Now or later? Predicting and Maximising Success of Navigation Actions from Long-Term Experience

Jaime Pulido Fentanes[1] Bruno Lacerda[2] Tomáš Krajník[1] Nick Hawes[2] Marc Hanheide[1]

***Abstract*— In planning for deliberation or navigation in real-world robotic systems, one of the big challenges is to cope with change. It lies in the nature of planning that it has to make assumptions about the future state of the world, and the robot's chances of successively accomplishing actions in this future. Hence, a robot's plan can only be as good as its predictions about the world. In this paper, we present a novel approach to specifically represent changes that stem from *periodic* events in the environment (e.g. a door being opened or closed), which impact on the success probability of planned actions. We show that our approach to model the probability of action success as a set of superimposed periodic processes allows the robot to predict action outcomes in a long-term data obtained in two real-life offices better than a static model. We furthermore discuss and showcase how this knowledge gathered can be successfully employed in a probabilistic planning framework to devise better navigation plans. The key contributions of this paper are (i) the formation of the spectral model of action outcomes from non-uniform sampling, the (ii) analysis of its predictive power using two long-term datasets, and (iii) the application of the predicted outcomes in an MDP-based planning framework.**

***Index Terms*— long term, topological map, path planning, mobile robotics, spatio-temporal representations**

I. Introduction

The performance of motion planning for a mobile robot is highly dependant on the quality of the knowledge about the environment, including the probability of successfully performing certain navigation actions. Traditionally, this knowledge has been considered as static (after an initial acquisition) as part of the robot's domain knowledge. However, for mobile robots in dynamic environments these conditions can vary regularly as a result of changes to the environments. Noticeable examples include doors being shut, stopping a robot from reaching a previously accessible place, or a crowd of people causing it to take longer to travel between two places temporarily. In a real world, some of these changes might be unpredictable and come as a surprise so the robot has to deal with it reactively (e.g. in [1]).

Complementary to such a reactive replanning, the contribution in this paper is to represent reoccurring *patterns* of environmental changes, gathered from *experience* by the robot itself, and using this representation for path avoiding or predicting situations that can block the robot. This representation shall enable a robot to exhibit anticipatory behaviour, such as avoiding certain routes at specific times because they are *expected* to be blocked.

[1]LCAS, University of Lincoln, UK {jpulidofentanes, tkrajnik, mhanheide}@lincoln.ac.uk

[2]School of Computer Science, University of Birmingham, UK {b.lacerda, n.a.hawes}@cs.bham.ac.uk

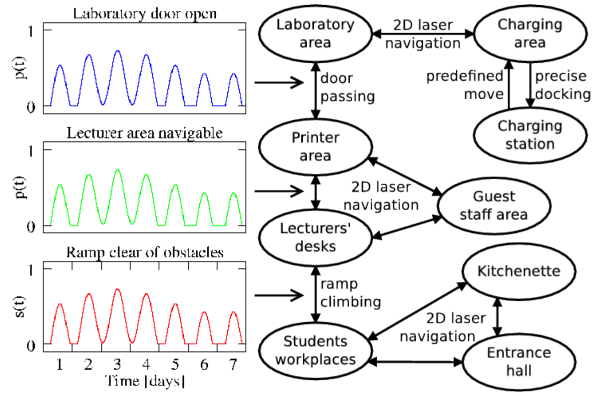

Fig. 1. Example topological map with temporal edge traversability models. Nodes are places in a map with edges linking them by movement actions. These actions can fail. i.e. `door passing` requires an open door, and will vary in duration depending on the environments state. The plots are illustrative, showing the predicted probability for given times $p(s|t)$.

Based on the hypothesis that a significant amount of the changes in indoor environments are actually following certain *routines*, creating mostly *periodic* patterns of change. We propose to represent these dynamics by augmenting a topological map (see Fig. 1), composed of nodes localised in a 2D map and connecting edges with a *spectral model* of their traversability and action execution times.

The spectral model allows to predict future states and their probability at specific times. The experiments presented show that relevant states can be predicted accurately from real-world data in office environments, proving our working hypothesis to be valid for this case. Based on these findings it will be furthermore showcased how such predictions can be used to perform robust motion planning by employing Markov decision process models of the environment. We show the robot's ability to achieve anticipatory behaviour in terms of routes taken in order to satisfy tasks specified in linear temporal logic. The evaluation is done using a mobile robot running autonomously in two different environments gathering data for a total of **10** weeks.

The main contributions in this paper are (i) the novel spectral analysis of world states gathered from robot navigation experience and the augmentation of a topological map; (ii) the study of the gathered long-term data to investigate the validity of the assumption that changes are indeed to a large extend periodic and predictions can be made; (iii) the extension of the method proposed in [2] to handle non uniform sampling and to automatically choose the best fitting model order; and (iv) the discussion and evaluation of employing this representation in high-level motion planning.

II. RELATED WORK

Reliable robot behaviour in real worlds can only be achieved by modelling the relevant domain precisely, and by employing planning algorithms that can handle the inherent uncertainty of the environments. Furthermore, in dynamic scenarios uncertainty grows with time, since the environment can be expected to be constantly changing. Some authors try to handle these changes by removing dynamic objects from the representation of the environment [3], or by tracking these objects and classifying them as dynamic landmarks [4].

Other approaches never assume the map to be complete and perform continuous mapping, adding new features to the map with every observation [5]. In general, such approaches can handle some problems of short-time mapping, but they cannot deal with long-term changes, because they are highly dependant on the persistence of the chosen features.

To tackle the long-term challenge, some approaches gather and maintain different temporal representations simultaneously and choose the best one according to its consistency with current perception of the world (e.g. [6], [7]). However, they are costly in terms of memory and computational effectiveness.

Considering this, many authors have tried these approaches at a discrete, topological, level where computational requirements are lower. At this level, most representations use visual appearance for place recognition [8], [9], having shown that visual features can be used for robust place identification. However, these present a decrease in robustness when facing long-term changes [10], as again they are prone to error when features disappear over long times.

All these methods tackle uncertainty at the node level of a topological map, namely to recognise where the robot actually is. But they do not model how the transitions between nodes are affected by the dynamics of the environment. For this reason [11] proposes a new representation that models occupancy grid maps in the wavelet space in order to optimize the amount of information that has to be processed for path planning, and [12] presents a representation of the environment which models transitions of dynamic objects in the environment, by learning motion patterns from a temporal signal of occupancy on cell. These approaches are closer related to the work presented in this paper, as they employ specific temporal models. Using spectral models to model planning domain knowledge is – to the best of our knowledge – novel on our work. However, the aim to learn situation-dependent costs to improve planning in dynamic domains has been studied for some time, e.g. in [13], who already looked at costs being dependent on the time of the day, but neither learned this from long-term experience nor where able to exploit the periodic nature explicitly. Since this early work, other researchers have incorporated probabilities into their planning domain with [14], [1], [15] being only three representatives of a group of works that employed (partially learned) probabilities to cope with the uncertainty of the world in robot planning.

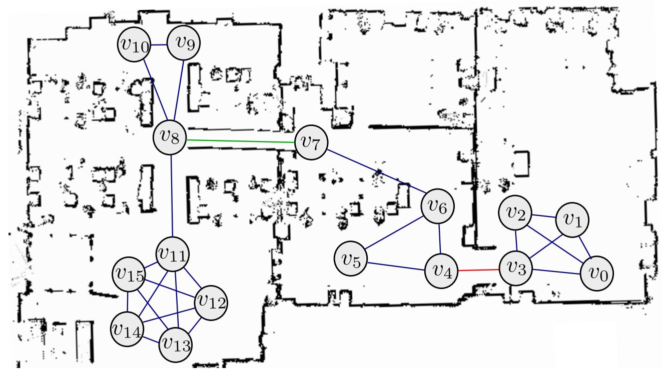

Fig. 2. A topological map. Edges in blue represent the default laser based navigation, the edge in green represents a specialized ramp traversal behaviour, and the edge in red represents a specialized door crossing behaviour. All edges are bi-directional.

III. SPATIO-TEMPORAL TOPOLOGICAL REPRESENTATION

A. Topological map representation

Topological representations typically consist of a set of nodes that represent physical locations of the environment and a set of edges representing the robot's ability to move between these locations. While the nodes are usually associated with descriptions of the locations (e.g. Way point), the edges usually represent a traversable connections between two of these places. In the proposed representation, each edge is associated with an particular continuous navigation action (e.g, laser based navigation, door crossing, taking a lift) that moves the robots between the associated nodes. This representation allows one to implement a unified system that can tackle different challenges in terms of continuous navigation. This fact will be further illustrated on the description of our application scenarios.

Thus, we define our topological map as a tuple $T = \langle V, E, N, nav, P_E \rangle$, where: (i) $V = \{v_1, ..., v_n\}$ is a set of possible robot poses in the environment; (ii) $E \subseteq V \times V$ represents possible continuous navigation actions between different robot poses; (iii) N is a set of possible continuous navigation actions that the robot can execute; (iv) $nav : E \rightarrow N$ is a function that maps each edge $e = (v_i, v_j)$ to the continuous navigation action to be executed to drive the robot from v_i to v_j; (v) $P_E = \cup_{e \in E}\{p_e\}$ is a set of probability functions $p_e : \mathbb{R}_+ \rightarrow \{0, 1\}$ that, for each edge e, map a given time t to the probability $p_e(t)$ of it being traversable.

The construction of the temporal edge models P_E will be described in the next subsection. In Figure 2, we show a depiction of a topological map, for the lab of the Lincoln Centre for Autonomous Systems. We will discuss the deployment in this environment in Section IV.

B. Frequency map enhancement

In this section, we show how to use the *Frequency Map Enhancement (FreMEn)* approach presented in [2], to build a traversability probability distribution p_e for each edge $e \in E$.

The temporal models for each edge are created under the assumption that the traversability of an edge can assume two distinct states (successful or failed) and is affected by a series of processes caused by human activity which might be periodical. The FreMEn approach is based on the Fourier transform [16] and states that if the influence and periodicity of any process is identifiable, it is possible to estimate its future state from its description.

Typically, the uncertainty of these states would be represented by means of probability value that is updated each time the robot measures the state of a particular environment component. FreMEn considers probability of each environment state or or particular action outcome being a function of time $p_j(t)$ and represents this function by its frequency spectrum.

The spectrum of each state can be obtained through its sparse, but long-term observations. The main advantage of the approach is that it can efficiently calculate the probabilities of the environment states for a given time in the future. This augments the robotic maps with predictive capabilities which allows for a better planning. Another advantage of the FreMEn is that since it is a mathematical tool for signal analysis, it can be applied to most environment models used in mobile robotics. Introducing dynamics to volumetric [17], landmark [18] and topological [19] environment models has shown to improve the efficiency of mobile robots that operate autonomously for long periods of time. In this work, we apply FreMEn to topological maps - in particular, we use it to model the probability that the robot will successfully traverse a particular map edge. Although the essential idea - i.e. representing the periodic changes of the environment by means of frequency spectrum - is the same, we had to deal with the fact that the robot cannot gather the information on a regular basis. For this reason a non-uniform sampling extension was developed specifically for this case, also since in this case it was impossible to manually choose the best model for each edge individually a method for automatically finding the best model order had to be introduced.

So let us assume that the outcome of the robot attempt to traverse a particular edge at time t can result either in a success or a failure. The outcome of the action is likely to be influenced by the state of the environment corresponding to the particular edge. Therefore, we consider that an edge observation can assume values of 0 (blocked) or 1 (traversable), depending on the state of the corresponding navigation action. We denote the observation of edge e at time t as $s_e(t) \in \{0, 1\}$. Thus, the probability $p_e(t)$ can be given by the model prediction for $s_e(t) = 1$. As stated before, we assume that the edges' traversabilities can be influenced by the routines of people within the work environment of the robot. This allows us to represent the probability of success when traversing the edge (traversability) of a particular edge by a combination of periodic functions

$$p_e(t) = \alpha_0 + \sum_{j=1}^{m} \alpha_j cos(\omega_j t + \varphi_j) \tag{1}$$

where $\alpha_j, \omega_j, \varphi_j$ correspond to influence, periodicity and time offset of the processes that affect $s_e(t)$. These parameters can be identified by a well-known tool used in signal processing, the Fourier Transform (FT). To ensure that the probability $p_e(t)$ is between 0 and 1, we can simply truncate the output of Eq.(1) to $[0, 1]$. The standard FT assumes uniformity on the observations of the states of the system. However, in our scenario, it is not possible to guarantee that observations are performed on a regular basis. Thus, one needs to apply a non-uniform version of the FT, described below.

Assume that the robot has attempted to traverse edge e n times, and let us denote the times of these attempts as t_i and the observed edge traversability as $s_e(t_i)$. The first spectral component α_0 is simply an arithmetic mean of $s_e(t)$. Let us denote a set of frequencies corresponding to the periodicities of the processes that might influence the edge's traversability as $\omega_k \in \Omega$. The influence α_k of these processes can be calculated by

$$\alpha_k = |\sum_{i=1}^{n} (s(t_i) - \alpha_0) e^{-j2\pi t_i \omega_k}|. \tag{2}$$

Thus, to build a spectral model of a state $s(t)$, one has to calculate Eq. (2) for all relevant $\omega_k \in \Omega$, and select m frequencies with highest influence α_k and calculate the φ_l in a similar fashion as α_k, where l is the index of the subset of α_k. Thus, one retains m triplets consisting of influence α_l, frequency ω_l and phase φ_l that are parameters of the spectral model (1). Previously the number of modelled processes m had to be chosen by hand. In our case, the parameter m is chosen to minimize a reconstruction error $e_r(m)$ defined as

$$e_r(m) = \sum_{i=1}^{n} |s(t_i) - \alpha_0 - \sum_{l=1}^{m} \alpha_l cos(\omega_l t_i + \varphi_l)|. \tag{3}$$

Typically, $m = 2$ which means that the spectral model consists of 7 numbers, representing two periodic and one static process.

It is important to mention that even though the method presented in [2] was originally presented for binary states equations (1,2,3) can be used even when the values of s(t) are real numbers. In later sections we will show how we can take advantage of this situation for time analysis and we will show how times are also affected by environment dynamics and that the predicted times obtained by this model correspond to real world data obtained during the experiments.

IV. Periodicity and Prediction Analysis

In order to validate the presented approach, we deployed mobile robots in two populated office-like environments. In each one of these environments a topological map was created, and different tasks had to be executed at different nodes of the map. The statistics for traversability and navigation times on these maps were recorded and used to create the dynamical models for the edges in these environments. Our experimental setup consists of two SCITOS-G5 Robots (Figure 3) equipped with RGB-D and laser sensors running the STRANDS system[1]. The robots were deployed in an

[1]The software employed to generate the results for this paper is open-source and available from https://github.com/strands-project

Fig. 3. The robot Linda at Open space office and Bob at office corridor

open plan office and on an office corridor for several weeks (see Table IV) using different patrol configurations and during this deployment this environments kept their usual working routines. The open plan office is a two level area

TABLE I
MAP SIZE AND DURATION OF ROBOT ACTIVITY AND DEPLOYMENT

Site	Duration (in days) of deployment	activity	Number of map nodes	edges
Open plan	76	35	17	62
Corridor	24	17	45	139

with a laboratory, as depicted in Figure 2. Every ten minutes a patrol run was scheduled on which the robot undocked from its charging station (node v_0) and had to go to several nodes around the office to gather data to return to its charging station after each patrol. This data gathering process was repeated 24 hours a day for several days. Because of lab regulations the lab door was closed when no one was at the office so the robot could not access the office area when no one was there.

The office corridor is an office space consisting on a long corridor with offices and meetings rooms along it. In this case the robot was set to execute a specific routine during the weekdays which meant that the robot was crossing the edges in a more periodical but less frequent way.

A. Reconstruction Analysis

To validate this methodology the first step was to build a model using the whole set of data and then comparing the real world observations against the best (by means of criterion (3)) dynamic model. In this case we obtained the mean reconstruction error for success rate and traversal time by averaging the $e_r(m)$ (see eq.3) for all edges. The errors are summarized in Table IV-A, which also shows how many edges use dynamic models with a given number of periodic processes. Note that the reconstruction error for the times is in seconds and represents the mean deviation between the predicted times and the real navigation times on all edges.

These results show that effectively the algorithm tries to fit the reconstruction to the dynamics of the outcome of the navigation action, and that the algorithm adapts its model order accordingly to the data gathered and the individual dynamics of each edge. It is also interesting to observe that despite there was a two weeks period on which the robot

TABLE II
MODEL ERRORS AND EDGE DYNAMICS DISTRIBUTION

Site	Model	Average error	Edge model order 0	1	2	3
Open plan	Success rate	0.14	35	18	4	5
	Time [s]	1.88	24	19	10	9
Corridor	Success rate	0.17	71	52	9	7
	Time [s]	3.94	62	64	10	3

did not traverse these edges the model construction was not affected which demonstrates that the approach can handle uneven sampling as stated in section III.

B. Periodicity and prediction

When comparing the periodicity to human routines some very interesting results have been found especially on the door crossing edge. The dynamics of its traversability showed two main periodicities: one corresponding to a daily routine and another one for a weekly routine. Figure 4 shows the reconstructed signal for the dynamic model along different time frames.

Results show how the traversability probability for edge v_3-v_4 varies greatly within one same day, meaning that there are times during which the door is very likely to be closed (early morning) and times when it is very likely to be open (evenings) which correspond to the usual routines of the researchers at the lab. It can also be seen how the signal amplitude is lower during the weekends (Saturdays to Mondays) meaning that the certainty of this state is lower.

The dynamics for the traversing times also show some interesting results. On most edges there is some periodicity detected which means that human activity effectively affects this model, however the amplitude of this signal doesn't vary greatly due to the fact that there only very few variations on the times. However, most predicted values are close to the average of all times, proving that the dynamic model tries to create a better prediction than the stationary model also in this case.

V. SPATIO-TEMPORAL TOPOLOGICAL REPRESENTATIONS FOR MOTION PLANNING

Having shown that our Frequency-enhanced topological representation has the potential to represent and predict re-occurring action outcomes better than just a static model, we briefly illustrate how this additional knowledge can be employed in probabilistic motion planning. We apply an approach based on [20], which allows planning for tasks specified in linear temporal logic (LTL), using Markov decision process (MDP) models of the environment. We present this approach mostly as a case study of how our frequency-based predictions can be employed; our work is not limited to this particular planning technique.

The key idea of our implementation is that we generate time-indexed *Navigation MDPs* that are based on predicted probabilities of the previously introduced method. These MDPs allow us to reason about the probability of successfully accomplishing actions at a given time and generate

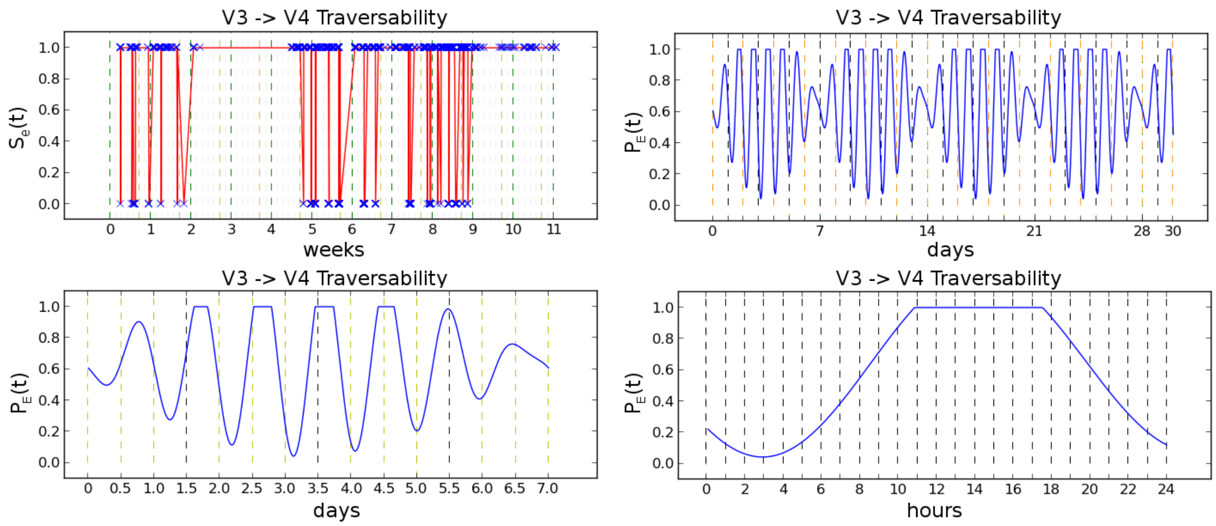

Fig. 4. reconstructed signal for traversability and time along different time periods, top left figure is the action outcomes used for the model building, the remaining three figures represent the predicted $p_e(t)$ state along different time frames, one month (top right), one week (bottom left) and one day (bottom right). Weekly and monthly periodicities are presented starting from Monday, the day depicted in the bottom left figure is a Thursday.

policies that maximise overall expected success of the LTL task.

In order to generate a policy at a given time t, we start by creating an MDP model based on the topological map $T = \langle V, E, N, nav, P_E \rangle$. This Navigation MDP at time t is defined as a tuple $\mathcal{M}_t = \langle S, \overline{s}, A, \delta \rangle$, where: (i) $S = V \cup \{s_f\}$ is a finite set of states, corresponding to the topological nodes, plus a *dump* state s_f, which is reached after a navigation action failure; (ii) $\overline{s} \in S$ is the initial state, corresponding to the current position of the robot in the environment; (iii) $A = E$ is a finite set of actions, corresponding to the edges in the topological map; (iv) $\delta : S \times A \times S \to [0, 1]$ is a probabilistic transition function, where $\sum_{s' \in S} \delta(s, a, s') \in \{0, 1\}$ for all $s \in S$, $a \in A$. For $v_i, v_j \in S$, if there is an edge $e = (v_i, v_j)$ in the topological map, we define $\delta(v_i, e, v_j) = p_e(t)$, $\delta(v_i, e, s_f) = 1 - p_e(t)$ and $\delta(v_i, e, v) = 0$ for all $v \in S \setminus \{v_j, s_f\}$.

In [20], it is shown how, given a *co-safe* LTL formula φ and a cost function defined over state-action pairs of the MDP (in our case, such function would be the expected time to navigate between two nodes in the environment), one can create policies that minimize the accumulated cost to generate a trace of the system that satisfies φ. Broadly speaking, LTL allows for the specification of goals that are not simply reaching a given target node in the environment, but can be *temporally extended* goals that require, for example, a set of nodes to be visited in a given order, or to visit a given node while avoiding a set of forbidden nodes. The co-safe fragment of LTL contains all the formulas that can be satisfied by a *finite* trace of the system. An example of such a task is a mail delivery robot that needs to distribute mail to different rooms in a building, and minimise the time spent in delivery so it can be available to do other tasks as soon as possible.

We adapted the approach in [20], and use the PRISM model checker [21] to generate a policy that maximizes the probability of satisfying a co-safe LTL formula, i.e., we generate the policy that fulfils the task while minimizing the probability of occurrence of a continuous navigation failure.

The fact that we can specify tasks that involve visiting more than one node in the topological map allows us to analyse the different choices taken by the robot at different times. More specifically, for the navigation MDPs obtained from the topological map depicted in Fig. 2, we analyse the policies obtained for formula $(\texttt{F}\, v_1 \vee \texttt{F}\, v_{14})$, i.e., "visit either node v_1 or node v_{14}". This task allows the policy to choose which node to try to visit first, taking into account the current position of the robot, and the traversability probabilities for the edges in the topological map. Furthermore, it is a type of task that is common for mobile robots. For example, a data gathering robot might want to unload its data, and in nodes v_1 and v_{14} there are data unloading stations it can use. Thus, to increase the robustness of the system, we want the robot to choose the station it can navigate to with the lowest probability of failure.

In Table III, we show the probabilities of being able to execute the task, starting on v_5, without any navigation failures, for different times of day. As expected, it is possible to see that during the times where it is more probable for people to be present in the office, the probability of fulfilling the task without navigation failures is higher. This is because the robot asks for human intervention when he has problems navigating, and the presence of people to help it increases the probability of fulfilling the navigation task (we do not consider these interactions with humans as failures). Furthermore, we also analyse the optimal action for v_5 at different times of day. This illustrates the choice the robot makes on which area of the environment to visit when at v_5. We depict the choice of visiting v_{14} in light gray, and the choice of visiting v_1 in dark gray. This choice is heavily

TABLE III

MAXIMUM PROBABILITIES OF FULFILLING TASK ($\mathrm{F}\, v_1 \vee \mathrm{F}\, v_{14}$) AT DIFFERENT DATES AND TIMES OF DAY.

Date\Time	0h	2h	4h	6h	8h	10h	12h	14h	16h	18h	20h	22h
Tuesday, 07-10-2014	0.646	0.654	0.681	0.721	0.795	0.891	0.884	0.916	0.938	0.938	0.901	0.777
Thursday, 09-10-2014	0.717	0.693	0.704	0.746	0.803	0.857	0.946	0.971	0.938	0.877	0.751	0.687
Sunday, 12-10-2014	0.736	0.716	0.750	0.798	0.840	0.898	0.914	0.889	0.871	0.849	0.793	0.703

influenced by the probability of the the door being open: the policy needs to choose between a longer path, where there are more possibilities of navigation failures, but each one with low probability, or a shorter path that goes through the door, which has a high probability of being closed at some times of day. One can see that the choice to go through the door happens roughly on weekday afternoons.

VI. CONCLUSION

In this paper we set out to investigate a spectral representation of action success probabilities in mobile robot topological navigation in order to improve task and navigation planning in a probabilistic framework. The novel topological representation integrates a model of the periodicity of action outcomes with the spatial representation of a standard topological map, facilitating spatio-temporal situation-dependent predictions basic on a model of periodic processes. The approach utilises the Frequency Enhanced Mapping framework which describes the dynamics of the environment by means of periodicity, amplitude and time shift of these underlying periodic processes, and augments the edges of the topological graph with this information.

It has been shown that indeed a large variability in the long-term data sets can be modelled as a set of very few superimposed periodic processes that actually coincide with usual daily routines of humans in the studied indoor office environments. The framework presented in this paper thus enables the robot to predict future action success to a high degree of accuracy, providing it with the augmented knowledge about the dynamics of the environment. It has been shown to be effective in coping with long-term data sets.

Based on these findings it has been discussed and analysed how the approach can be used to build a MDP model of the environment. This model can then be used to predict global probabilities of success of tasks specified as co-safe LTL formulas, and to generate policies that maximize those probabilities. Furthermore, we showed how this time dependence influences the order for which LTL tasks are executed in different times of day. In future work, we plan to extend the MDP planning approach to allow the addition of cost optimal policy generation, along with maximization of the probability of satisfying the LTL task. Also, in order to have a more accurate model of the environment, we will model certain structures present in office-like environments (e.g., doors) explicitly. Finally, we plan to have these global probabilities, along with the expected execution times, used by a scheduling mechanism that uses the information to decide the best times for the robot to try to execute given tasks.

ACKNOWLEDGMENTS

The research leading to these results has received funding from the EU FP7 under grant agreement No. 600623.

REFERENCES

[1] M. Hanheide *et al.*, "Exploiting Probabilistic Knowledge under Uncertain Sensing for Efficient Robot Behaviour," in *IJCAI*, 2011.
[2] T. Krajník, J. P. Fentanes, G. Cielniak, C. Dondrup, and T. Duckett, "Spectral analysis for long-term robotic mapping," in *ICRA*, 2014.
[3] D. Wolf and G. Sukhatme, "Mobile robot simultaneous localization and mapping in dynamic environments," *Autonomous Robots*, 2005.
[4] L. Montesano *et al.*, "Modeling dynamic scenarios for local sensor-based motion planning," *Autonomous Robots*, 2008.
[5] F. Dayoub, G. Cielniak, and T. Duckett, "Long-term experiments with an adaptive spherical view representation for navigation in changing environments," *Robotics and Autonomous Systems*, 2011.
[6] P. Biber and T. Duckett, "Experimental analysis of sample-based maps for long-term slam," *International Journal of Robotics Research*, 2009.
[7] W. Churchill and P. Newman, "Practice makes perfect? managing and leveraging visual experiences for lifelong navigation," in *ICRA*, 2012.
[8] T. Krajník, S. Pedre, and L. Přeučil, "Monocular Navigation System for Long-Term Autonomy," in *Proc. of ICAR*, 2013.
[9] C. Cadena, D. Gálvez-López, J. D. Tardós, and J. Neira, "Robust place recognition with stereo sequences," *IEEE Trans. on Robotics*, 2012.
[10] F. Dayoub and T. Duckett, "An adaptive appearance-based map for long-term topological localization of mobile robots," in *IROS*, 2008.
[11] M. Yguel *et al.*, "Wavelet occupancy grids: a method for compact map building," in *Field and Service Robotics*, 2006.
[12] T. Kucner *et al.*, "Conditional transition maps: Learning motion patterns in dynamic environments," in *Int. Robots and Systems*, 2013.
[13] K. Z. Haigh and M. M. Veloso, "Learning situation-dependent costs: Improving planning from probabilistic robot execution," *Robotics and Autonomous Systems*, 1999.
[14] M. Sridharan, J. Wyatt, and R. Dearden, "Planning to see: A hierarchical approach to planning visual actions on a robot using POMDPs," *Artif. Intell.*, vol. 174, pp. 704–725, July 2010.
[15] J. J. Velez *et al.*, "Modelling observation correlations for active exploration and robust object detection," *JAIR*, 2012.
[16] R. N. Bracewell and R. Bracewell, *The Fourier transform and its applications*. McGraw-Hill New York, 1986, vol. 31999.
[17] T. Krajník, J. Santos, B. Seemann, and T. Duckett, "Froctomap: An efficient spatio-temporal environment representation," in *Advances in Autonomous Robotics Systems*, 2014.
[18] 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 *International Conference on Intelligent Robots and Systems (IROS)*, 2014.
[19] 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.
[20] B. Lacerda, D. Parker, and N. Hawes, "Optimal and dynamic planning for Markov decision processes with co-safe LTL specifications," in *Int. Conf. on Intelligent Robots and Systems (IROS)*, 2014.
[21] M. Kwiatkowska, G. Norman, and D. Parker, "PRISM 4.0: Verification of probabilistic real-time systems," in *Computer Aided Verification (CAV'11)*, ser. LNCS, vol. 6806. Springer, 2011, pp. 585–591.