

Life-Long Spatio-Temporal Exploration of Dynamic Environments

Tomáš Krajník

João M. Santos

Tom Duckett

Abstract— We propose a new idea for life-long mobile robot spatio-temporal exploration of dynamic environments. Our method assumes that the world is subject to constant change, which adds an extra, temporal dimension to the explored space and makes the exploration task a never-ending data-gathering process. To create and maintain a spatio-temporal model of a dynamic environment, the robot has to determine not only where, but also when to perform observations. We address the problem by application of information-theoretic exploration to world representations that model the environment states’ uncertainties as probabilistic functions of time.

We compare the performance of different exploration strategies and temporal models on real-world data gathered over the course of several months and show that combination of dynamic environment representations with information-gain exploration principles allow to create and maintain up-to-date models of constantly changing environments.

Index Terms— mobile robotics, spatio-temporal exploration

I. INTRODUCTION

As robots gradually leave the well-structured worlds of factory assembly lines and enter natural, human-populated environments, new challenges appear. One of the first problems was to operate in less structured and more uncertain environments. This challenge gave birth to the field of probabilistic mapping, which enables the representation of incomplete world knowledge obtained through noisy sensory measurements [1]. Initially, the environment models had to be created during a human-guided procedure, but later, the combination of probabilistic mapping and planning methods allowed the robots to create the environment models by themselves by means of autonomous exploration [2]. However, as robots became gradually able to operate autonomously for longer periods of time, a new challenge appeared – that the natural environments are subject to change.

These changes manifest themselves through sensory measurements – every perceived environment change causes the sensory data to disagree with the original model obtained during the exploration phase. Although the probabilistic mapping methods can deal with the conflicting measurements, their approach is rooted in the idea that these are caused by inherent sensor noise rather than by structural environment change. Thus, these conflicting measurements are generally treated as outliers caused by unwanted noise. This has a negative impact on the ability of the mapping methods to deal with environment dynamics and provide support for long-term mobile robot autonomy.

Lincoln Centre for Autonomous Systems, University of Lincoln, UK
tkrajnik@lincoln.ac.uk

The work has been supported by the EU ICT project 600623 ‘STRANDS’.

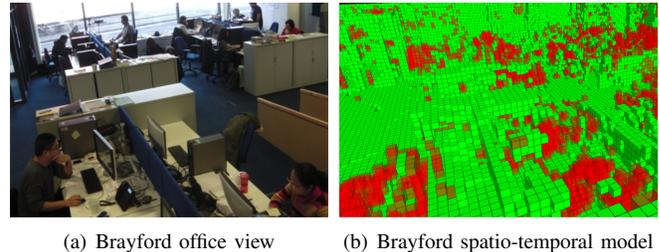


Fig. 1. Spatio-temporal occupancy grid of the Brayford office of the Lincoln Centre for Autonomous Systems. The static cells are in green and cells that exhibit daily periodicity are in red.

Recently, some authors proposed to exploit these conflicting measurements in order to obtain information about the world dynamics and proposed representations that model the environment dynamics explicitly. These dynamic representations have shown their potential by improving mobile robot localization in changing environments [3], [4], [5], [6].

Similarly to traditional robotic mapping, introduction of spatio-temporal mapping naturally requires spatio-temporal exploration. Unlike the classic exploration strategies, where the finite size of the explored space causes the exploration task to be finite, exploration of dynamic environment is never finished. Rather, the spatio-temporal exploration becomes a part of the robot’s daily routine that has to be carried out along with other tasks that the robot is required to perform. A typical spatio-temporal exploration routine would consist of repeated observations of different locations spread throughout the robot’s operational time.

We present a novel exploration method that integrates sensory data captured at different times and locations into a dynamic spatio-temporal model and uses the model to determine where and when to perform future observations. We show that application of information-theoretic planning principles to environment models that represent uncertainties of environment states in the frequency domain results in an intelligent and continuously improving exploratory behaviour, which evolves as the environment knowledge becomes more refined over time. The proposed method allows the mobile robot to create, maintain and refine its environment models as a part of its daily routine, which enables efficient long-term operation in changing environments. To demonstrate the advantages of the approach presented, we perform an experimental evaluation of its performance on two datasets collected over a period of several months.

II. RELATED WORK

In order to explore the environment in an efficient way, the robot has to be able not only to create maps from its sensory inputs, but also to use the map to plan its path so that it can reach previously unknown areas of the environment. Therefore, mobile robot exploration is an iterative process in which the robot integrates its real world observations into its world model, interprets the world model to determine which parts of the environment are unknown, and plans a path to visit and observe these unknown areas. Therefore, an efficient exploration system consists of three essential components: mapping, goal generation and path planning. For the purpose of spatio-temporal exploration, we have to use mapping methods that allow to represent dynamic environments and goal generation methods that can determine not only the position, but also the times of observations. Moreover, the planning has to take into account the time domain – i.e. it has to schedule the observations in such a way that the robot can perform its other tasks as well.

A. Exploration methods

One of the earliest and well-known methods is frontier-based exploration [7]. These approaches [8], [9], represent the environment as an occupancy grid which is processed to obtain boundaries (frontiers) between the known and unknown parts of the environment. The robot movement is then planned so that these frontiers are visited and removed. The advantage of this approach is its scalability – the frontiers can be distributed among a number of robots that can explore the environment in a cooperative manner [10].

Another class of exploration methods is based on the notion of entropy. These methods generate a set of candidate observations and estimate the amount of information these are expected to provide. The information gain is calculated as a reduction of entropy of the environment model, which requires a probabilistic representation of the environment states.

An example of an information-gain-based approach that integrates localization, mapping and exploration is presented in [2]. The method uses a Rao-Blackwellized particle filter to build the map of the environment and an entropy estimation method to plan the next location to be visited by the robot.

However, the candidate observations are not evaluated simply by their information gain. Rather, the evaluation takes into account other criteria, such as the time to reach the position [11]. An advantage of these methods is that they do not only attempt to cover the entire environment as quickly as possible, but also plan re-observations of previously visited locations to increase the quality of the resulting map [12].

B. Dynamic environment representations

As soon as robots had attained the ability to operate for longer time periods, the effects of the environment changes had to be taken into account. The first approaches were aimed at short-term dynamics. These methods identify moving objects and remove them from the environment representations [13], [14] or use them as moving landmarks [15]

for self-localization. However, some dynamic objects do not move at the time of the mapping session and so the robot needs long-term observations to identify them. [16] propose to process several 3d point clouds of the same environment obtained over a period of several weeks to separate movable objects and refine the model of static environment structure at the same time.

Other approaches do not attempt to explicitly identify movable objects, but rely on less abstract environment representations. [17] and [18] represent the environment dynamics by multiple temporal models with different timescales, and [19] use a ranking scheme that allows to identify environmental features that are more likely to be stable in long-term. Churchill and Newman [3] propose to cluster similar observations at the same spatial locations to form ‘experiences’ which are then associated with a given place and show that this approach improves autonomous vehicle localization. The authors of [5] represent the states of the environment components (cells of an occupancy grid) with a hidden Markov model and show that their representation improves the localization robustness as well. Kucner’s method [20] learns conditional probabilities of neighbouring cells of an occupancy grid to model typical motion patterns in dynamic environments. Another team proposed a method that can learn appearance changes based on an across-seasons dataset and use the learned model to predict the environment appearance for a given time [6].

Finally, [21] proposes to represent the environment dynamics in the spectral domain and apply this approach to image features to improve localization [4], to occupancy grids to reduce memory requirements [22] and to topological maps to improve path planning [23] and robotic search [24].

While being applicable to most environment models used in mobile robotics, the aforementioned method suffers from a major drawback that is caused by its reliance on the traditional Fast Fourier Transform (FFT) method, which requires that the environment observations are taken on a regular and frequent basis. This means that the robot’s activity has to be divided into a learning phase, when it would frequently visit individual locations to build its dynamic environment model, and a deployment phase when it would use its model to perform useful tasks. This division means that while the robot can create dynamic models which are more suitable for long-term operation, it cannot maintain them. Thus, the robot does not adapt to dynamics that were not present during the learning phase which leads to deterioration of its efficiency over time.

III. SPATIO-TEMPORAL EXPLORATION

The primary purpose of robotic exploration is to automatically acquire a complete and precise model of the robot’s operational environment. To explore efficiently, the robot has to direct its attention to environment areas that are currently unknown. If the world was static, these areas would simply correspond to previously unvisited locations. In the case of dynamic environments, visiting all locations only once is not enough, because they may change over time. Thus,

dynamic exploration requires that the environment locations are revisited and their (re-)observations are used to update a dynamic environment model. However, revisiting the individual locations with the same frequency and on a regular basis is not efficient because the environment dynamics will, in general, not be homogeneous, (i.e. certain areas change more often and the changes occur only at certain times). Similarly to the static environment exploration, the robot should revisit only the areas whose states are unknown at the time of the planned visits. Thus, the robot has to use its environment model to predict the uncertainty of the individual locations over time and use these predictions to plan observations that improve its knowledge about the world’s dynamics.

To tackle the problem of predicting environment uncertainty over time, we propose to model the probabilities and entropies of the environment states as functions of time. While the main idea lies in the fact that some of the environment’s mid- to long-term dynamics are periodic [21], the underlying mathematical representation had to be reformulated. Unlike the method [21] that requires frequent and regular environment observations, the method proposed in this paper allows to incrementally and continuously update the spatio-temporal model from sparse observations taken at different locations and times. This eliminates the need for a separate training and deployment phase, and allows integration of spatio-temporal exploration into the robot’s daily routine. Thus, the robot can continuously refine its internal environment model and improve its efficiency from the experience gathered over long periods of time.

A. Problem definition

Let us represent the environment as a set \mathcal{S} of n discrete non-stationary independent binary states $s_i(t)$ that are observable by a mobile robot through its sensors. The states $s_i(t)$ might represent the occupancy of individual cells in an occupancy grid, the traversability of edges in a topological map, the visibility of environmental features, etc. Since these states are dynamic and the robot cannot observe all the states all the time, it maintains an internal environment model that we denote as a set \mathcal{S}' , where each element $s'_i(t)$ corresponds to the real-world state $s_i(t)$. To represent the fact that the currently unobserved states are uncertain, we associate each state with a probability value $p_i(t)$ such that $p_i(t) = P(s_i(t) = 1)$. We refer to the probability function $p_i(t)$ and the way it is calculated from the past observations of $s_i(t)$ as a **temporal model**.

Let us define a **location** as a set of environment states that can be observed simultaneously, i.e. a location \mathcal{L}_j is a subset of \mathcal{S} such that by visiting location \mathcal{L}_j at time t , observations of the states that belong to \mathcal{L}_j are obtained. Given that the robot location at time t is $l(t)$, the states of the robot’s internal environment model are

$$s'_i(t) = \begin{cases} s_i(t) & \text{if } s_i \in \mathcal{L}_j \text{ and } l(t) = j \\ p_i(t) \geq 0.5 & \text{otherwise.} \end{cases} \quad (1)$$

The purpose of the exploration process is to obtain and maintain as faithful an environment model as possible, i.e. to

minimize the difference of the states of the real environment \mathcal{S} and its model \mathcal{S}' . Technically, this corresponds to minimization of the model error $\epsilon(T)$ calculated as the difference of real and estimated states over the time period $[0, T)$ as

$$\epsilon(T) = \frac{1}{T} \sum_{t=0}^{T-1} \sum_{i=1}^n |s'_i(t) - s_i(t)|. \quad (2)$$

Although the reduction of the error $\epsilon(T)$ can be partially achieved by visiting the relevant locations as often as possible, the robot has to perform other tasks and the number of observations is typically limited. Thus, the robot has to carefully **plan** where and **when** to perform observations so that it obtains the relevant data to create, maintain and refine its spatio-temporal models of the environment. From a technical point of view, the robot has to use its internal temporal models $p_i(t)$ to determine a sequence of locations $l(t)$. We refer to the way the robot plans the sequence of $l(t)$ from the $p_i(t)$ as its **exploration strategy**.

IV. SPATIO-TEMPORAL MODELS

The underlying spatial environment representations that we will use to test our approach are occupancy grids, topological and feature-based maps. The elementary states of these models represent the occupancy of individual cells, the presence of people at the particular areas and the visibility of image features. Unlike in classic environment models that represent the probabilities of the elementary states $s(t)$ by constant values p , we represent the probability of each elementary state as a function of time $p(t)$. In particular, we model each $p(t)$ as a combination of harmonic functions that correspond to hidden periodic processes in the environment.

A. Spectral maps

The idea of identifying periodic patterns in the measured states and using them for future predictions was originally presented in [21]. These methods process the sequences of the measured state $s(t)$ by the Fast Fourier Transform (FFT) to obtain the corresponding frequency spectrum $s(\omega)$ and extract its most prominent spectral components $s'(\omega)$. Then, they employ the Inverse Fast Fourier Transform (IFFT) method to recover a sequence of state probabilities $p_i(t)$, which can be used for anomaly detection [21] or state prediction [4]. However, the reliance of these methods on the Fast Fourier Transform (FFT) algorithm makes their real-world application impractical. First, the FFT method can transform only the complete sequence of a state $s(t)$ or its full spectral representation $s(\omega)$. Thus, updating the spectral representation with new measurements or prediction of a single probability requires to recalculate the entire sequence of observations, which becomes computationally expensive as the observations accumulate. Most importantly, the FFT algorithm requires that the environment observations are taken on a regular basis which is contrary to the idea of spatio-temporal exploration.

B. Frequency map enhancement (FreMEn)

Similarly to the aforementioned spectral representation, our method still aims to identify the periodic patterns of the environment states and use them for predictions. Unlike the spectral maps, the method proposed here allows to update the underlying dynamic models incrementally from sparse, irregular observations. The proposed method represents each state by a number of performed measurements n , its mean value μ , and two sets \mathcal{A} , \mathcal{B} of complex numbers α_k and β_k that correspond to the set Ω of periodicities ω_k that might be present in the modelled environment. Initially, the mean value μ is set to 0.5 and all α_k, β_k are set to 0, which corresponds to a completely unknown state.

1) *Addition of a new measurement*: Each time a state $s(t)$ is observed at time t , we update its representation, i.e. the number of measurements n , the mean μ and values of \mathcal{A} , \mathcal{B} , which are actually a sparse spectral representation of $s(t)$, as follows:

$$\begin{aligned} \mu &\leftarrow \frac{1}{n+1} (n\mu + s(t)), \\ \alpha_k &\leftarrow \frac{1}{n+1} (n\alpha_k + s(t)e^{-jt\omega_k}) \quad \forall \omega_k \in \Omega, \\ \beta_k &\leftarrow \frac{1}{n+1} (n\beta_k + \mu e^{-jt\omega_k}) \quad \forall \omega_k \in \Omega, \\ n &\leftarrow n + 1. \end{aligned} \quad (3)$$

The proposed update step is analogous to incremental averaging – the absolute values of $|\alpha_k - \beta_k|$ actually correspond to the average influence of a periodic (with a frequency of ω_k) process on the values of $s(t)$. Note that the size of the representation of the state (i.e. the number of elements in \mathcal{A} , \mathcal{B}) is independent of the number of observations, which means that the memory requirements of the proposed representation do not grow over time. Note also that if the times of observations t and frequencies ω_k were equally spaced, i.e. $t = i\Delta_t$ and $\omega_k = i\Delta_\omega$, then (3) corresponds to the traditional Discrete Fourier Transform.

2) *Performing predictions*: To predict the value of state $s(t)$ for a future time t , we first create a set \mathcal{C} consisting of $\gamma_k = \alpha_k - \beta_k$ and then sort it descendingly according to the absolute values $|\gamma_k|$. Then, we extract the first m elements γ_l along with their corresponding frequencies ω_l and calculate the state's probability over time as

$$p(t) = \varsigma\left(\mu + \sum_{l=1}^m |\gamma_l| \cos(\omega_k t + \arg(\gamma_l))\right), \quad (4)$$

where $\varsigma(\cdot)$ ensures that $p(t) \in [0, 1]$. The choice of m determines how many periodic processes are considered for prediction. Setting m too low would mean that we omit to take into account some environment processes that actually influence the state, while setting m too high might include components of \mathcal{C} that are caused by sensor noise. In both cases of m being too low or too high, the prediction of $p(t)$ would lose accuracy. To estimate the optimal value of m , one can compare the predictions performed by (4) to the actually measured values by means of (2), and select the value of m that minimizes the prediction error ϵ .

One of the main advantages of the proposed representation is that the state is modelled probabilistically. This allows to

calculate the time intervals when the particular states are uncertain, which is crucial to direct the robot's attention during the exploration.

C. Alternative temporal models

To evaluate the proposed approach for temporal modelling, we will compare it with three other methods that allow to handle changing environments.

The most popular way to deal with uncertainty of the environment states is based on Bayesian filtering, which updates the environment states based on the sensor noise characteristics. Since the typical measurement rate of the robot sensors exceeds the mid- to long-term environment dynamics we are concerned with, the Bayesian update scheme causes the probabilities of the observed states to quickly converge towards the latest observed values. Thus, the traditional Bayesian filtering tends to reflect the latest state measurements and acts as a **short-term memory** (SM).

Another way to reflect the uncertainty of the observed states in the long-term is to implement a **long-term memory** (LM). Our long-term memory model calculates the probability of a given state simply as an arithmetic mean of all its past observations.

Both of the memory-based models are actually static – the probabilities of the modelled states change only when these are directly observed by the robot. An alternative, representation of the environment dynamics might simply assume that the states exhibit daily periodicity and model the probability of an event at a given time of a day by means of **Gaussian Mixture Models** (GM).

V. EXPLORATION STRATEGIES

As noted in Section III-A, an exploration strategy is defined as a process that determines when and which locations to visit. One has to assume that a real mobile robot has to perform other tasks as well and it can spend only a fraction of the total time on actual exploration. We refer to this fraction as the exploration ratio e , e.g. $e = 0.2$ means that the robot can spend 20% of its operational time on exploration.

Thus, given an exploration ratio e and a set \mathcal{T} of time intervals $[t_s, t_{s+1})$, the exploration algorithm has to determine a sequence $l(t_s)$ of locations to visit. To represent situations where the time slot $[t_s, t_{s+1})$ is allocated to an unrelated activity, the value of $l(t_s)$ is set to zero, whereas a non-zero value of $l(t_s)$ signifies the location to be observed during $[t_s, t_{s+1})$.

A. Information-gain strategies

The information-gain strategies take into account the experiences the robot has gathered so far to plan when and which location to visit. These strategies attempt to reduce the uncertainty of the environment models by planning the observations that maximize the potential information gain. To estimate how much information is gained by a particular observation, we will use the notion of entropy. We assume that direct observation of particular states at a given time

reduces the entropy of these states to zero. Thus, the information gained by a particular observation can be estimated as the sum of the entropies of the states observed at a given location as

$$I(\mathcal{L}, t) = - \sum_{i \in \mathcal{L}} (p_i(t) \log(p_i(t)) + (1 - p_i(t)) \log(1 - p_i(t))). \quad (5)$$

The **Greedy** strategy calculates the potential information gains for all given time slots and locations, then assigns the best location to visit at each time slot. Then, it selects a subset \mathcal{T}' of time slots with the highest information gain such that $e = |\mathcal{T}'|/|\mathcal{T}|$. The remaining time slots are assigned for exploration-unrelated tasks. Thus, this strategy maximizes the potential information gain obtained over the time slots in the set \mathcal{T} .

The **Monte Carlo** strategy chooses the locations randomly, but the probability of selecting a given location at a given time is proportional to the estimated information gain. At first, it estimates the $I(l, t_s)$ for all given time slots and locations and sums these values to I' . Then, it calculates the value of $I(0, t_s) = I'(1 - e)/(ne)$. Finally, it calculates the probabilities of each $l(t_s)$ as

$$P(l(t_s) = j) = \frac{I(j, t_s) + \iota}{\sum_{i \in \mathcal{L}} I(i, t_s) + \iota}. \quad (6)$$

Here, the value of $I(0, t_s)$ does not represent actual information gain, but is added to ensure that the exploration ratio e is satisfied by ensuring sufficient chance of assigning the time slots to exploration-unrelated tasks. The positive constant ι ensures that the locations will be occasionally visited even at times when the spatio-temporal model predicts their state with absolute certainty. This allows to detect unexpected changes of environment dynamics.

B. Uninformed strategies

For comparison purposes, we include strategies which select the places to visit regardless of the environment dynamics. These calculate the sequence of visits $l(t_s)$ simply from the values of the ratio e , number of locations n and number of time slots m .

The **Round-Robin** strategy visits all areas of the environment with the same frequency, interleaving the observations with other tasks so that the exploration ratio e is satisfied.

The **Random** strategy also attempts to visit all areas with the same frequency, but the sequence of $l(t_s)$ is not deterministic, but random. The probability of a given slot being assigned to a non-exploration task is equal to $1 - e$ and the probability of visiting the individual locations is uniform and equal to e/n .

VI. EVALUATION DATASETS

To evaluate the ability of the various temporal models and exploration strategies, we performed a comparison on two datasets gathered over several weeks. The first, ‘Aruba’ dataset was gathered by a team of the Center for Advanced Studies in Adaptive Systems (CASAS) to support their research concerning smart environments [25]. The second,

‘Brayford’ dataset was created at the Lincoln Centre for Autonomous System Research (LCAS) for their research on long-term mobile robot autonomy [4]. The aforementioned datasets were processed so that the dynamics of these environments are represented as visual-feature-based, topological and metric maps.

A. The Aruba dataset

The ‘Aruba’ dataset consists of maps capturing 16 week long dynamics of a large apartment that was occupied by a single, house-bound person who occasionally received visitors. An occupancy grid and a topological map were created for every minute of a 16 week long period – the resulting dataset contains over 160 000 metric and topological maps. Since the original dataset [25] is simply a year-long collection of measurements from 50 different sensors spread over an eight-room apartment, these maps had to be created by means of simulation.

First, we processed the events from the original dataset’s motion detectors to establish the location of the people in the flat for every minute of the 16 weeks. Then, we partitioned the flat into ten different areas, where eight areas represent the rooms and two correspond to corridors. This allowed us to create a topological map that indicates the presence of people in these locations. To obtain the metric representation,



Fig. 2. Aruba environment simulation.

we created a simulated environment with the same structure as the ‘CASAS’ apartment, see Figure 2. Then the simulation was provided with a sequence of person locations recovered in the previous step. As a result, the simulated environment contains physical models of people at locations provided by the real-world dataset, and thus it reflects the dynamics of the real apartment. A virtual, RGB-D camera equipped robot was also introduced into the virtual environment. Every time the configuration of the simulated environment (i.e. locations of the people) changed, the robot used its 3D sensors to create occupancy grids of the flat’s individual rooms. Thus, we obtained occupancy grids that reflect the real environment dynamics minute-by-minute for 16 weeks.

B. The Brayford dataset

The Brayford dataset was originally collected for the purpose of benchmarking long-term mobile robot localization algorithms in dynamic environments [4]. This time, the data

collection was performed by a humanoid-like robot equipped with an RGB-D camera in a large, open-space office of the Lincoln Centre for Autonomous Systems. The robot was set-up to obtain RGB-D images of eight designated areas every 10 minutes for a period of one week. Representative examples of the captured images are shown in Figures 3 and 1. While the high-level environment model of this dataset



Fig. 3. Examples of Brayford dataset images.

contains information about people presence at the individual locations, the states of the low-level model represent the visibilities of visual features [26]. The resulting dataset contains more than 8000 feature-based and 8000 semantic maps collected over a period of one week.

VII. EXPERIMENTAL RESULTS

We assume that the aforementioned datasets reflect the real state of the environments they have been captured in and thus we use the sequence of the observations in the datasets as ground truth. To evaluate how the various temporal models and different exploration strategies affect the robot’s ability to create and update its internal environment models, we emulate the exploration process using the datasets gathered. We assume that the exploration can be performed only half of the robot’s operational time (i.e. $e = 0.5$) and that a single observation takes 10 minutes including the time to move between locations.

This exploration procedure corresponds to the situation when the robot updates its spatio-temporal model and generates a new observation schedule every midnight. The robot starts with an empty environment model that has all probabilities constant and equal to 0.5.

First, the entropy functions of the individual locations are calculated and 72 observations for the following day are scheduled. Then, these 72 observations are retrieved from the given dataset and the temporal models of the environment states are updated. The updated temporal models are used to recalculate the spatio-temporal entropy and the next day’s observation schedule is generated. These steps are repeated for every day of the given dataset.

A. Evaluating environment model error

To compare the performance of the temporal models and exploration strategies described in Sections IV and V, the resulting world model is compared to the actual dataset by Equation 2, which estimates the environment model error. Since there are 4 temporal models and 4 exploration strategies, each comparison consists of 16 numbers that characterize the ratio of incorrectly estimated states to the

total number of environment states. One dataset evaluation consist of two comparisons, each corresponding to the given environment representation. The results of the ‘Aruba’

TABLE I

THE ARUBA DATASET RESULTS: MODEL ERRORS ([%]) FOR DIFFERENT EXPLORATION STRATEGIES AND TEMPORAL MODELS

Strategy	Spatio-Temporal model							
	People presence				Occupancy grids			
	Static		Dynamic		Static		Dynamic	
	SM	LM	FT	GM	SM	LM	FT	GM
Round-Robin	9.9	9.7	6.5	7.5	11.6	11.0	7.4	8.8
Random	8.0	9.5	9.2	7.5	9.0	10.2	10.8	9.1
Greedy	9.8	8.7	7.0	9.4	21.0	13.0	8.8	8.3
Monte Carlo	9.9	8.9	5.8	6.4	11.1	10.2	6.3	8.7

dataset summarized in Table I show that the exploration method which combines the Frequency Map Enhancement and Monte Carlo exploration strategy reduces the model error by more than 40%. Since more than 99% of the cells in the ‘Aruba’ occupancy grids represent empty space or static objects, the model error (2) is calculated for the cells that change their occupancy at least once.

TABLE II

THE BRAYFORD DATASET RESULTS: MODEL ERRORS ([%]) FOR DIFFERENT EXPLORATION STRATEGIES AND TEMPORAL MODELS.

Strategy	Spatio-Temporal model							
	People presence				Visual features			
	Static		Dynamic		Static		Dynamic	
	SM	LM	FT	GM	SM	LM	FT	GM
Round-Robin	20.3	23.7	17.9	21.1	19.0	23.8	9.9	20.8
Random	19.3	23.8	23.4	23.5	13.4	24.0	23.7	21.9
Greedy	21.7	22.5	21.1	20.4	21.8	24.0	15.8	19.5
Monte Carlo	21.3	23.5	16.7	19.6	19.0	23.8	9.2	16.4

The model errors of the ‘Brayford’ dataset as shown in Table II again indicate that the most faithful environment representation is based on Frequency-enhanced temporal models (see Section IV-B), which are obtained through Monte Carlo exploration. The improvement is more prominent in the case of visual feature-based maps. The reason for this might be caused by the fact that the visual features’ visibility tends to follow more regular patterns than the working habits of the office researchers.

B. Exploration vs. Exploitation

In the above experiments, the robot exploration ratio e was set to 0.5. Thus, the robot could spend 50% of its time gathering data about its operational environment. However, such a ratio is unrealistic – the robot has to spend some time replenishing its batteries and we have to assume that it has to perform other tasks as well. Moreover, we have to assume that the purpose of the robot is not in creating precise environment models, but to perform useful tasks. Thus, the exploration task it just an instrument to obtain and maintain knowledge which improves the robot’s performance. If the robot would spend too much time on exploration, it would

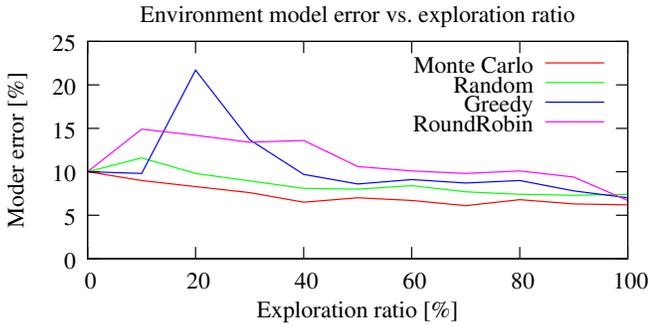


Fig. 4. Exploration vs. exploitation analysis: The influence of the fraction of time spend with exploration on the performance of the individual exploration strategies.

not be able to exploit the obtained knowledge in its everyday activities.

Thus, the *exploration versus exploitation dilemma* means that the robot has to find a balance between the time spent exploring and the quality of its internal model. We evaluate the efficiency of the individual exploration strategies with different exploration ratios on the Aruba dataset semantic map. We combine the Frequency Map Enhancement models with four different exploration strategies, fix the exploration ratio to a value between 0 and 1, and let the robot explore the Aruba environment for two consecutive weeks. The resulting error of the model obtained is shown in Figure 4. The results indicate that if the fraction of the time that the robot can spend on actual exploration is low, the dynamic models might make wrong assumptions about the environment changes and perform worse than their static counterparts – this is especially notable with the Greedy and RoundRobin strategies. However, this effect can be mitigated by a proper exploration strategy – the graph shows that Monte Carlo strategy improves the model even if the robot cannot spend too much time on exploration.

Note, that the initial model error is 10% – this is caused by the fact that the semantic map of the Aruba flat represents the presence of people in 10 different areas and the flat has only one inhabitant. Without any observations, the robot simply assumes that the flat is empty, which results in 10% error.

C. Qualitative evaluation

To allow an insight into the robot’s actual exploratory behaviour, we interpret the data gathered during the exploration of the ‘Aruba’ topological map. Here, the robot’s task was to create spatio-temporal model of person presence in the individual rooms of a small apartment. For the purpose of this explanation, let us focus on the dynamics of the three rooms only – the bedroom, the kitchen and a storage room. Let our robot use the best-performing exploration method that combines the FreMEn temporal models and the Monte Carlo exploration strategy. Applying the proposed spatio-temporal exploration method to this dataset produced the behaviour in Figure 5. The top part of Figure 5 shows the real state of the environment – the three binary functions $s_i(t)$ represent the room’s occupancies over time. The second part shows the robot’s internal model of the environment,

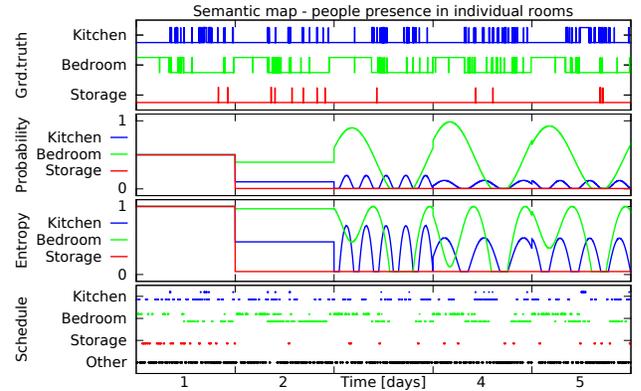


Fig. 5. Spatio-temporal exploration behaviour: The robot uses its probabilistic world model (second row) and spatio-temporal entropy estimates (third row) to schedule its observations (bottom graph) and learn the environment dynamics (top). As the environment knowledge improves over time, the scheduled observations provide more information which allows for further refinement of the environment model.

i.e. the probabilities $p_i(t)$. The third graph displays the information that is expected to be obtained by visiting these three locations at a given time. Finally, the bottom graph shows which locations have been visited at a particular time – we assume that the exploration ratio $e = 0.5$, which reflects the situation where the robot has to spend half of its time on its charging station. Now let us explain how the robot’s understanding of the environment changes over time and how this affects its exploratory behaviour day by day.

1) *Day one:* Initially, the robot has no knowledge of the environment and therefore the probabilities $p_i(t)$ of the world states $s(t)$ are equal to 0.5. This means that the expected information gain from visiting any of the three rooms equals 1 bit at any time of the first day. Thus, the robot has no room or time preference when scheduling the first day’s observations.

2) *Day two:* After performing the first day’s observations, the environment models provide enough evidence that the three rooms are not occupied with the same probability. This is reflected in the second day’s environment model, see the probability functions $p_i(t)$ of the second day in Figure 5. Thus the robot expects to gain more information by visiting the bedroom and kitchen than by going to the storage room. This is reflected by the second day’s observation schedule – the last row of Figure 5 shows that the first two rooms are visited more often.

3) *Day three:* The additional observations obtained during the second day provide information about the rooms’ dynamics: the robot assumes that the bedroom has a daily periodicity that the kitchen is visited five times per day. This causes the expected information gain to be time-dependent – the third day of the third row of Figure 5 shows that evening and morning observations of the bedroom will provide more information than in the afternoon. This fact is rather intuitive: visiting the room at the time of its state transition allows to refine the room’s state periodicity. Thus, on the third day, the bedroom is visited mostly in the evening and morning, while the afternoon visits are scheduled to the kitchen.

4) *Days four and five*: Based on the data gathered during the third day, the robot modifies its hypothesis about the periodicity of activities in the kitchen and assumes that it is visited three times per day. During the following days, the robot tends to visit the kitchen and bedroom, and checks the storage room only occasionally. While the kitchen is visited mostly in the early afternoon, the bedroom is visited late evenings and mornings, which allows to refine the robot's model of the person's daily habits.

This example indicates that the combination of a probabilistic temporal model with an information-based strategy not only allows the robot to obtain knowledge about the environment dynamics, but the observations are scheduled in a seemingly logical way: At first, all the locations are visited often and with the same frequency. As the spatio-temporal environment model become more refined, the robot tends to visit the particular locations only at times when their states are uncertain.

VIII. CONCLUSION

In this paper, we presented a method for life-long spatio-temporal exploration of dynamic environments. We assume that the robot's operational environment is undergoing constant change, which requires a method that can model and predict these variations. The purpose of spatio-temporal exploration is not only to obtain the environment structure and keep it up-to-date with any changes, but also to allow the robot to observe and understand the world dynamics.

We hypothesize that the problem of spatio-temporal exploration can be tackled by combining information-gain-based exploration strategies with probabilistic dynamic environment models. To verify our approach, we compare the performance of four exploration strategies and four temporal models on real-world data gathered over the course of several months. We show that the combination of spectral-based temporal models with information-gain-based Monte Carlo planning results in an intelligent exploration behaviour that improves as the environment knowledge becomes more refined.

Analysis of the robot behaviour shows that when introduced to a new environment, the robot prefers to explore unknown locations. After it has obtained the spatial models, it starts to revisit these locations in order to learn about their dynamics. Finally, the learned dynamics allow the robot to determine which locations to visit at which times.

The evaluations performed in this paper involved several assumptions to simplify the problem. The first assumption was that the time the robot spends moving to a particular location is negligible compared to the time it takes to make an observation. The second assumption was that the observations locations were predefined and fixed in time. While these assumptions were needed for validation purposes in this work due to the known difficulties of ground-truthing when comparing exploration strategies, our future work will involve relaxing these assumptions to achieve a more general and more autonomous spatio-temporal exploration method.

REFERENCES

- [1] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics (Intelligent Robotics and Autonomous Agents)*. The MIT Press, 2005.
- [2] C. Stachniss, G. Grisetti, and W. Burgard, "Information gain-based exploration using Rao-Blackwellized particle filters," in *Proc. of Robotics: Science and Systems (RSS)*, Cambridge, MA, USA, 2005.
- [3] W. S. Churchill and P. Newman, "Experience-based navigation for long-term localisation," *The Int. Journal of Robotics Research*, 2013.
- [4] T. Krajník, J. P. Fentanes, O. M. Mozos, T. Duckett, J. Ekekrantz, and M. Hanheide, "Long-term topological localization for service robots in dynamic environments using spectral maps," in *Proc. of Int. Conference on Intelligent Robots and Systems (IROS)*, 2014.
- [5] G. D. Tipaldi *et al.*, "Lifelong localization in changing environments," *The International Journal of Robotics Research*, 2013.
- [6] P. Neubert, N. Sünderhauf, and P. Protzel, "Superpixel-based appearance change prediction for long-term navigation across seasons," *Robotics and Autonomous Systems*, no. 0, 2014.
- [7] B. Yamauchi, "A frontier-based approach for autonomous exploration," in *Proc. of the IEEE Int. Symposium on Computational Intelligence in Robotics and Automation*, 1997.
- [8] S. Koenig, C. Tovey, and W. Halliburton, "Greedy mapping of terrain," in *Proc. of Int. Conference on Robotics and Automation (ICRA)*, 2001.
- [9] D. Holz, N. Basilico, F. Amigoni, and S. Behnke, "Evaluating the efficiency of frontier-based exploration strategies," *ISR/ROBOTIK 2010*.
- [10] B. Yamauchi, "Frontier-based exploration using multiple robots," in *Proc. of the 2nd Int. Conf. on Autonomous agents*, 1998.
- [11] C. Stachniss and W. Burgard, "Exploring unknown environments with mobile robots using coverage maps," in *Proceedings of the International Conference on Artificial Intelligence (IJCAI)*, 2003.
- [12] J. Fentanes, R. F. Alonso, E. Zalama, and J. G. García-Bermejo, "A new method for efficient three-dimensional reconstruction of outdoor environments using mobile robots," *Journal of Field Robotics*, 2011.
- [13] D. Hähnel, D. Schulz, and W. Burgard, "Mobile robot mapping in populated environments," *Advanced Robotics*, 2003.
- [14] D. Wolf and G. Sukhatme, "Mobile robot simultaneous localization and mapping in dynamic environments," *Autonomous Robots*, 2005.
- [15] C. C. Wang *et al.*, "Simultaneous localization, mapping and moving object tracking," *International Journal of Robotics Research*, 2007.
- [16] R. Ambrus, N. Bore, J. Folkesson, and P. Jensfelt, "Meta-rooms: Building and maintaining long term spatial models in a dynamic world," in *Proceedings of the International Conference on Intelligent Robots and Systems (IROS)*, 2014.
- [17] P. Biber and T. Duckett, "Dynamic maps for long-term operation of mobile service robots," in *Proc. of Robotics: Science and Systems*, 2005, pp. 17–24.
- [18] D. Arbuckle, A. Howard, and M. Mataric, "Temporal occupancy grids: a method for classifying the spatio-temporal properties of the environment," in *Proc. of Int. Conference on Intelligent Robots and Systems (IROS)*, vol. 1, 2002, pp. 409–414 vol.1.
- [19] F. Dayoub and T. Duckett, "An adaptive appearance-based map for long-term topological localization of mobile robots," in *Proc. of Int. Conference on Intelligent Robots and Systems (IROS)*, 2008.
- [20] T. Kucner, J. Saarinen, M. Magnusson, and A. J. Lilienthal, "Conditional transition maps: Learning motion patterns in dynamic environments," in *Proc. of Int. Conf. on Intelligent Robots and Systems (IROS)*, 2013.
- [21] T. Krajník, J. P. Fentanes, G. Cielniak, C. Dondrup, and T. Duckett, "Spectral analysis for long-term robotic mapping," in *Proc. of Int. Conference on Robotics and Automation (ICRA)*, 2014.
- [22] T. Krajník, J. Santos, B. Seemann, and T. Duckett, "Froctomap: An efficient spatio-temporal environment representation," in *Advances in Autonomous Robotics Systems*. Springer, 2014, pp. 281–282.
- [23] J. Pulido Fentanes, B. Lacerda, T. Krajník, N. Hawes, and M. Hanheide, "Now or later? predicting and maximising success of navigation actions from long-term experience," in *International Conference on Robotics and Automation (ICRA)*, 2015.
- [24] T. Krajník, M. Kulich, L. Mudrová, R. Ambrus, and T. Duckett, "Where's waldo at time t? using spatio-temporal models for mobile robot search," in *Int. Conf. on Robotics and Automation (ICRA)*, 2015.
- [25] D. J. Cook, "Learning setting-generalized activity models for smart spaces," *IEEE Intelligent Systems*, no. 99, p. 1, 2010.
- [26] M. Calonder, V. Lepetit, C. Strecha, and P. Fua, "Brief: Binary robust independent elementary features," in *Proc. of European Conference on Computer Vision (ECCV)*. Springer, 2010, pp. 778–792.