

Frequency Map Enhancement: Introducing Dynamics into Static Environment Models

Tomáš Krajník

Jaime Pulido Fentanes

João Santos

Tom Duckett

Abstract—We present applications of the Frequency Map Enhancement (FreME_n), which improves the performance of mobile robots in long-term scenarios by introducing the notion of dynamics into their (originally static) environment models. Rather than using a fixed probability value, the method models the uncertainty of the elementary environment states by their frequency spectra. This allows to integrate sparse and irregular observations obtained during long-term deployments of mobile robots into memory-efficient spatio-temporal models that reflect mid- and long-term pseudo-periodic environment variations. The frequency-enhanced spatio-temporal models allow to predict the future environment states, which improves the efficiency of mobile robot operation in changing environments. In a series of experiments performed over periods of weeks to years, we demonstrate that the proposed approach improves mobile robot localization, path and task planning, activity recognition and allows for life-long spatio-temporal exploration.

I. INTRODUCTION

As robots gradually enter human-populated environments, they have to deal with the fact that the environments are uncertain because they change over time. While the probabilistic mapping methods used in robotics can handle uncertain and incomplete environment knowledge, their theoretical foundations assume that the uncertainty is caused by sensor noise rather than by natural processes that govern the environment changes. The assumption of a static world negatively impacts the ability of these models to reflect the environment changes and effectively support long-term autonomous operation of mobile robots. However, several studies [1], [2], [3], [4], [5] indicated that explicit modeling of the environment changes improves localization robustness.

Biber and Duckett [5] proposed to represent the world dynamics by multiple maps with different timescales, which are switched on the fly based on their consistency with the current observations. Dayoub et al. [6] present a system that evaluates the persistence of visual features over time in order to identify features that are more likely to be stable. Churchill and Newman [1] demonstrated that clustering spatially-close observations into ‘experiences’ improves long-term localization. The article [3] associates each cell of an occupancy grid with a hidden Markov model, which improves the localization robustness as well. Kucner’s method [7] assumes that occupancies of grid cells are influenced by a moving objects, which allows to infer typical motion patterns in a given environment. Sünderhauf’s method [4] proposes to learn typical appearance changes caused by seasonal factors

and use this knowledge for long-term predictions of environment appearance. Finally, Rosen et al. [8] use Bayesian-based survivability analysis to predict which landmarks will be visible after some time and which are going to disappear.

Our approach to environment change modeling is based on the assumption that some of the mid- to long-term processes that cause the environment changes are (pseudo-)periodic, e.g. seasonal foliage variations, daily illumination cycle or routine human activities. To reflect this assumption, we represent the probability of each local environment state not by a single value, but by a probabilistic function of time composed of several harmonic functions whose periodicities and amplitudes relate to the frequencies and influences of these hidden processes.

II. METHOD DESCRIPTION

The proposed method, coined the Frequency Map Enhancement (FreME_n), represents the probability of each environment state by a function of time

$$p(t) = p_0 + \sum_{j=1}^n p_j \cos(\omega_j t + \varphi_j), \quad (1)$$

where n is the number of environment processes taken into account, ω_j , φ_j and p_j relate to the frequencies, time offsets and influences of these processes, and p_0 is the mean probability of the state. To obtain the parameters ω_j , φ_j and p_j , we first obtain the frequency spectra $S(\omega)$ of long-term observations of each environment state $s(t)$ by means of a (non-uniform) Fourier transform [9], i.e. $S(\omega) = \mathcal{FT}(s(t))$. The parameters ω_j , φ_j and p_j in Equation (1) are equal to the amplitudes, phases and frequencies of the n most prominent spectral components of the spectrum $S(\omega)$. To deal with the fact that robots might observe the environment on an irregular basis, we employ a non-uniform Fourier Transform scheme [9] similar to the one used in [10].

The approach, which was originally presented in [11], can be applied to all environment models that represent the world as a set of independent component with binary states. In this short overview, we will show its use does not improve only long-term localisation in changing environments [2], [12], but that it also improves robotic search [13], path planning [14] and activity recognition. We will also show that the time-dependent probability of the environment states expressed by Equation (1) allows the calculation of the spatio-temporal environment entropy, which, combined with information-theoretic planning, results in life-long spatio-temporal exploration of dynamic environments [9], [15].

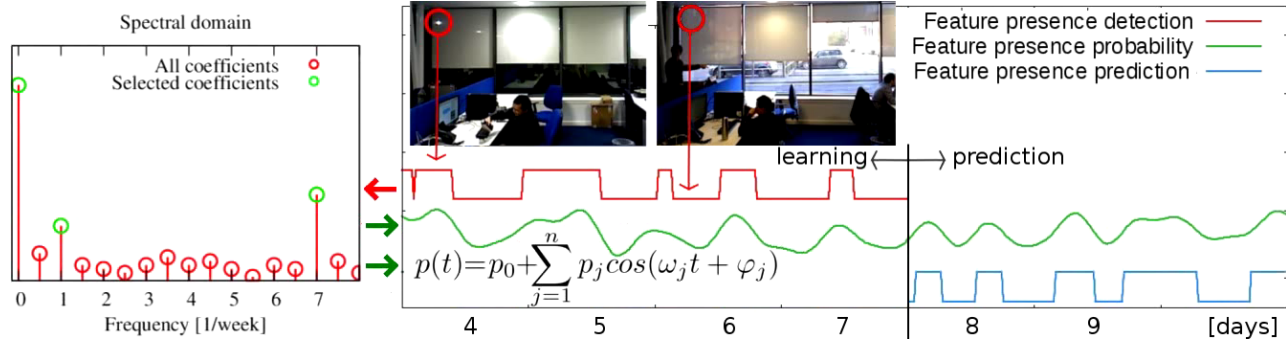


Fig. 1: Frequency-enhanced feature map [2] for visual localization: The observations of image feature visibility (centre, red) are transferred to the spectral domain (left). The most prominent components of the model (left, green) constitute an analytic expression (centre, bottom) that represents the probability of the feature being visible at a given time (green). This allows to predict the feature visibility at a time when the robot performs self-localization (blue).

III. VISUAL LOCALIZATION

The problem of visual place recognition in changing environments has received considerable attention during recent years [16]. We propose to represent the variations in appearance of different locations by modeling the visibility of individual image features in the frequency domain. Thus, we can predict which visual features are going to be visible at which time and use these time-specific features to localise the robot [2]. To evaluate our approach, we performed both indoors and outdoors experiments. The indoor experiment was performed at the Lincoln Centre for Autonomous Systems, where a SCITOS-G5 robot captured images of 8 areas every 10 minutes for one week and used the models created to localize itself after one week, three months and one year [2]. The outdoor experiment was performed in the Stromovka park in Prague, where a P3AT robot captured images of designated places on a monthly basis for one year and used the FreMEEn feature map to determine its location during three testing runs during the following year. While the indoor dataset was mainly affected by the daily illumination cycle and human activities, the outdoor dataset captured seasonal variations of foliage.




Fig. 2: The dependency of localization error on the number of features used.

The dependence of the localization error on the number of features used is shown on Figure 2, which indicates that the frequency-enhanced models that generate a set of likely-to-be-visible features for a particular time outperform the ‘static’ approach that relies on the most stable features, i.e. modelling the appearance variations by our approach improves the robustness of localization.

IV. CONTINUOUS LASER-BASED LOCALIZATION

While the advantages of using dynamic representations for visual localisation are clear, because the environment appearance variations are significant due to the passive nature of the cameras, the usefulness of dynamic maps for laser-based, 2d localisation was demonstrated only in highly dynamic environments, such as parking lots [3]. However, maps are not useful only for localization, but also for planning. Thus, a mobile robot that operates in the long-term should be able to reason about the nature of the environment changes it encountered during its deployment: knowing that obstacles which were blocking a corridor a hours ago are likely to be gone by now or that during noon, the cafeteria is too crowded, is beneficial when planning the robot’s path.

However, the major part of the changes observed in 2D occupancy grids build by lasers are not periodic. Typically, they are caused by temporary stationary objects that tend to disappear after some time. Since modeling periodic changes alone is not sufficient, we combined FreMEEn with the approach presented in [3] and created a 2D occupancy grid that models both the periodicity and stationarity of the changes [12]. This FreMEEn 2D occupancy grid was



Environment change example (cubboard doors open/closed)

Fig. 3: Example of the regular variations and corresponding map sections predicted for morning (left) and evening (right).

integrated into the ROS navigation stack of our SCITOS-G5 robot, which operated in a populated open plan office for 2 weeks. During these two weeks, we observed that the efficiency of the robot navigation gradually improved (e.g. path re-planning was triggered less often), while the reduction in the localization error was only marginal [12].

V. TOPOLOGICAL PATH PLANNING

Imagine a robot operating in an office like environment 24/7, performing different user-defined tasks at different locations. The robot needs to schedule not only these tasks, but also has to determine when to visit its charging station. To create the schedule, the robot needs to predict which areas of the environment are going to be accessible at which times.

To address this problem, we represented the environment as a topological map, where the traversability of individual edges was modelled by FreMEn [14], see Figure 4. Using this topological map, the robot can not only plan its path, but it can also determine what is the chance of the path’s successful traversal, i.e. the chance of reaching a given destination at a particular time. To evaluate our approach,

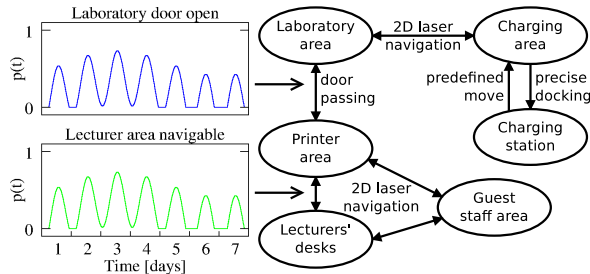


Fig. 4: Partial map of the operational environment with temporal edge traversability models. Nodes represent locations and edges movement actions which can fail, e.g. *door passing* requires an open door. The (illustrative) graphs show the predicted probability of successful edge traversal.

we let our SCITOS-G5 robot operate in an populated office-like environment for more than 10 weeks, during which the robot learned that two edges of the topological map exhibit periodic changes to their traversability. The first edge corresponded to a laboratory safety door that was kept closed at night and the second edge led through a narrow area behind lecturer’s desks, which was occasionally blocked by chairs. Using its knowledge about the dynamics of these edges, the robot could infer that the best time to perform its activities in the office areas was afternoon during the weekdays, because at other times, it would risk that after completing these tasks, it would not be able to return to the laboratory and reach its charging station, see the map in Figure 4.

VI. ROBOT SEARCH

Another combination of topological representations with FreMEn was used to model object and people presence in a robotic search scenario. Here, a mobile robot has to find a certain object as quickly as possible. For the sake of simplicity, we assume that the object is detected as soon as the robot arrives at the location of the searched object.

To plan an efficient search path through the individual locations, the robot needs to take into account the probability of object occurrence. Improved knowledge about possible object location leads to more efficient plans and hence, shorter times to locate the desired object. We formulated

the search as a path planning problem in a graph where the probability of object occurrences at particular nodes is a function of time represented by FreMEn [13]. To




Fig. 5: Aruba ‘CASAS’ [17] apartment with probabilities of person presence in two rooms.

evaluate our approach, we used three datasets collected over several months containing person and object occurrences in residential and office environments. Several types of spatio-temporal models were created for each of these datasets and the efficiency of the search method was assessed by measuring the time it took to locate a particular object. The experimental results indicated that representing the dynamics of object occurrences by FreMEn reduced the search time by 25% to 65% compared to maps that consider the probability of object occurrences as independent of time.

VII. LIFE-LONG EXPLORATION

In the previous scenarios, the map was created from sensory data which were gathered when the robot was performing user-motivated tasks. This passive approach to mapping is not only slow, but it typically results in an incomplete knowledge that might lead to misinterpretations of the processes that govern the environment changes. To deal with this, the robot has to explore its environment in an active way, which does not only mean that it should visit all the relevant locations at least once, but it should also revisit them to understand how they change over time.

The crucial issue is to determine where and when to perform observations in order to refine and complete the spatio-temporal model. Since the FreMEn model predicts the probability of the environment states, we use it to calculate the states’ entropy, which directly corresponds to the amount of information obtained by observing these states at a particular time. Application of information-based exploration methods to the spatio-temporal entropy predicted by FreMEn resulted in intelligent and continuously improving exploratory behaviour, which evolves as the environment knowledge becomes more refined over time [9], [15].

Figure 6 illustrates the exploratory behaviour on the Aruba [17] dataset, where the robot created a spatio-temporal




Fig. 6: 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.

model of person presence used for the robot search [13]. During the first day, the robot has no knowledge of the environment and it has no room or time preference when scheduling its observations. After the first day, it schedules more visits of the rooms where the person presence changes more often. The second day observations provide information about the rooms’ dynamics: the robot assumes that the bedroom has a daily periodicity and that the kitchen is visited five times per day. This causes the expected information gain to be time-dependent – e.g. evening and morning observations of the bedroom provide more information than in the afternoon, which is reflected by the exploration schedule, see the last row of Figure 6.

VIII. ACTIVITY RECOGNITION

Using the office and household datasets from Sections III and VI respectively, we also tested the use of the FreMEn as a model providing temporal-based priors for human activity recognition. The FreMEn models were created incrementally by a Bayesian update scheme, which was performed every time an activity was recognized. The method gradually learned about the typical rhythms of the people’s activities, effectively reducing the error in activity classification, see Figure 7 and article [18].




Fig. 7: Activity recognition error over time

IX. CONCLUSION

We presented applications of a method that improves the efficiency of long-term mobile robot operation in changing environments. The method assumes that in a mid-term perspective, the environment is influenced by processes which might be periodical and that the evolution of the some environment states can be described by the periodicity, amplitude and time shift of these underlying processes. To identify the parameters of these processes and to predict the environment’s local state we use techniques based on the Fourier transform. We gave an overview of the methods’ applications so far, showing that in long-term scenarios, it reduces localisation error [2], [12], speeds-up robotic search [13], improves path planning [14], activity recognition. and allows for active, life-long environment exploration [9], [15]. To facilitate the use of the method by other researchers, we published its ROS-compatible source code at <http://fremen.uk>.

REFERENCES

- [1] W. S. Churchill and P. Newman, “Experience-based navigation for long-term localisation,” *IJRR*, 2013.
- [2] 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 *IROS*, 2014.
- [3] G. D. Tipaldi *et al.*, “Lifelong localization in changing environments,” *The International Journal of Robotics Research*, 2013.
- [4] P. Neubert, N. Sünderhauf, and P. Protzel, “Superpixel-based appearance change prediction for long-term navigation across seasons,” *Robotics and Autonomous Systems*, 2014.
- [5] P. Biber and T. Duckett, “Dynamic maps for long-term operation of mobile service robots,” in *Robotics: Science and Systems*, 2005.
- [6] 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.
- [7] T. Kucner, J. Saarinen, M. Magnusson, and A. J. Lilienthal, “Conditional transition maps: Learning motion patterns in dynamic environments,” in *Proceedings of the International Conference on Intelligent Robots and Systems (IROS)*, 2013.
- [8] D. M. Rosen, J. Mason, and J. J. Leonard, “Towards lifelong feature-based mapping in semi-static environments,” in *International Conference on Robotics and Automation (ICRA)*. IEEE, 2016, to appear.
- [9] T. Krajník, J. Santos, and T. Duckett, “Life-long spatio-temporal exploration of dynamic environments,” in *ECMR*, 2015.
- [10] S. Ferraz-Mello, “Estimation of periods from unequally spaced observations,” *The Astronomical Journal*, vol. 86, p. 619, 1981.
- [11] T. Krajník, J. P. Fentanes, G. Cielniak, C. Dondrup, and T. Duckett, “Spectral analysis for long-term robotic mapping,” in *ICRA*, 2014.
- [12] T. Krajník, J. P. Fentanes, M. Hanheide, and T. Duckett, “Persistent Localization and Life-long Mapping in Changing Environments using the Frequency Map Enhancement,” in *International Conference on Intelligent Robots and Systems (IROS)*, 2016, in review.
- [13] 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 *ICRA*, 2015.
- [14] 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 *ICRA*, 2015.
- [15] J. M. Santos, T. Krajník, J. Pulido Fentanes, and T. Duckett, “Lifelong information-driven exploration to complete and refine 4d spatio-temporal maps,” *Robotics and Automation Letters*, 2016.
- [16] S. Lowry, N. Sunderhauf, P. Newman, J. Leonard, D. Cox, P. Corke, and M. Milford, “Visual place recognition: A survey,” *IEEE Transactions on Robotics*, vol. PP, no. 99, pp. 1–19, 2015.
- [17] D. J. Cook, “Learning setting-generalized activity models for smart spaces,” *IEEE Intelligent Systems*, no. 99, p. 1, 2010.
- [18] C. Coppola, T. Krajník, N. Bellotto, and T. Duckett, “Learning temporal context for activity recognition,” in *European Conference on Artificial Intelligence (ECAI)*, 2016, to be submitted.