

I want my coffee hot! Learning to find people under spatio-temporal constraints

Gian Diego Tipaldi

Kai O. Arras

Abstract—In this paper we present a probabilistic model for spatio-temporal patterns of human activities that enable robots to blend themselves into the workflows and daily routines of people. The model, called *spatial affordance map*, is a non-homogeneous spatial Poisson process that relates space, time and occurrence probability of activity events. We describe how learning and inference is made and present a novel planning algorithm that produces paths which maximize the probability to encounter a person. We show that the problem is a special class of the orienteering problem that can be solved as a finite horizon Markov decision process.

We develop a simulator of populated office environments to validate the model and the planning algorithm. The simulated agents follow activity patterns learned by administering a questionnaire to 27 colleagues over two weeks. The experiments shows that the model is statistically valid with respect to both the Anderson-Darling test and the expected waiting time estimation. They further show that the proposed algorithm is able to find optimal paths.

I. INTRODUCTION

Robots that operate in human environments require the ability to sense people and recognize their activities. But beyond that, they also need the ability to reason about the places and times when and where people are engaged into which activity. Such a model – over large time and space scales – enables a robot to coordinate its motions, tasks and schedules with the patterns of human activities, giving it the ability to smoothly blend itself into the workflows and daily routines of people. We believe that this ability is key in the attempt to build socially acceptable robots for many domestic and service applications.

In this paper, we approach this problem by the *spatial affordance map*, a model that represents human activity events as a rate function of a non-homogenous spatial Poisson process. We present how learning and inference is made with the spatial affordance map and propose an algorithm able to find paths that maximize the probability of finding a person. This is a under-explored planning problem to our knowledge, relevant in a number of applications: a health care robot that needs to find a nurse fast, a surveillance robot that must find a patrolling human colleague, a receptionist robot that can answer the query of finding a particular person at a given time, or a delivery robot that brings urgent goods to people such as coffees getting cold (Fig. 1).

The problem of planning maximum encounter probability paths has not yet been addressed in this sense to



Fig. 1. Example scenario. The picture shows an example path computed by the robot in our simulated environment.

our knowledge. Most related is the work by Roy *et al.* [1] in which the authors approach a similar problem as with a partially observable Markov decision process where the position of a person is not observable and modeled with a sample-based distribution. In contrast to our approach, they have no prior information on how the environment is used by the people, a vital information especially in large scale environments.

Learning models of human behavior has been addressed by several researchers. Kruse and Wahl [2] propose statistical grids whose cells hold temporal occupancy probabilities of people and stochastic trajectories which are paths of dynamic objects along which their appearance probability is modeled by a Poisson process. The goal is to assess and plan minimal collision probability paths. The grid and the trajectories are learned from ceiling-mounted cameras.

Bruce and Gordon [3] learn goal locations in an environment from trajectories obtained by a laser-based people tracker. Based on the assumption that people move in a goal-oriented fashion, paths are planned from the location of a person being tracked to the goal locations.

Bennewitz *et al.* [4] learn typical motion patterns that people follow in an environment. The approach collects trajectories of people with multiple statically mounted laser scanners and combine similar trajectories to motion patterns using EM clustering. From each pattern a Hidden Markov Models is derived which enables a mobile robot to predict the motion of people and to adapt its navigation behavior accordingly.

Not only focused on human motion is the work by Ihler and Smiyth [5]. The authors presents a non parametric approach to learn time profiles of human activities. The

rate function of a Poisson process is learned using non parametric Bayesian models: the infinite mixture model with a Dirichlet process prior. Although interesting and related to our work, their approach does not consider the spatial variation of activities.

All these works consider either the special case of human motion or lack the ability to make inference in both time and space. In contrast, the spatial affordance map is a single representation for inference about spatio-temporal behavior of people that coherently relates time, space and occurrence probability of activity events. Accordingly, it is able to process the following three queries that ask for occurrence probability, time and space, respectively:

- What is the probability of an event within a given time interval and area of space?
- How long do I have to wait until an event happens with a given confidence in an area of space?
- Given a deadline, what is the path along which the encounter probability of an event is optimized?

We will present the theory and give application examples how the map is able to produce answers to these queries. The name of the map lends itself from the view of human activities as affordances that an environment offers to human agents.

The paper is structured as follows: the next section gives the theory of the spatial affordance map followed by section III that explains how inference is made and section IV that develops the planning algorithm. Section V describes the people simulator we used in the experiments, while section VI contains the experimental results. Finally, section VII concludes the paper.

II. SPATIAL AFFORDANCE MAP

The spatial affordance map is a non-homogeneous spatial Poisson process. In this section, we explain its theory and how learning is done in this case of a Poisson process.

Under the assumption that events in time occur independently of one another, a *Poisson process* can deal with distributions of time intervals between events. Concretely, let $N(t)$ be a discrete random variable to represent the number of events occurring up to time t with rate λ . Then we have that $N(t)$ follows a Poisson distribution with parameter λt

$$P(N(t) = k) = \frac{e^{-\lambda t} (\lambda t)^k}{k!} \quad k = 0, 1, \dots \quad (1)$$

In general, the rate parameter may change over time. In this case, the generalized rate function is given as $\lambda(t)$ and the expected number of events between time a and b is

$$\lambda_{a,b} = \int_a^b \lambda(u) du. \quad (2)$$

A homogeneous Poisson process is a special case of a non-homogeneous process with constant rate $\lambda(t) = \lambda$.

The *spatial* Poisson process introduces a spatial dependency on the rate function given as $\lambda(\vec{x}, t)$ with $\vec{x} \in X$

where X is a vector space such as \mathbb{R}^2 or \mathbb{R}^3 . For any subset $S \subset X$ of finite extent (e.g. an area in space), the number of events occurring inside this area can be modeled as a Poisson process with associated rate function $\lambda_S(t)$ such that

$$\lambda_S(t) = \int_S \lambda(\vec{x}, t) d\vec{x}. \quad (3)$$

In the case that this generalized rate function is a separable function of time and space, we have:

$$\lambda(\vec{x}, t) = f(\vec{x})\lambda(t) \quad (4)$$

for some function $f(\vec{x})$ for which we can demand

$$\int_X f(\vec{x}) d\vec{x} = 1 \quad (5)$$

without loss of generality. This particular decomposition allows us to decouple the occurrence of events between time and space. Given Eq. 5, $\lambda(t)$ defines the occurrence rate of events, while $f(\vec{x})$ can be interpreted as a probability distribution on where the event occurs in space.

Learning the spatio-temporal distribution of events in an environment is equivalent to learn the generalized rate function $\lambda(\vec{x}, t)$. However, learning the full continuous function is a highly expensive process. For this reason, we approximate the non-homogeneous spatial Poisson process with a piecewise homogeneous one. The approximation is performed by discretizing the environment into a tridimensional grid, where each cell represents a local – in terms of space and time – homogeneous Poisson process with a constant rate,

$$P_{ij\tau}(N(t) = k) = \frac{e^{-\lambda_{ij\tau}(t-t_\tau)} (\lambda_{ij\tau}(t-t_\tau))^k}{k!} \quad (6)$$

with $k = 0, 1, \dots$ and $t_\tau \leq t < t_{\tau+1}$ and where $\lambda_{ij\tau}$ is assumed to be constant. Finally, the spatial affordance map represents the generalized rate function $\lambda(\vec{x}, t)$ using a grid approximation,

$$\lambda(\vec{x}, t) \simeq \sum_{ij\tau} \lambda_{ij\tau} \mathbf{1}_{ij\tau}(\vec{x}, t) \quad (7)$$

with $\mathbf{1}_{ij\tau}(\vec{x}, t)$ being the indicator function.

Instead of a grid approximation, other tessellation schemes in space and time such as octrees, regions of homogeneous Poisson rates or function approximators [5] can equally be used. Subdivision of space and time into regions of fixed Poisson rates has two interesting properties. First, having intervals of constant rate over time, the preferable decomposition in Eq. 4 holds. Second, we can instantly infer properties of the environment without computing expensive integrals as to be shown in the next section.

We take a Bayesian learning approach using Gamma priors to estimate the Poisson rate parameter of each cell. We discard the use of a maximum likelihood approach since, without priors, it cannot properly initialize never observed cells. The map is learned from human activity observations $k_{1,\dots,n}$ that can be obtained either from ceiling-mounted

cameras [2], wearable devices or the exteroceptive sensors of the mobile robot as in Luber *et al.* [6]. Here they will come from a simulator of people in office environments described in Section V. The Bayesian approach allows for learning of the spatial affordance map by counting in a grid. This makes life-long learning particularly simple as new information can be added at any time by one or multiple robots. More details on learning the spatial affordance map can also be found in [6].

III. INFERENCE IN THE SPATIAL AFFORDANCE MAP

This section describes how inference in the spatial affordance map is made. For this purpose we will only consider activity events k_i that correspond to observations of people in the environment (in [6] this has been implemented using track confirmations from a people tracker). We will show how the probability of encountering at least one person along a path can be computed and, as a side effect, how long does the robot need to wait until this probability is above a confidence.

Waiting Time Estimation

One of the key properties of the Poisson process is that it is able to explicitly model waiting time. For notation simplicity, we assume that the robot arrives at a location and starts waiting at time t_0 . Let us consider T_1 as the time when the first person arrives. Clearly, this person arrives at time $T_1 > t$ if and only if the number of arrivals between time t and t_0 is 0. Combining this property with the probability of a number of homogeneous Poisson events in an interval gives

$$p(T_1 > t; t_0) = p(N(t) - N(t_0) = 0) \quad (8)$$

$$= e^{-\int_{t_0}^t \lambda(u) du}, \quad (9)$$

where for simplicity we have ignored the spatial dimension. Then, the waiting time until the first arrival T_1 has an exponential distribution with parameter $\lambda(t)$

$$p(t; \lambda(t)) = \begin{cases} \lambda(t)e^{-\int_{t_0}^t \lambda(u) du} & \text{if } t \geq t_0, \\ 0 & \text{if } t < t_0. \end{cases} \quad (10)$$

Defining the area of interaction $I(\vec{x})$ as a function of the robot pose (which e.g. can be a room), we have that the waiting time in the position \vec{x} has an exponential distribution with parameter

$$\lambda_I(\vec{x}, t) = \int_{I(\vec{x})} \lambda(\vec{\chi}, t) d\vec{\chi} \quad (11)$$

where $\lambda(\vec{x}, t)$ is the general rate function.

The waiting time is then estimated considering the cumulative density function of the exponential distribution Eq. 10. Let ρ be a confidence value such as 0.95 or 0.99, the waiting time is the value of t such that $p(T_1 \leq t) \geq \rho$. Considering the cumulative density function of the exponential distribution we have

$$p(T_1 \leq t) = 1 - e^{-\int_{t_0}^t \lambda_I(\vec{x}, u) du} \quad (12)$$

where we only consider the positive axis for notation simplicity. This leads to

$$p(T_1 \leq t) \geq \rho \implies 1 - e^{-\int_{t_0}^t \lambda_I(\vec{x}, u) du} \geq \rho \quad (13)$$

and

$$\int_{t_0}^t \lambda_I(\vec{x}, u) du \geq -\log(1 - \rho). \quad (14)$$

A closed form solution to this integral inequality depends on the form of the rate function employed and is not always available. In the case of our grid, we can exploit the piecewise constant approximation and the solution for the homogeneous process case. In this way, we obtain a recursive formula to compute the minimum waiting time $w(t_0, \rho)$ for a confidence value ρ and a starting time t_0 as:

$$w(t_0, \rho) \geq \begin{cases} t_0 - \frac{\log(1-\rho)}{\lambda_{ij\tau}} & \text{if } p(T_1 \leq t_{\tau+1}) \geq \rho, \\ w(t_{\tau+1}, \hat{\rho}) & \text{otherwise} \end{cases} \quad (15)$$

$$\hat{\rho} = 1 - e^{\log(1-\rho) + (t_{\tau+1} - t_{\tau})\lambda_{ij\tau}} \quad (16)$$

This is the expected waiting time given a confidence value.

Knowledge about waiting times can help a robot to coordinate its activities with the learned patterns of human activities in an environment. However, if a robot needs to find a person *proactively*, planning is required which is considered hereafter.

Encounter Probability Paths

Formally, we seek to find the path, \mathcal{P} , that maximizes the encounter probability within a given time t_{max} . Since a path is a mapping from time to space, $\mathcal{P} : t \rightarrow \vec{x}$, we have that the number of people encountered in a certain path follows a non-homogeneous Poisson process whose rate function depends on the path itself

$$\lambda^{\mathcal{P}}(t) = \lambda(\mathcal{P}(t), t). \quad (17)$$

The probability of encountering at least a person along a path is obtained by considering the probability of not encountering anyone and using the law of total probability

$$\begin{aligned} p(N^{\mathcal{P}}(t_{max}) > 0) &= 1 - p(N(t_{max}) = 0) \\ &= 1 - e^{-\int_{t_0}^{t_{max}} \lambda^{\mathcal{P}}(u) du}, \end{aligned} \quad (18)$$

resulting in the best path being

$$\mathcal{P}^* = \underset{\mathcal{P}}{\operatorname{argmax}} p(N^{\mathcal{P}}(t_{max}) > 0). \quad (19)$$

IV. MAXIMUM ENCOUNTER PROBABILITY PLANNING

Two aspects have to be addressed in order to obtain valid paths \mathcal{P} . A first aspect is to compute the path rate described in Eq. 17. Considering the grid approximation and assuming the robot moves at constant speed, the path can be approximated with the sequence of grid cells the robot traverses. The integral then becomes the sum of the Poisson rate over those cells. A second problem arises in the maximization step of Eq. 19. This problem is an instance of the *orienting problem* that has been shown

Algorithm 1: Encounter Probability Planning

In: Rate $\lambda(\vec{x}, t)$; time t_{max} ; initial state s_0 ;**Out:** The best path \mathcal{P}^* ;

```

1 Compute the horizon  $N$ ;
2  $J_N(s) \leftarrow \lambda_{ij\tau} \forall s$ ;
3 for  $k \leftarrow N - 1$  to 0 do
4    $J_k(s) \leftarrow \max_a \left[ R(s, a) + \sum_{s'} p(s'|s, a) J_{k+1}(s') \right]$ ;
5    $A_k^*(s) \leftarrow \operatorname{argmax}_a \left[ R(s, a) + \sum_{s'} p(s'|s, a) J_{k+1}(s') \right]$ ;
6 end
7  $\mathcal{P}^*(0) \leftarrow s_0$ ;
8 for  $k \leftarrow 1$  to  $N$  do
9    $s \leftarrow \mathcal{P}^*(k - 1)$ ;
10   $\mathcal{P}^*(k) \leftarrow \mathbb{E} [p(s'|s, A_{k-1}^*(s))]$ ;
11 end
12 return  $\mathcal{P}^*$ ;
```

to be NP-hard [7]. The essence of the problem is to find a path that maximize the sum of certain rewards within a limited amount of time, where in our setting, the rewards are the individual Poisson rates of each cell.

In our case, it can be shown that in our setting a finite-horizon Markov decision process (MDP) can solve the problem in polynomial time using dynamic programming and the Bellmann equation. This is possible since the map consists in a regular grid in Cartesian space that allows us to propagate the path rewards in a flooding manner.

Formally, an MDP is a probabilistic model for sequential decision problems. At each time step, the process is in some state $s \in \mathcal{S}$, and the decision maker may choose any action $a \in \mathcal{A}(s)$ that is available in state s . The process responds at the next time step by randomly moving into a new state s' , and giving the decision maker a corresponding reward $R(s, a)$. The probability that the system will move into s' is influenced by the chosen action according to the state transition distribution $p(s'|s, a)$ but it is conditionally independent from any previous state or action taken, obeying to the Markov property.

Our state space is represented by the cells of the spatial affordance map that are within the free space of the environment and the time interval of the day the robot is in that cell, $s = \{i, j, \tau\}$. In each cell, a set of (maximum) nine actions are defined and they account for movements in the 8-neighborhood and waiting in the current cell. The state transition distribution may be derived from the robot odometry reflecting its accuracy. The reward function is defined as the reward of the state s' , obtained by performing action a in state s and is defined as

$$R(s, a) = \sum_{s'} \lambda_{i'j'\tau'} p(s'|s, a). \quad (20)$$

To obtain a maximum encounter probability path given an initial state s_0 and a deadline t_{max} we seek to find a deterministic policy $\pi = \{s_0, a_0, \dots, s_n, a_n, \dots\}$ such that

it maximizes the expected total reward. By discretizing the time in a finite number of steps called horizon N , this minimization problem is solved using backwards induction, solving N single stage problems of increasing horizon. The resulting procedure is given in Algorithm 1.

V. SIMULATOR

For evaluating the model and the planner, we developed a people simulator for office-like environments. The simulator models the floor of our university building (see Fig. 1). Simulation in this case is needed as with real humans, experiments cannot be reproduced and simulated agents do not change their behavior in the presence of robots.

The simulator models a generic work day from 8 am to 7 pm in which a number of agents perform typical office activities (ten in our case). To learn realistic activity patterns, an anonymous questionnaire has been handed out to 27 colleagues. The subjects filled in their activities over a work day (e.g. arriving to work, working, eating, smoking, drinking coffee, going to the restroom, etc.) including the time and duration of each activity. From this information we learned a discrete distribution of when each activity is performed over a two week period. To generate the actual activities of the simulated agents we sample from these distributions.

The engine follows the three-layered agent architecture from [8] that in our case consists in the layers *activity scheduler*, *activity executor* and *action executor*. At the beginning of the day, the activity scheduler randomly generates a fixed schedule for each agent. Every activity is composed of a set of actions such as enter, move, stay or leave, which in turn are activated and deactivated by the activity executor. Once an action is activated, the action executor takes care of its progress and signals back when it reached its final state. The actual plans are generated using A* with action costs that are randomly perturbed to simulate some motion variability. Each time an agent is engaged into an activity, its type and place form an activity observation k_i to learn the map.

VI. EXPERIMENTS

For the experiments, we come back to the three types of queries given in section I that ask for occurrence probability, time and space. The first query is the forward application of the Poisson process. It allows us to understand if the model is statistically sound, that is, if a spatial Poisson process is an appropriate model for patterns of human activities. The second query is related to waiting time estimation. By testing the prediction accuracy of the map we again obtain experimental evidence for its statistical validity. The maximum encounter probability planner is then evaluated to demonstrate the third type of query.

For the experiment, the map is learned from activities of ten agents over ten days, followed by a series of ten testing days. We use a grid resolution of $0.25m$ in x and y and 1 hour in time resulting in 11 time slots per working day.

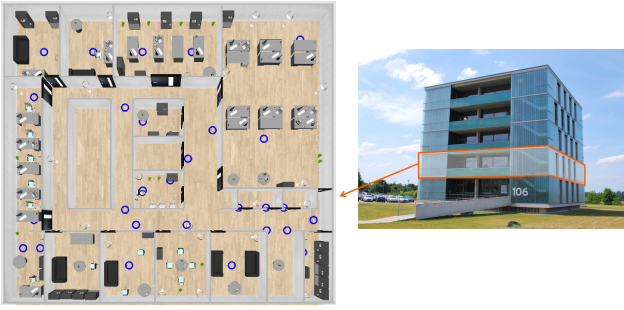


Fig. 2. The simulator that models behavior of people in office environments describing one floor in our building. The picture shows a top view of the simulated environment.

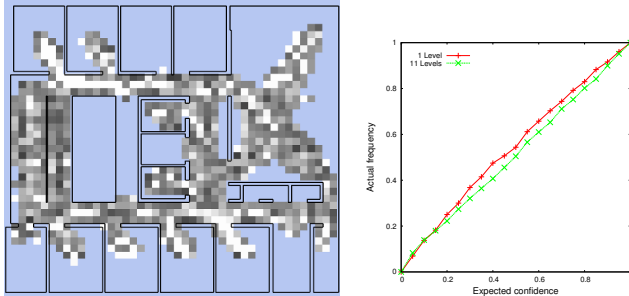


Fig. 3. **Left:** Anderson-Darling test shown for the simulated environment. Non-white cells pass the test (the darker the better), light blue cells have never been visited by agents and are excluded from the test. Overall, 95% of the cells pass the test. **Right:** Waiting time experiment. The diagram shows the actual frequency versus the expected confidence for the predicted waiting time. The plot shows close correspondence to the ideal result which is the diagonal.

A. Model Evaluation: Anderson-Darling Test

The first experiment quantifies the goodness of fit of the map with respect to the learned rate and its prediction capability. The goodness of fit is tested using the Anderson-Darling test. The test looks for evidence that a given sample of data did not arise from a given probability distribution.

In our setting, we tested each cell of the spatial affordance map and checked if the interarrival times of the people follow the learned exponential distribution or not. The results of this test are shown in Fig. 3. The map shows the Anderson-Darling test score for the cells in the environment. White cells are places where the test failed, darker cells are places where the test was successful (the darker the better). Light blue cells have never been used by the simulated agents and are ignored.

Overall, the test was successful on 95% of the cells. To study the effect of the time-variable Poisson rate, that is, the case of a temporally homogeneous Poisson process, we have also tested a map with a single, eleven hour time slot. The success rate for this map is only 54%. Although this may seem a low value, the next test will show that it will not compromise the prediction performances. The main reason is that when the model is not accurate enough, it tends to yield conservative probability estimates.




Fig. 5. Encounter probability path experiment. The picture shows the success rate versus the path length for the different strategies. As expected, the MDP strategy outperforms the heuristic planners.

B. Model Evaluation: Waiting Time

In this test, we check the ability of the model to actually predict the waiting time until a person is encountered up to a certain confidence. We perform the test by sampling 1000 random positions in space and time for a set of confidence values ranging from 5% to 95% with an increment of 5%. In each of these space/time-positions, the waiting time for the first arrival event has been computed according to Eq. 15. During the testing phase, we check the frequency of encountering a person during this interval and compared it to the expected confidence. In the perfect model these two values coincide. Fig. 3 shows the corresponding curve, for both the regular map and the map with a single eleven hour slot, where we added the trivial values for 0% and 100% confidence. As can be seen, the actual frequency closely matches the predicted confidence, resulting in very precise waiting time predictions. Even the single time slot map is accurate despite the low rate in the Anderson Darling test. The insight is that the model is intrinsically conservative, also shown by the curve being higher than the diagonal.

C. Maximum Encounter Probability Planning

The experiment that corresponds to the third query demonstrates the ability of the MDP planner to generate optimal paths in terms of encounter probability. We compare the MDP paths with four other heuristic strategies: local goal (LOC), global goal (GLO), random walk (RND) and waiting (WAT). The local and global goal strategies are informed in that they use the spatial affordance map. They are both greedy as they plan the shortest paths to the cell with the highest encounter probability using A*. This cell is either computed in a local neighborhood (LOC), or in the entire map (GLO). If the deadline is shorter than the length of the resulting path, the path is truncated, otherwise the path waits at the local or global maximum, respectively. The RND strategy follows a random walk from the start location and the WAT strategy simply waits on the spot until a person drops by. We gen-




Fig. 4. Maximum encounter probability planning. The red path is the solution of the MDP strategy, the blue path the result of the GLO strategy. Figures a), b) and c) show the MDP paths from the same start location with increasing deadlines. The resulting paths lead the robot to regions of increasing Poisson rates. Figure d) compares the MDP and the LOC strategies for the same start location and deadline. While the LOC strategy greedily aims for the shortest path, the MDP algorithm correctly finds a path across a high-density area that maximizes the encounter probability.

erate 500 random positions in space-time, for each of them we compute the paths with increasing lengths ranging from 25 to 361 with an increment of 24 steps – corresponding to deadlines from 1 to 15 minutes with increments of 1 minute at the speed of $0.4m/s$. For different speeds more steps may be considered but the curves will show the same behavior. Fig. 5 shows the result of the comparison while Fig. 4 gives example paths.

As can be seen, the MDP procedure is the best strategy, waiting in the spot is the worse strategy. Interestingly, there is a clear difference between the *informed* strategies (MDP, LOC, GLO) and the *uninformed* strategies (RND, WAT), showing the importance of the spatial affordance map for this type of task. The small difference between the LOC and the GLO strategies is due to the size of the environment, where after some time the global and local maximum are in the same place. Both strategies are not much worse than the MDP which is because they have been developed to imitate its behavior. The MDP strategy also tries to reach the global maximum and waits there, but it does so by trading off path length with encounter probability. A clear difference can be seen in Fig. 4 d), where the MDP solution (red) takes the longer path to the maximum across a high density area, where the LOC path (blue) simply takes the shortest path through a low density region. Such differences will become more accentuated with larger environments and more complex topologies.

These results clearly show the importance of prior knowledge of human space usage, an information not used in prior work [1] representing one of the major contributions of this paper.

VII. CONCLUSIONS

In this paper, we presented the *spatial affordance map* as a model to learn and reason about spatio-temporal patterns of human activities and demonstrated its use to estimate waiting times to activity events and to plan maximum encounter probability paths.

We performed statistical tests to evaluate the validity of the spatial affordance map. To this end, we developed a simulator in which agent behavior is derived from a

questionnaire study with 27 subjects over two weeks. The tests show that the model is a sound statistical model as well as a precise and effective predictor of waiting times.

The approach of a spatial Poisson process enabled us to formulate the encounter probability planning problem as a finite horizon Markov decision process. We compared the proposed MDP planner with two informed and two uninformed heuristic planners where informedness refers to the use of the map. As expected, the MDP planner outperforms all other planners. However, one heuristic planner has a comparable performance to the MDP algorithm. It is greedy but scales better with environment size representing a viable alternative in large-scale applications.

ACKNOWLEDGMENT

The authors would like to thank Malte Helmert for valuable input, Luciano Spinello, Jens Silva and Matthias Frorath for their help with the simulator. This work has been supported by the German Research Foundation (DFG) under contract number SFB/TR-8.

REFERENCES

- [1] N. Roy, G. Gordon, and S. Thrun, “Finding approximate POMDP solutions through belief compression,” *Journal of Artificial Intelligence Research*, vol. 23, pp. 1–40, 2005.
- [2] E. Kruse and F. M. Wahl, “Camera-based observation of obstacle motions to derive statistical data for mobile robot motion planning,” in *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, Leuven, Belgium, 1998.
- [3] A. Bruce and G. Gordon, “Better motion prediction for people-tracking,” in *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, Barcelona, Spain, 2004.
- [4] M. Bennewitz, W. Burgard, G. Cielniak, and S. Thrun, “Learning motion patterns of people for compliant robot motion,” *Int. Journal of Robotics Research*, vol. 24, no. 1, 2005.
- [5] A. T. Ihler and P. J. Smyth, “Learning time-intensity profiles of human activity using non-parametric Bayesian models,” in *Neural Information Processing Systems 19*. Cambridge, MA: MIT Press, 2007, pp. 625–632.
- [6] M. Luber, G. D. Tipaldi, and K. O. Arras, “Place-dependent people tracking,” in *Proc. of the Int. Symposium on Robotics Research (ISRR)*, Lucerne, Switzerland, 2009.
- [7] B. Golden, L. Levy, and R. Vohra, “The orienteering problem,” *Naval Research Logistics*, vol. 34, pp. 307–318, 1987.
- [8] R. P. Bonasso, D. Kortenkamp, D. P. Miller, and M. Slack, “Experiences with an architecture for intelligent, reactive agents,” *Journal of Experimental and Theoretical Artificial Intelligence*, vol. 9, pp. 237–256, 1995.