our knowledge, only a few methods explicitly highlight the performance in new environments outside the training scenario [10], [28]-[30]. As we will later describe, our approach explicitly uses only semantic maps as input and there will be no need to adapt the learned models to new environments, described with the same semantic classes.

Modeling the effect of surface classes on human motion was mainly used in reactive approaches such as [31] and planningbased approaches [9], [10]. These methods leverage semantic segmentation tools for understanding and detecting the semantics in the environment.

Several approaches exist to segment available semantics [32]– [34] and to build semantic maps of the environment [35], [36] – a prerequisite to the methods presented in this letter. We build on those methods to predict areas frequently used by pedestrians based on the semantic class of the surface. Several Inverse Reinforcement Learning (IRL, or Inverse Optimal Control, IOC) approaches make use of semantic maps for predicting future human motion [9], [10], [30]. In particular, they use the semantic maps for encoding the features of the reward function. However, these IRL approaches are limited to one weight per feature and thus do not generalize well to new environments or heteregenous datasets with different geometries [37], [38]. In this work, we review an adaption of the Maximum Entropy IRL algorithm [39] to the task of occupancy priors estimation, and compare our CNN-based approach to it.

III. OCCUPANCY PRIORS ESTIMATION IN URBAN ENVIRONMENTS

In this work, we study the problem of estimating occupancy priors of walking humans in semantically-rich urban environments. The problem is formulated as follows: given a grid-map of the environment \mathcal{M} with associated feature responses $f(s) = [f_1(s), \ldots, f_K(s)], \sum_{k=1}^K f_k(s) = 1$ for each state $s \in \mathcal{M}$ over the set of K semantic classes, we seek to estimate the probability p(s) of a walking human being observed in this state.

If we assume having access to a large set of trajectories $\mathcal{T}_{\mathcal{M}}$ in \mathcal{M} , the problem of estimating p(s) can be solved by counting visitation frequencies in each state:

$$p(s) = \frac{D(s)}{\sum_{s' \in \mathcal{M}} D(s')} \Big|_{\mathcal{T}_{\mathcal{M}}},\tag{1}$$

where D(s) is visitation count of state s over all trajectories $\mathcal{T}_{\mathcal{M}}$. In this letter, we estimate this distribution in environments where no trajectory data are available. One natural way to overcome the lack of trajectories is to simulate them, in particular using learned human walking preferences. To this end, in Section III-A we first review an Inverse Optimal Control (IOC) method [10]

Algorithm 1: Backward Pass.		
1:	function BackwardPass($\mathcal{T}_{\mathcal{M}}^{i}, \boldsymbol{\theta}$)	
2:	$V(s) \leftarrow -\infty$	
3:	for $n = N, \ldots, 1$ do	
4:	$V^{(n)}(s_g) \leftarrow 0$	
5:	$Q^{(n)}(s,a) \leftarrow \mathcal{R}(s,\boldsymbol{\theta}) + E_{P^{s'}_{s,a}}[V^{(n)}(s')]$	
6:	$V^{(n-1)}(s) \leftarrow \operatorname{softmax}_a Q^{(n)}(s,a)$	
7:	end for	
8:	$\pi(a s) \leftarrow e^{\alpha(Q(s,a) - V(s))}$	
9:	return π	

for predicting motion trajectories in semantic environments, and discuss its applicability and limitations. Then, in Section III-B, we propose a novel approach based on Convolutional Neural Networks, which is an extension of the occupancy prior estimation method by Doellinger *et al.* [13]. For this task, we assume a semantic map of the environment $f(\mathcal{M})$, or a method to extract it, to be available. Without loss of generality and for the sake of visual clarity, the states in this letter are represented with one-hot vectors, i.e. $\forall s \exists ks.t.f_k(s) = 1$ and $\forall j \neq k : f_j(s) = 0$.

A. Inverse Optimal Control on Multiple Maps (IOCMM)

Both Reinforcement Learning (RL) and Inverse RL or Inverse Optimal Control (IOC) frameworks deal with modeling optimal behavior of an agent, operating in a stochastic world S and collecting rewards \mathcal{R} on the way to their goal state $s_g \in S$. An agent's behavior is encoded in a policy $\pi(a|s)$, which maps the state $s \in S$ to a distribution over actions $a \in A$. When the reward function is not known beforehand, which is the case in many real-world applications, one possibility is to learn it from a set of observations \mathcal{T} with an IOC method. In this case, the reward function is parametrized by a set of parameters θ .

Modeling the behavior of an agent navigating in the environment, which is described with a set of features f(s) for each state, suits the problem of recovering occupancy priors from semantic map inputs well. Prior art, however, has not dealt with abstract quantities, such as occupancy expectations, focusing rather on the policy of an individual agent [40] or multiple agents jointly [41]. In our work we adapt the IOC framework to this task. As our IOC implementation is based on [10], we give a short summary of their approach in this section.

MDP-based Maximum Entropy Inverse Reinforcement Learning (MaxEnt IRL) [39] assumes that the observed motion of agents is generated by a stochastic motion policy, and seeks to estimate this policy with maximum likelihood to the available demonstrations. The reward an agent gets in state s is linear with respect to the feature responses in that state: $\mathcal{R}(s, \theta) = r_0 + \theta^T f(s)$, where $r_0 > 0$ is the base reward of a transition and θ is a set of weights or *costs* of the semantic classes: $\sum_{k=1}^{K} \theta_k = 1, \theta_k \in [0, 1]$. Given \mathcal{R} , the distribution over the sequence of states s is defined as

$$p(\boldsymbol{s},\boldsymbol{\theta}) = \frac{\prod_{t} e^{\mathcal{R}(s_{t},\boldsymbol{\theta})}}{Z(\boldsymbol{\theta})} = \frac{e^{\sum_{t} r_{0} + \boldsymbol{\theta}^{T} \boldsymbol{f}(s_{t})}}{Z(\boldsymbol{\theta})}.$$
 (2)

Finding the optimal θ^* vector is equivalent to maximizing the entropy of $p(s, \theta)$ in Eq. 2 while matching the semantic

Algo	Algorithm 2: Forward Pass.		
1:	function ForwardPass($\mathcal{T}_{\mathcal{M}}^{i}, \pi$)		
2:	$D \leftarrow 0$		
3:	for $n = 1, \ldots, N$ do		
4:	$s \leftarrow s_0$		
5:	while $s \neq s_g$ do		
6:	$D(s) \leftarrow D(s) + 1$		
7:	$s' \leftarrow \pi(a s)$		
8:	$s \leftarrow s'$		
9:	$\hat{\mathbf{f}}_{ heta} \leftarrow \sum_{s} \boldsymbol{f}(s) D(s)$		
10:	return $\hat{\mathbf{f}}_{\theta}$		

Algorithm 3: IOCMM.

1:	$\boldsymbol{\theta} \leftarrow 1/K$
2:	repeat
3:	$\hat{\mathbf{f}}_{\theta} \leftarrow 0, \bar{\mathbf{f}} \leftarrow 0$
4:	Batch B_m maps
5:	for $m = 1, \ldots, B_m$ do
6:	$oldsymbol{\mathcal{T}} \leftarrow ext{Batch} \ B_t$ trajectories from \mathcal{M}_m
7:	$ar{\mathbf{f}} \leftarrow ar{\mathbf{f}} + rac{1}{ \mathcal{T} }\sum_i^{ \mathcal{T} } oldsymbol{f}(\mathcal{T}^i)$
8:	$\mathcal{R}(s, \boldsymbol{\theta}) \leftarrow r_0 + \boldsymbol{\theta}^T \boldsymbol{f}(s)$
9:	for $i = 1, \ldots, B_t$ do
10:	$\pi \leftarrow \texttt{BackwardPass}(s_g)$
11:	$\hat{\mathbf{f}}_{ heta} \leftarrow \hat{\mathbf{f}}_{ heta} + extsf{ForwardPass}(s_0,\pi)$
12:	$ar{\mathbf{f}} \leftarrow \texttt{normalize}(ar{\mathbf{f}})$
13:	$\hat{\mathbf{f}}_{ heta} \gets \texttt{normalize}(\hat{\mathbf{f}}_{ heta})$
14:	$ abla \mathcal{L}_{ heta} \leftarrow ar{\mathbf{f}} - \hat{\mathbf{f}}_{ heta}$
15:	$oldsymbol{ heta} \leftarrow oldsymbol{ heta} e^{\lambda abla \mathcal{L}_{ heta}}$
16:	until $ \nabla \mathcal{L}_{\theta} < \epsilon$

class feature counts of the training trajectories. An iterative procedure based on the exponentiated gradient descent of the log-likelihood $\mathcal{L} \triangleq \log p(\boldsymbol{s}|\boldsymbol{\theta})$ is described by Kitani *et al.* in [10]. The gradient $\nabla \mathcal{L}_{\theta}$ is computed as the difference between the *empirical* mean feature count $\mathbf{\bar{f}} = \frac{1}{|\mathcal{T}|} \sum_{i}^{|\mathcal{T}|} f(\mathcal{T}^{i})$, i.e. the average features accumulated over the \mathcal{T} training trajectories in the map \mathcal{M} , and the *expected* mean feature count $\mathbf{\hat{f}}_{\theta}$, the average features accumulated by trajectories generated by the current parameters $\boldsymbol{\theta}: \nabla \mathcal{L}_{\theta} = \mathbf{\bar{f}} - \mathbf{\hat{f}}_{\theta}$. The weight vector is then updated as

$$\boldsymbol{\theta} \leftarrow \boldsymbol{\theta} e^{\lambda \nabla \mathcal{L}_{\boldsymbol{\theta}}},\tag{3}$$

where λ is the learning rate, and the expected mean feature count $\hat{\mathbf{f}}_{\theta}$ is computed using an iterative algorithm described below.

The algorithm iterates backward and forward passes, detailed in Algorithm 1 and 2 respectively. The *backward pass* uses the current θ vector to compute the value function V(s) for each state s in \mathcal{M} given the goal state s_g – the final state of the trajectory $\mathcal{T}_{\mathcal{M}}^i \in \mathcal{T}_{\mathcal{M}}$. A stochastic motion policy $\pi_{\theta}(a|s)$ to reach s_g in \mathcal{M} under $\mathcal{R}(s, \theta)$ is then computed and used in the *forward pass* to simulate several trajectories from s_0 to s_g , where s_0 is the initial state of $\mathcal{T}_{\mathcal{M}}^i$. The expected mean feature count is computed as a weighted sum of feature counts $\hat{\mathbf{f}}_{\theta} =$ $\sum_{s} f(s)D(s)$ in the simulated trajectories, and the backwardforward iteration is repeated for a batch of trajectories in $\mathcal{T}_{\mathcal{M}}$. The θ vector is updated as in Eq. 3 using the cumulative $\hat{\mathbf{f}}_{\theta}$ for the trajectories in the batch, and the algorithm is iterated until the gradient $\nabla \mathcal{L}_{\theta}$ reaches zero. In order to learn from multiple maps, including those where only a subset of K features is present, we run the backward and forward passes for a batch of trajectories *in a batch of maps*, accumulating the visitation counts D(s)across several maps. The resulting Inverse Optimal Control on Multiple Maps (IOCMM) method is detailed in Algorithm 3.

Having obtained the optimal θ^* weights, it is possible to compute the reward $\mathcal{R}(s, \theta^*)$ and simulate trajectories in any environment which is described by a subset of K semantic features. By simulating semantic-aware trajectories, an average visitation count for each state, normalized across all states in \mathcal{M} , yields the occupancy probability p(s), as in Eq. 1. Apart from the θ^* vector, this simulation depends on the distributions from which the initial and goal states s_0 and s_a (hereinafter denoted $s_{0,q}$) are drawn: since the algorithm is inherently unaware of the semantics behind classes, omitting this step may result in $s_{0,q}$ generation inside of obstacles or other high-cost areas. To counteract this issue, we consider two strategies: (1) directly learn probabilities to sample the start or goal position in a state s, conditioned on the semantic class f(s) of the state: $p(s_{0,g}|\boldsymbol{f}(s))$, and (2) generate the $s_{0,g}$ only from low-cost regions with the softmax function over the estimated cost of the state: $p(s_{0,g}) \sim e^{-\mathcal{R}(s,\theta^*)/\tau}$. Furthermore, to generate long trajectories spanning across the map, both distributions (1) and (2) are scaled linearly proportional to the distance between the $s_{0,q}$ and the center of the map.

1) Analysis and Discussion: While delivering adequate results in our experiments, as we show in Section IV, the IOCMM approach to occupancy priors estimation has an inherent drawback. With rigid costs of a semantic class k, defined by the corresponding θ_k weight, the IOCMM method cannot produce flexible estimations for a spatial region given its position in a wider topological structure. For instance, if we assume that the grass surface is walkable, then it will have a low cost and predicted people would largely ignore paved paths in a park. However, this behavior is probably not confirmed in the training data, which will increase the costs of the grass regions, potentially making them not traversable in some cases where such behavior is expected. Controlled by one θ_k parameter, the cost of the semantic class stays constant over the entire map. Similarly, learning to step on the road surface in places where this behavior is unavoidable will inevitably lead to decreasing the costs of the road surface everywhere in the map. Learning such flexible behavior requires reasoning over the local context and interconnections between different semantic attributes and the surface. To this end, we propose our Convolutional Neural Network-based approach "Semantic Map-Aware Pedestrian Prediction" (semapp), described in the next section.

B. Semantic Map-Aware Pedestrian Prediction (semapp)

Convolutional Neural Networks (CNNs) have shown great successes for operations on map data, such as semantic segmentation [32] or value function estimation for deep reinforcement learning [42]. For our task of predicting occupancy distributions of walking humans in semantic environments, we need a method to map the feature responses $f(\mathcal{M})$ to probabilities p(s). To this end, we extend the network to predict occupancy values in semantics-free geometric environments [13]. This network, based on the FC-DenseNet architecture [43], has reasonably few parameters which helps to avoid overfitting when training on limited amounts of data. Experiments with different architectures have been made in [13] but the authors have found their results to be very robust to such changes. We thus decided to perform no further optimization on the network architecture. The method in [13] is referred to as Map-Aware Pedestrian Prediction (mapp), therefore we call our extension "Semantic mapp," or semapp. Extending the architecture from [13] to semantic inputs by changing the input from one binary input channel to one channel for each semantic class allows the network to differentiate between pedestrians walking on grass, sidewalks and streets additional to avoiding obstacles. The architecture is outlined in Fig. 2.

The network directly outputs the map-sized tensor with the occupancy distribution, so, unlike IOCMM, semapp requires no trajectory simulation for inference. Consequently, for training we convert the trajectories $\mathcal{T}_{\mathcal{M}}$ in each map \mathcal{M} into the occupancy distribution using Eq. 1. This conversion itself is not without meaning: trajectories, as compared to the processed occupancy distribution, contain additional temporal information. However, this information is not necessarily relevant for the task at hand - in fact, we are deliberately discarding the temporal aspect of human motion, inferring instead the generalized prior of observing a person in any state of the environment. Using directly the distribution emphasises the relationship between the topology of the environment and the desired occupancy priors. Furthermore, it relaxes the requirements to the data itself: detections are sufficient, and there is no need for continuous tracks.

Since the network operates on map crops of fixed size, we decompose a larger input image into a number of random crops of appropriate size and then rebuilt the final distribution p(s) for state s as an average of predicted occupancy values of s in all crops which include that state (see Fig. 3). In this case random crops, as compared to regular grids, remove the aliasing issues from combining adjacent crops.

IV. EXPERIMENTS

In this section, we give an overview of the training data (Section IV-A) and the experiments' design, as well as details on the training and baseline implementation (Section IV-B).

A. Datasets

We evaluate all methods on two datasets of human trajectories in semantically-rich environments: the Stanford Drone Dataset [14] and a set of simulated maps. Both datasets are summarized in Table I.

To prove the concept of learning occupancy distributions from semantic maps, we created the "U4" dataset which includes 80 hand-crafted maps of simulated <u>U</u>rban environments with <u>four</u> semantic classes (sidewalk, grass, road and obstacle) and manually marked \sim 30 trajectories in each map. In this dataset, we pay particular attention to "illegal crosswalk" detection, i.e. such



Fig. 3. (a) CNN training on crops from larger maps in the Stanford Drone Dataset, (b) corresponding crops from the ground truth distribution. (c) Inference for a large map by averaging the inference for random crops in the test image, (d) full predicted occupancy distribution from 500 crops.

TABLE I DATASETS SUMMARY

Dataset	U4	Stanford Drone
Number of maps	80	25
Map size in pixels	32×32	$\sim 146 \times 152$
Resolution	simulation	0.4 m
Number of trajectories per map	~ 30	~ 132
Number of semantic classes in the dataset	4	9



Fig. 4. Training examples from the U4 dataset: each pair shows the semantic map on the left and the ground truth occupancy distribution on the right. Semantic classes include vehicle road in **dark gray**, pedestrian areas and sidewalks in **light gray**, unpaved areas and grass in **green** and obstacles in **red**.

scenes where global topology of the environment encourages people to step onto the driveway and cross it. Additionally, as people often tend to cut sharp corners by walking over grass, such behavior is also included in this dataset. Several scenes from U4 are shown in Fig. 4.

The *Stanford Drone Dataset* (SDD) [14] was recorded on the Stanford University grounds, which include a wide variety of environments and semantic classes, e.g. shared roads for cyclists and vehicles, pedestrian areas, college buildings, vegetation and parking lots. The dataset includes 51 top-down scenes with bounding boxes for various agents, from which we extracted trajectories of people, approximating the position by the center of the bounding box. We chose 25 scenes sufficiently covered by trajectories and scaled the maps to the constant physical



Fig. 5. Semantic segmentation, obtained by our proof-of-concept pipeline in the Stanford Drone dataset. Left: input images. Middle: ground truth labels. Right: predicted semantic map.

resolution of 0.4 m per cell. We manually segmented each scene into nine semantic classes: pedestrian area, vehicle road, bicycle road, grass, tree foliage, bulging, entrance, obstacle and parking. Some example scenes from SDD are shown in Fig. 1, 3 and 5.

B. Training and Evaluation

To our knowledge, there exist quite some works on human motion prediction [1] but none of the existing methods predicts prior occupancy distribution of walking people in urban environments only based on semantic information – a task considerably different from trajectory prediction [9], [10]. Therefore in our experiments we mainly compare the IOC and CNN solutions to the problem against the ground truth distributions and the following baseline methods:

- 1) uniform distribution over \mathcal{M}
- 2) uniform distribution over the walkable states in \mathcal{M}
- 3) semantics-unaware *mapp* network [13]

For quantitative evaluation we measure *Kullback-Leibler divergence* (KL-div) between the predicted and the ground truth distribution:

$$D_{\mathrm{KL}}(P_{\mathrm{GT}}||Q_{\mathrm{Pred.}}) = \sum_{x \in \mathcal{M}} P_{\mathrm{GT}}(x) \log \frac{P_{\mathrm{GT}}(x)}{Q_{\mathrm{Pred.}}(x)}.$$
 (4)

We train and evaluate IOCMM, semapp and the baselines separately on the U4 and Stanford Drone datasets. Training and inference parameters are summarized in Table II.

We optimized the hyperparameters for IOCMM prior to the main experiments on a small portion of data from both datasets. For KL-div benchmarking we learn the θ weights from leave-one-out maps in the dataset, and validate the result in the remaining map, iterating over all maps in the respective dataset. Furthermore, we evaluate the impact of the number of simulated trajectories for inference in the new map (as in Eq. 1), measuring runtime and solution quality. Finally, we estimate the two $s_{0,g}$ sampling strategies, detailed in Section III-A, separately.

Training the mapp and semapp networks on the U4 dataset is straightforward as the size of the maps (32 by 32 pixels) is equivalent to the network input. From the larger images in the SDD (on average 146 by 152 pixels) we take 500 random crops of size 64 by 64 pixels. Each crop in the training data is augmented 7 times by rotation and mirroring. In both cases, leave-one-out maps are used for training and validation (50/50), and the remaining map is used for evaluation. We followed the

 TABLE II

 IOCMM AND SEMAPP PARAMETERS USED FOR TRAINING AND INFERENCE IN THE U4 AND STANFORD DRONE DATASETS

		U4 dataset			Stanford Drone dataset	
	Training		Inference	Training		Inference
IOCMM	Traj. batch B_t : 10 Map batch B_m : 7 Base reward r_0 : 0.01	Stoch. policy α : 0.1 Learning rate λ : 1.0	$s_{0,g}$ sampl. τ : 0.01 Num. sim. traj.: 500	Traj. batch B_t : 50 Map batch B_m : 5 Base reward r_0 : 0.01	Stoch. policy α : 0.1 Learning rate λ : 1.0	$s_{0,g}$ sampl. τ : 0.01 Num. sim. traj.: 1000
Semapp	Pooling layers: 1 Growth rate: 2 Layers per block: 5 Num. conv. layers: 19 Num. param.: 6.5 k Weight decay: 4.5 × 1	Crop size: 32 Batch size: 32 Dropout prob.: 0.35 Learning rate: 0.01 Learn. rate decay: 0.9985 0^{-5}	Num. crops: 1	Pooling layers: 1 Growth rate: 2 Layers per block: 5 Num. conv. layers: 19 Num. param.: 6.5 k Weight decay: 4.5 × 1	Crop size: 64 Batch size: 32 Dropout prob.: 0.35 Learning rate: 0.01 Learn. rate decay: 0.9985 0^{-5}	Num. crops: 500

TABLE III AVERAGE KL-DIV IN THE U4 AND STANFORD DRONE DATASETS

Method	U4 dataset	Stanford Drone
Uniform $p(s)$ over \mathcal{M}	1.21 ± 0.26	1.40 ± 0.31
Uniform $p(s)$ over walkable states in \mathcal{M}	0.97 ± 0.28	1.69 ± 0.42
Uniform $p(s f(s))$ learned from $\mathcal{T}_{\mathcal{M}}$	0.53 ± 0.09	1.09 ± 0.26
mapp CNN	0.93 ± 0.27	1.03 ± 0.19
IOCMM with learned $p(s_{0,q})$	0.42 ± 0.07	1.04 ± 0.24
IOCMM with modeled $p(s_{0,g})$	0.43 ± 0.07	1.21 ± 0.32
semapp CNN	0.37 ± 0.11	0.71 ± 0.13

training procedure and hyperparameters from [13] which turned out robust enough for our application. In particular, we trained the networks for 100 epochs with binary cross-entropy loss using the Adam optimizer [44], and stopped the training when the performance on the validation set did not improve for 15 epochs. Further parameters are detailed in Table II.

C. Semantic Segmentation

As a proof-of-concept that the semantic maps, required for the methods presented in this letter, can be obtained from images during runtime, we have set up a preliminary pipeline for semantic segmentation using UNet implementation in Keras and TensorFlow [45]. We trained the CNN with the 19 images from the Stanford Drone dataset, augmented 3 times with rotation, and tested on the remaining 6 images. Even with such negligible amount of data, our experiments (see Fig. 5) reached 0.64 frequency-weighted IoU in the training dataset (78% accuracy), and 0.53 IoU (69% accuracy) in the test dataset. We are confident, that using state-of-the-art semantic segmentation techniques [46] the performance, necessary for the application of our method, will be reached. Combining these two pipelines is of prime priority for our future work.

V. RESULTS

We report the mean and standard deviations of the KLdivergences for both datasets in Table III. In the U4 dataset, both IOCMM and semapp outperform the other baselines, furthermore both proposed sampling strategies for IOCMM show similar performance after appropriate hyperparameter optimization. Semapp, in addition to the quantitative improvement of minimum 14% over the closest baseline (IOCMM), offers a clear qualitative improvement in identifying crucial non-linearities in the predicted priors, as displayed in Fig. 6. This figure shows the extent to which semantics of the environment impact the distribution prediction - all semantics-unaware methods in our comparison, e.g. uniform p(s) over \mathcal{M} and mapp, perform poorly. On the contrary, in both datasets semapp outperforms all baselines, due to its ability to reason over spatially-connected regions using convolutions, learning not only local contexts where motion probability is high, but also which locations are usually avoided by pedestrians. Interestingly, in the SDD dataset, due to incomplete ground truth coverage of the scenes (as seen in Fig. 1 and 3), removing unwalkable spaces from the uniform distribution over all states in \mathcal{M} only decreases performance of this baseline. The reason here is that for many walkable states, where no motion is recorded in the ground truth, probabilities increase, resulting in worse KL-div scores. Despite using this imperfect training material, semapp consistently outperforms all baselines with the best KL-div score and smaller standard deviation between maps.

In Fig. 7, we visualize the generalization capabilities of the IOCMM method. To this end, we show the optimal θ^* weights in each individual map in both datasets, and compare them to the globally optimal set of weights, learned from training on all maps in the respective dataset. Here lies one benefit of the Inverse Optimal Control strategy to find occupancy priors: IOCMM does not only generalize well on a large amount of maps, but also retains high performance when learning from small amounts of data. In fact, when training on a fraction of maps from the dataset (e.g. as little as 10 random maps), IOCMM on average still converges to the globally-optimal θ^* costs for semantic classes, and thus the KL-div scores do not drop. This property is not shared by semapp, which needs a large selection of maps sufficiently covered by trajectories to generalize across various local contexts.

Both methods, IOCMM and semapp, depend on the number of random samples during inference. IOCMM samples trajectories between random start and goal positions, while semapp samples random crops from the larger semantic map. In Fig. 8 we show the relation between the number of samples, performance and inference time in the large maps of the SDD dataset. Runtimes were measured for Python implementations of both algorithms on an ordinary laptop with Intel Xeon 2.80 GHz ×8 CPU and 32 GB of RAM. The CNN is implemented using Theano on the built-in GPU Quadro M2000 M. It is worth mentioning that the inference time of semapp on one crop of size 64 by 64 pixels (equivalent to 25.6×25.6 m) is ~ 0.054 seconds, appropriate for real-time application.



Fig. 6. Qualitative comparison of results in the U4 dataset. A binary *occupancy map* of the environment highlights the amount of structure imposed by semantics in urban scenes. There is little surprise that the semantics-unaware mapp approach for learning occupancy priors [13] is not learning any meaningful behaviors apart from the fact that people (often) tend to be found close to obstacles. On the contrary, IOCMM correctly estimates the priors in different walkable areas. On top of that, CNN-based semapp is capable of detecting all "illegal crosswalks" in these scenes, as well as cutting over grass in such places where the topology of the environment encourages to do so, e.g. see the sharp corner in the third row.



Fig. 7. Optimal θ costs of various semantic classes, learned by IOCMM in each individual map, are shown in gray. Globally optimal weights for the entire dataset are overlaid in black. Left: U4 dataset. Right: Stanford Drone dataset.



Fig. 8. Runtime and performance of IOCMM (left) and semapp (right) in the SDD dataset as a function of the number of sampled trajectories/crops during inference. Both methods' runtime scales linearly with the number or random samples, while performance improves exponentially.

VI. CONCLUSION

In this letter, we research the problem of learning human occupancy priors in semantically-rich urban environments using only the semantic map as input. Considering two established classes of approaches to this end (Inverse Optimal Control and Convolutional Neural Networks), we show that our CNN-based semapp approach is outperforming all baselines already with limited training data. The IOCMM approach, on the other hand, can be used to reasonably estimate the costs of semantic classes from several maps and few trajectories. However, it is limited to constant weights, which may not reflect the behavior of people in all local contexts of semantically-complex environments. This approach lacks reasoning on spatial relevance of surfaces to infer cases where people may prefer to walk on one surface class over another, or not walk at all.

In future work, we intend to further investigate the possibilities of applying advanced IOC techniques, for instance, non-linear IRL with complex features [47], non-linear reward modeling [48], automated feature extraction to exploit local correlations in the environment [49], [50], IOC with multiple locally-consistent reward functions [51], and with CNN-based reward function approximator [52]. Furthermore, we plan to validate semapp with on-the-fly semantics estimation and extend it to first-person view for application in automated driving to infer potential pedestrians' entrance points to the road.

REFERENCES

- A. Rudenko, L. Palmieri, M. Herman, K. M. Kitani, D. M. Gavrila, and K. O. Arras, "Human motion trajectory prediction: A survey," *Int. J. Robot. Res.*, vol. 32, pp. 895–935, 2020.
- [2] D. Ellis, E. Sommerlade, and I. Reid, "Modelling pedestrian trajectory patterns with Gaussian processes," in *Proc. 12th IEEE Int. Conf. Comp. Vis. Workshop*, 2009, pp. 1229–1234.
- [3] Z. Wang, P. Jensfelt, and J. Folkesson, "Modeling spatial-temporal dynamics of human movements for predicting future trajectories," in *Workshop Proc. AAAI Conf. Artif. Intell. Knowl., Skill, Behav. Transfer Auton. Robots*, 2015.
- [4] T. P. Kucner, M. Magnusson, E. Schaffernicht, V. H. Bennetts, and A. J. Lilienthal, "Enabling flow awareness for mobile robots in partially observable environments," *IEEE Robot. Automat. Lett.*, vol. 2, no. 2, Apr. 2017, pp. 1093–1100.
- [5] E. Rehder and H. Kloeden, "Goal-directed pedestrian prediction," in *Proc. IEEE Int. Conf. Comput. Vis. Workshops*, 2015, pp. 50–58.
- [6] J. Wu, J. Ruenz, and M. Althoff, "Probabilistic map-based pedestrian motion prediction taking traffic participants into consideration," in *Proc. IEEE Intell. Veh. Symp.*, 2018, pp. 1285–1292.
- [7] E. A. Pool, J. F. Kooij, and D. M. Gavrila, "Context-based cyclist path prediction using recurrent neural networks," in *Proc. IEEE Intell. Veh. Symp.*, 2019, pp. 824–830.
- [8] Y. Hu, W. Zhan, and M. Tomizuka, "Probabilistic prediction of vehicle semantic intention and motion," in *Proc. IEEE Intell. Veh. Symp.*, 2018, pp. 307–313.
- [9] V. Karasev, A. Ayvaci, B. Heisele, and S. Soatto, "Intent-aware long-term prediction of pedestrian motion," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2016, pp. 2543–2549.
- [10] K. M. Kitani, B. D. Ziebart, J. A. Bagnell, and M. Hebert, "Activity forecasting," in *Proc. Europ. Conf. Comp. Vis.*, 2012, pp. 201–214.
- [11] N. Rhinehart and K. M. Kitani, "First-person activity forecasting from video with online inverse reinforcement learning," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 42, Feb. 2018, pp. 304–317.
- [12] T. van der Heiden, N. S. Nagaraja, C. Weiss, and E. Gavves, "SafeCritic: Collision-aware trajectory prediction," 2019, arXiv:1910.06673.
- [13] J. Doellinger, M. Spies, and W. Burgard, "Predicting occupancy distributions of walking humans with convolutional neural networks," *IEEE Robot. Automat. Lett.*, vol. 3, no. 3, pp. 1522–1528, 2018.
- [14] A. Robicquet, A. Sadeghian, A. Alahi, and S. Savarese, "Learning social etiquette: Human trajectory understanding in crowded scenes," in *Proc. Europ. Conf. Comp. Vis.*, 2016.
- [15] S. Pellegrini, A. Ess, K. Schindler, and L. van Gool, "You'll never walk alone: Modeling social behavior for multi-target tracking," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2009, pp. 261–268.
- [16] B. Lau, K. O. Arras, and W. Burgard, "Tracking groups of people with a multi-model hypothesis tracker," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2009.
- [17] T. Linder, S. Breuers, B. Leibe, and K. O. Arras, "On multi-modal people tracking from mobile platforms in very crowded and dynamic environments," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2016.
- [18] A. F. Foka and P. E. Trahanias, "Probabilistic autonomous robot navigation in dynamic environments with human motion prediction," *Int. J. Social Robot.*, vol. 2, no. 1, pp. 79–94, 2010.
- [19] H. Bai, S. Cai, N. Ye, D. Hsu, and W. S. Lee, "Intention-aware online POMDP planning for autonomous driving in a crowd," in *Proc. IEEE Int. Conf. Robot. Automat.*, May 2015.
- [20] L. Palmieri, T. P. Kucner, M. Magnusson, A. J. Lilienthal, and K. O. Arras, "Kinodynamic motion planning on Gaussian mixture fields," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2017, pp. 6176–6181.
- [21] B. Okal and K. O. Arras, "Learning socially normative robot navigation behaviors with bayesian inverse reinforcement learning," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2016.
- [22] S.-Y. Chung and H.-P. Huang, "Incremental learning of human social behaviors with feature-based spatial effects," in *Proc. IEEE Int. Conf. Intell. Robots Syst.*, 2012, pp. 2417–2422.
- [23] A. Rudenko, L. Palmieri, A. J. Lilienthal, and K. O. Arras, "Human motion prediction under social grouping constraints," in *Proc. IEEE Int. Conf. Intell. Robots Syst.*, 2018.
- [24] P. A. Lasota, T. Fong, and J. A. Shah, "A survey of methods for safe human-robot interaction," *Found. Trends Robot.*, vol. 5, no. 4, pp. 261–349, 2017.
- [25] A. Alahi, K. Goel, V. Ramanathan, A. Robicquet, L. Fei-Fei, and S. Savarese, "Social LSTM: Human trajectory prediction in crowded spaces," in *Proc. IEEE Conf. Comp. Vis. Pat. Rec.*, 2016, pp. 961–971.

- [26] S.-Y. Lo, S. Alkoby, and P. Stone, "Robust motion planning and safety benchmarking in human workspaces," in *SafeAI*, AAAI, 2019.
- [27] P. Trautman and A. Krause, "Unfreezing the robot: Navigation in dense, interacting crowds," in *Proc. IEEE Int. Conf. Intell. Robots Syst.*, 2010, pp. 797–803.
- [28] L. Ballan, F. Castaldo, A. Alahi, F. Palmieri, and S. Savarese, "Knowledge transfer for scene-specific motion prediction," in *Proc. Europ. Conf. Comp. Vis.*, 2016, pp. 697–713.
- [29] S. Srikanth, J. A. Ansari, K. R. Ram, S. Sharma, J. K. Murthy, and K. M. Krishna, "INFER: INtermediate Representations for Future Prediction," in *Proc. RSJ Int. Conf. Intell. Robots Syst.*, 2019, arXiv:1903.10641.
- [30] M. Shen, G. Habibi, and J. P. How, "Transferable pedestrian motion prediction models at intersections," in *Proc. IEEE RSJ Int. Conf. Intell. Robots Syst.*, 2018, arXiv:1804.00495.
- [31] P. Coscia, F. Castaldo, F. A. N. Palmieri, A. Alahi, S. Savarese, and L. Ballan, "Long-term path prediction in urban scenarios using circular distributions," *Image Vis. Comput.*, vol. 69, pp. 81–91, 2018.
- [32] D. Munoz, J. A. Bagnell, and M. Hebert, "Stacked hierarchical labeling," in Proc. Eur. Conf. Comput. Vis., Berlin, Germany: Springer, 2010, pp. 57–70.
- [33] V. Badrinarayanan, A. Kendall, and R. Cipolla, "SegNet: A deep convolutional encoder-decoder architecture for image segmentation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 39, no. 12, pp. 2481–2495, Dec. 2017.
 [34] Z. Wu, C. Shen, and A. van denHengel, "Wider or deeper: Revisiting
- [34] Z. Wu, C. Shen, and A. van denHengel, "Wider or deeper: Revisiting the resnet model for visual recognition," *Pattern Recognit.*, vol. 90, pp. 119–133, 2019.
- [35] L. Gan, R. Zhang, J. W. Grizzle, R. M. Eustice, and M. Ghaffari, "Bayesian spatial kernel smoothing for scalable dense semantic mapping," *IEEE Robot. Automat. Lett.*, vol. 5, no. 2, pp. 790–797, Apr. 2020.
- [36] H. Grimmett *et al.*, "Integrating metric and semantic maps for visiononly automated parking," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2015, pp. 2159–2166.
- [37] L. Yu, T. Yu, C. Finn, and S. Ermon, "Meta-inverse reinforcement learning with probabilistic context variables," in *Proc. 33rd Conf. Neural Inf. Process. Syst.*, 2019.
- [38] K. Xu, E. Ratner, A. Dragan, S. Levine, and C. Finn, "Learning a prior over intent via meta-inverse reinforcement learning," in *Proc. 36th Int. Conf. Mach. Learn.*, 2019, pp. 6952–6962.
- [39] B. D. Ziebart, A. Maas, J. A. Bagnell, and A. K. Dey, "Maximum entropy inverse reinforcement learning," in *Proc. AAAI Conf. Artif. Intell.*, 2008, pp. 1433–1438.
- [40] B. D. Ziebart et al., "Planning-based prediction for pedestrians," in Proc. IEEE Int. Conf. Intell. Robots Syst., 2009, pp. 3931–3936.
- [41] M. Kuderer, H. Kretzschmar, C. Sprunk, and W. Burgard, "Feature-based prediction of trajectories for socially compliant navigation," in *Proc. Int. Conf. Robot.: Sci. Syst.*, 2012.
- [42] A. Tamar, Y. Wu, G. Thomas, S. Levine, and P. Abbeel, "Value iteration networks," in Adv. Neural Inf. Proc. Syst., 2016, pp. 2154–2162.
- [43] S. Jégou, M. Drozdzal, D. Vazquez, A. Romero, and Y. Bengio, "The one hundred layers tiramisu: Fully convolutional denseNets for semantic segmentation," in *Proc. IEEE Conf. Comput. vis. Pattern Recognit. Workshops*, 2017, pp. 11–19.
- [44] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, arXiv:1412.6980.
- [45] D. Gupta, "Semantic segmentation using keras." [Online]. Available: https: //github.com/divamgupta/image-segmentation-keras
 [46] Cityscapes Dataset, "Cityscapes Pixel-Level Semantic Labeling
- [46] Cityscapes Dataset, "Cityscapes Pixel-Level Semantic Labeling Benchmark." [Online]. Available: https://www.cityscapes-dataset.com/ benchmarks/.
- [47] S. Levine, Z. Popović, and V. Koltun, "Nonlinear inverse reinforcement learning with gaussian processes," in *Proc. 24th Int. Conf. Neural Inf. Process. Syst.*, 2011, pp. 19–27.
- [48] J. Mainprice, A. Byravan, D. Kappler, D. Fox, S. Schaal, and N. Ratliff, "Functional manifold projections in Deep-LEARCH," in *NIPS Workshop Neurorobotics*, 2016.
- [49] S. Levine, Z. Popovic, and V. Koltun, "Feature construction for inverse reinforcement learning," in Adv. Neural Inf. Process. Syst., 2010, pp. 1342– 1350.
- [50] A. M. Metelli, M. Pirotta, and M. Restelli, "Compatible reward inverse reinforcement learning," in *Adv. Neural Inf. Process. Syst.*, 2017, pp. 2050– 2059.
- [51] Q. P. Nguyen, B. K. H. Low, and P. Jaillet, "Inverse reinforcement learning with locally consistent reward functions," in *Adv. Neural Inf. Process. Syst.*, 2015, pp. 1747–1755.
- [52] M. Wulfmeier, P. Ondruska, and I. Posner, "Maximum entropy deep inverse reinforcement learning," 2015, arXiv:1507.04888.