

Spectral Analysis for Long-Term Robotic Mapping

Tomáš Krajník Jaime Pulido Fentanes Grzegorz Cielniak Christian Dondrup Tom Duckett

Abstract—This paper presents a new approach to mobile robot mapping in long-term scenarios. So far, the environment models used in mobile robotics have been tailored to capture static scenes and dealt with the environment changes by means of ‘memory decay’. While these models keep up with slowly changing environments, their utilization in dynamic, real world environments is difficult.

The representation proposed in this paper models the environment’s spatio-temporal dynamics by its frequency spectrum. The spectral representation of the time domain allows to identify, analyse and remember regularly occurring environment processes in a computationally efficient way. Knowledge of the periodicity of the different environment processes constitutes the model predictive capabilities, which are especially useful for long-term mobile robotics scenarios.

In the experiments presented, the proposed approach is applied to data collected by a mobile robot patrolling an indoor environment over a period of one week. Three scenarios are investigated, including intruder detection and 4D mapping. The results indicate that the proposed method allows to represent arbitrary timescales with constant (and low) memory requirements, achieving compression rates up to 10^6 . Moreover, the representation allows for prediction of future environment states with $\sim 90\%$ precision.

Index Terms—long-term autonomy, mobile robotics, spatio-temporal mapping

I. INTRODUCTION

Long-term robotic autonomy has not yet been achieved due to many challenges, one of them being the fact that robotic mapping is vulnerable to changes in the environment. However, many future tasks performed by mobile robots will be carried out in places where humans perform their usual activities, causing the environment to change. Many of these activities consist of daily routines that are performed on a regular basis. This regularity can be exploited by the robots to build more robust representations of their surroundings. To address this issue, we propose to treat the observed environment states as signals over which a spectral analysis can be performed to identify regularities of environmental variations occurring on timescales from days to months. Knowledge of these regularities allows efficient representation of the world model’s temporal domain as well as prediction of the environment properties at a specific time.

This ability would allow to improve many tasks that the mobile robots have to perform. For example, a robot can predict the appearance of a particular location at a certain time in order to achieve more robust localization. Moreover, higher-level path planning can use the approach to predict which doors will be open or which areas will be difficult to navigate because they will be crowded at particular times.

Mapping a static environment is a problem that has been widely studied for a long time [1], however mapping dynamic environments is still an open problem. So far, mapping dynamic environments has been done by removing moving objects from the representation of the environment [2], [3] or by tracking these objects and classifying them as moving landmarks [4], [5], [6]. In general, separation-based approaches can handle some problems of dynamic mapping, but they assume that most of the environment is static, which makes them unsuitable for scenarios where the environment appearance varies significantly. In [7] a new map type that represents local maps at different time scales is presented, where the best map for localisation is chosen by its consistency with current readings, which provided improved localisation over large time scales.

Adaptive approaches never assume the map to be complete and perform continuous mapping, adding new features to the map every time the robot observes its environment. In these approaches, the key problem is managing map size [8], [9], [10]. [11] presents a feature persistence system based on temporal stability in sparse vision-based maps. Some authors propose systems that learn a fixed set of possible states for the dynamic objects in the environment, e.g. corresponding to open and closed doors [12], [13], but this approach is limited in the real world, where the number of states is unpredictable and this approach does not offer state prediction capabilities.

[14] 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. [15] presents a representation which models transitions of dynamic objects in the environment. Spectral analysis has also been used for 2D [16] and 3D [17] registration. Finally, [18] has shown state prediction can be useful on long-term robotic tasks by proposing an implementation of an appearance change prediction (ACP) method to improve the performance of place recognition on a large-scale dataset under extremely different conditions.

The representation proposed in this paper models the environment’s spatio-temporal dynamics by its frequency spectrum. We propose to model local states of an environment by means of a probability function which is a superposition of periodical functions. The model includes recurrent environment changes, which improves the robot’s knowledge of its surroundings and the events that take place in them. We hope that this proposal becomes a useful tool for modelling the environment and leads to developing new representations of the environment that consider the spatio-temporal dynamics that take place there.

II. SPECTRAL REPRESENTATION FOR SPATIO-TEMPORAL ENVIRONMENT MODELS

Most mapping approaches assume that the principal components of a particular environment model can be in two distinct states. For example, cells of an occupancy grid are occupied or free, edges of a topological map are traversable or not, doors are opened or closed, rooms are vacant or occupied, landmarks are visible or occluded, etc. In a typical situation, the state of each model component is uncertain, because it is measured indirectly by means of sensors which are affected by noise. A common way to represent the uncertainty in the state estimate of the j^{th} world model component is by its associated probability p_j . This allows to counter the effect of noisy measurements by employing statistical methods, such as Bayesian filtering. While Bayesian filtering methods allow to keep up with a changing environment, the mathematical foundations they are based on assume a static world, i.e. the p_j of the world components are assumed to be constant. As a result, a change in the environment causes the old state to be ‘forgotten’ over time.

Once we assume that p_j is a function of time, we need to outline a suitable representation for $p_j(t)$. Although one could simply store the whole history of the environment model, such an approach would quickly face memory limitations. Typical static 3D models of complex environments contain millions of distinct components [19]. Storing the entire model history is infeasible. Moreover, in the context of robotic mapping, it is not clear how to utilize the past estimates of the environment models, i.e. what is the relation of the past models to the current state of the world.

In our approach we assume that the variations of the environment are caused by a number of unknown processes, which might be periodical. If we could identify the influence and periodicity of these processes, we could then calculate the probabilities $p_j(t)$ from the description of these processes. To identify these periodical processes, we propose to use a frequency transform, namely the Fourier transform [20].

A. An introduction to the Fourier Transform

The Fourier Transform (FT) is a well-established mathematical tool widely used in the field of statistical signal processing. It transforms a function of time $f(t)$, into a function of frequency $F(\omega)$, i.e. $F(\omega) = \mathcal{F}(f(t))$. The function $F(\omega)$ is commonly referred to as the frequency spectrum of $f(t)$. The Fourier transform is invertible, and therefore, one can recover the function $f(t)$ from its spectrum $F(\omega)$, i.e. $f(t) = \mathcal{F}'(F(\omega))$. If one wants to analyze or alter the periodic properties of a process characterized by a function $f(t)$, it is reasonable to calculate its spectrum $F(\omega)$, perform the analysis or alteration in the frequency domain, and then transform the altered spectrum $F'(\omega)$ back to the temporal domain. Such a process is referred to as spectral analysis.

Typically, $F(\omega)$ is a complex-valued function, whose absolute values and arguments correspond to the amplitudes and phase shifts of the frequency components ω . Considering

that $f(t)$ is a real periodical discrete function, the spectrum $F(\omega)$ can be represented by a finite set of complex numbers.

B. The proposed representation

Let us represent the environment as a set of independent components, which can be in two distinct states. Without loss of generality, we will explain our approach with an occupancy grid. Let us assume that each grid cell state $s_j = \{\text{occupied}, \text{free}\}$ is not constant, but is a function of time, i.e. $s_j(t)$. The uncertainty of the state $s_j(t)$ is represented by its probability $p_j(t)$. Now, let us assume that the occupancy of each grid cell is affected by a set of unknown periodical processes, which can be identified by the Fourier Transform. Since we assume the occupancy of the individual cells to be independent, we can explain the use of the Fourier transform on the state $s(t)$ of a single cell.

1) *The spectral model:* The main idea behind the proposed model is to measure the temporal sequence of states $s(t)$ and calculate their frequency spectrum by means of a Fourier Transform as $\mathcal{S}(\omega) = \mathcal{F}(s(t))$. Then, we select the l most prominent (i.e. of highest absolute value) coefficients S_i of the spectrum \mathcal{S} and store them along with their frequencies ω_i in a set \mathcal{P} . The coefficients stored in \mathcal{P} are then used to recover the smoothed signal $p(t)$ by means of the Inverse Fourier Transform $p(t) = \varsigma(\mathcal{F}'(P(\omega)))$, where ς denotes a saturation function that ensures that $p(t) \in \langle 0, 1 \rangle$. One can easily verify that $P(s(t) = \text{occupied}) = p(t) \geq 0$ and $P(s(t) = \text{occupied}) + P(s(t) = \text{free}) = p(t) + 1 - p(t) = 1$, i.e. $p(t)$ satisfies Kolmogorov’s axioms and therefore is a probability. Thresholding the probability $p(t)$ allows us to calculate an estimate $s'(t)$ of the original state $s(t)$. In order not to lose any information of the original signal, the differences between $s'(t)$ and $s(t)$ are stored in an outlier set \mathcal{O} , which is Δ -encoded, see Figure 1.

Thus, our model of the state consists of two finite sets \mathcal{P} and \mathcal{O} . The set \mathcal{P} consists of l triples $\text{abs}(P_i)$, $\text{arg}(P_i)$ and ω_i , which describe the amplitudes, phase shifts and frequencies of the model spectra. Each such triple might be interpreted as the importance, time offset and periodicity of one particular periodical process influencing the state $s(t)$. We will refer to the number of modeled processes l (i.e. to the number of triples in \mathcal{P}) as the ‘order’ of the spectral model. The set \mathcal{O} represents a set of k time intervals, during which the state $s(t)$ did not match the state $s'(t)$ calculated from $p(t)$. Internally, the set \mathcal{O} is implemented as a sequence of values, indicating the starts and ends of time intervals when the predicted and observed state did not match, i.e. $s'(t) \neq s(t)$. Each interval is represented by its limits $\langle t_{2k}, t_{2k+1} \rangle$.

2) *Model adaptation:* To be able to build, maintain and use this representation, we define four operations: reconstruction of the original state $s(t)$, addition of a new measurement, model update and prediction of the future state with a given confidence level. The aforementioned representation allows us to retrieve the cell state $s(t)$ by means of the following equation:

$$s(t) = (\mathcal{F}'(\mathcal{P}) > 0.5) \oplus (t \notin \mathcal{O}), \quad (1)$$

where \oplus is a XOR operation. The idea behind this equation is to reconstruct the probability $p(t)$ from the spectrum \mathcal{P} , set $s(t)$ to 1 if $p(t)$ exceeds 0.5 and finally to negate $s(t)$ if t belongs to the set of outliers \mathcal{O} .

Whenever a real state $s^m(t)$ is measured, we calculate $s(t)$ by means of Equation (1) and if it differs from $s^m(t)$, the current time t is added to the set \mathcal{O} :

$$s^m(t) \neq ((\mathcal{F}'(\mathcal{P}) > 0.5) \oplus (t \notin \mathcal{O})) \rightarrow \mathcal{O} = \mathcal{O} \cup t. \quad (2)$$

Since $p(t)$ does not predict $s(t)$ with perfect accuracy, the set \mathcal{O} is likely to grow larger as measurements are added.

To update the model, we reconstruct $s(t)$ in the desired time interval $\langle t_{start}, t_{end} \rangle$ and calculate its spectrum \mathcal{P} . Again, we select the l coefficients with highest absolute values $|P_i|$ and reconstruct the outlier set \mathcal{O} by means of Equation 2. In a typical situation, the updated spectrum \mathcal{P} would reflect $s(t)$ more accurately, causing reduction of the set \mathcal{O} . Note, that the spectral model order l can be changed prior to the update step without causing any loss of information.

3) *Prediction*: Note, that the Equation (1) allows for calculating $s(t)$ for any t and that the threshold value of 0.5 can be set arbitrarily. In fact, a threshold c such that $P(s(t) = occupied) > c$ represents a confidence level of the grid cell being full at time t . Therefore, we can use Equation (1) for future prediction of $s(t)$ with a certain confidence level c . In the case of prediction, the outlier set \mathcal{O} is not included in the calculation and the predicted state might not match the real state, so we denote it as $s'(t, c)$. To simplify notation, we also define $s'(t)$ as $s'(t, 0.5)$. Therefore, $s'(t, c)$ and $s'(t)$ can be calculated as follows:

$$s'(t, c) = \mathcal{F}'(\mathcal{P}) > c. \quad (3)$$

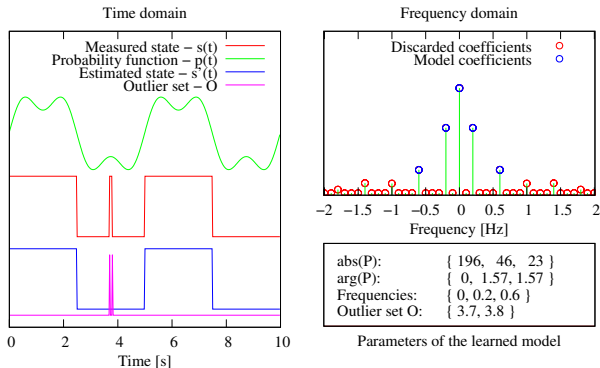


Fig. 1. An example of the measured state and its spectral model. The left part shows the time series of the measured state $s(t)$, probability estimate $p(t)$, predicted state $s'(t)$ and outlier set \mathcal{O} . The upper right part shows the absolute values of the frequency spectrum of $s(t)$ and indicates the spectral coefficients, which are included in the model.

An example of the third-order spectral model which represents a quasi-periodic function is provided in Figure 1. Since the observed processes are not identified perfectly, one can expect that the prediction becomes less and less accurate

over time. However, modelling the uncertainty of the model prediction is outside the scope of this paper.

C. Notes on notation

Throughout the rest of the article, we will keep the notation introduced in Section II-B. Therefore, $s_j(t)$ represents the measured state of the j^{th} world component, $P_{i,j}(\omega_i)$ is a complex number representing amplitude and phase shift of the state's frequency component ω_i , (i.e. i_{th} coefficient of \mathcal{P}), the value of $p_j(t)$ is the state's probability estimate given by the inverse Fourier transform of \mathcal{P} , $s'_j(t)$ is the most probable state at a given time t and $s'_j(t, c)$ is the estimate of $s_j(t) = 1$ with a confidence level c . For the sake of simplicity, the index j will be omitted in the case where our description concerns a single world model component only.

D. Spectrum and periodical processes

In the rest of the article, we will try to examine the tractability of using the Fourier transform as a core component of spatio-temporal models for mobile robotics. In particular, we will investigate the following questions:

- How many parameters of the spectrum typically have to be stored to represent and predict the environment state?
- What is the accuracy and efficiency of the spectral representation?
- Are the predictions of this approach good enough to detect anomalous situations?
- How does the approach deal with sensor noise and localization uncertainty?

To answer these questions, we analysed several types of environment models gathered by a mobile robot, which was continuously operating for more than one week in an indoor environment shared with humans.

III. DATA COLLECTION

Our experimental platform consists of a SCITOS-G5 mobile robot (see Figure 2) equipped with RGB-D and laser sensors. The robot gathered two datasets in two different environments: a staff office and the robotics lab of the Lincoln Centre for Autonomous Systems.

A. The office dataset

The first dataset, considered for basic evaluation, was gathered from a stationary position with sensors aimed at the door of one of the building's offices. The range measurement was used to establish the occupancy of a single cubic cell located in the middle of the room entrance. Since the office has an open door policy, the door remained open whenever the office was occupied by a person. Therefore, the measured state $s(t)$ not only indicated if the particular area of the environment was occupied or free, but also corresponded to the presence of people inside the office. Every time someone went through the door, the monitored area was briefly occupied and the room was considered empty, which introduced noise on the measured state $s(t)$. The measurements were taken continuously for one week (July 23-29 2013) at a rate

of 30Hz, so $s(t)$ consisted of 18 million values. After this week, two additional full-day datasets (July 31 and August 2 2013) were gathered.

B. The laboratory datasets

To gather the second dataset the robot was programmed to visit three designated areas of the robotics lab every five minutes. Each time an area was visited, the robot created a 2D and 3D point cloud and analysed the onboard color camera image to check for the presence of people. Thus, the robot created three datasets of different dimensionality, where each dataset contains observations from three different places, see Figure 2. We will refer to these datasets as Lab-2D, Lab-3D and Lab-1D respectively.



Fig. 2. The robot and its view of two locations of the ‘Lab’ dataset.

The data gathering process started on August 2013 and is still in progress. For this study, we use the data collected during the first week of September consisting of approximately 12,000 point clouds and 6,000 results of people detection.

The autonomous patrolling has been based on combination of the ROS nav stack and the visual localization method proposed in [21]. The robot reports its status regularly [22], so the occasional failures can be dealt with immediately.

IV. ALGORITHM PERFORMANCE

To answer the questions set in Section II-D, we analyse the performance of the proposed representation on the datasets described in Section III.

A. Model accuracy and efficiency

Knowing the coefficients $P_i(\omega_i)$ of the spectrum \mathcal{P} allows us to calculate an estimate $s'(t)$ of the original state $s(t)$. A natural concern is the accuracy of reconstruction of $s'(t)$, which affects the prediction capabilities of the model and the size of the outlier set \mathcal{O} . One can expect that increasing the number of spectral parameters will increase the reconstruction accuracy. However, as the number of parameters grows, the model becomes more adjusted to the specific time series of $s(t)$ and loses its generality. This loss of generality would hamper the ability of the model to predict the environment state in the future.

We define the accuracy of the spectral model $q(t_a, t_b)$ as the ratio of the correctly estimated signal $s'(t)$ on a given time interval $t \in \langle t_0, t_1 \rangle$ to the length of the interval:

$$q(t_a, t_b) = \frac{1}{t_b - t_a} \int_{t_a}^{t_b} |s'(t) - s(t)| dt. \quad (4)$$

In our case, $q = q(0, T)$ can be directly calculated from the values t_i stored in the outlier set \mathcal{O} by

$$q = \frac{1}{T} \sum_{k=0}^{|\mathcal{O}|/2-1} (t_{2k+1} - t_{2k}). \quad (5)$$

Suppose that the spectrum \mathcal{P} is estimated for $s(t)$ within interval $\langle t_c, t_d \rangle$ and q is calculated for an interval $\langle t_a, t_b \rangle$. If $t_c \leq t_a$, then $q(t_a, t_b)$ relates to the accuracy of model prediction and if $\langle t_a, t_b \rangle \in \langle t_c, t_d \rangle$, then q relates to the accuracy of reconstruction.

To estimate the dependence of the accuracy of reconstruction q_r and prediction q_p on the number of model parameters, we built a spectral model of the one-week-long ‘Office’ dataset. The accuracy of reconstruction q_r was calculated as the difference between the original and reconstructed signal. Moreover, we calculated the accuracies of prediction q_{p1} and q_{p2} for two days of the following week, see Figure 3. The

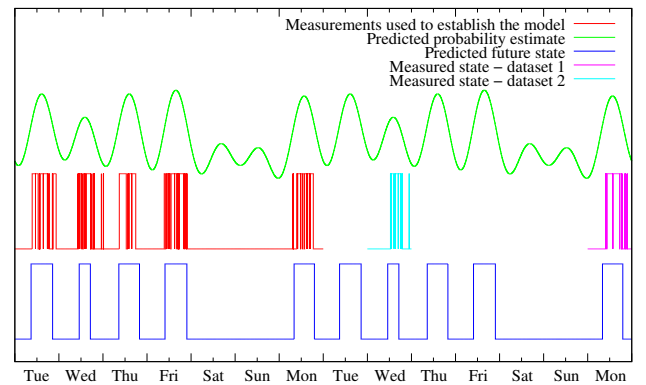


Fig. 3. Comparison of predicted and real values - office dataset.

dependence of the reconstruction and prediction accuracy on the number of parameters of the spectral model is shown in Figure 4. The Figure shows that the spectral model order

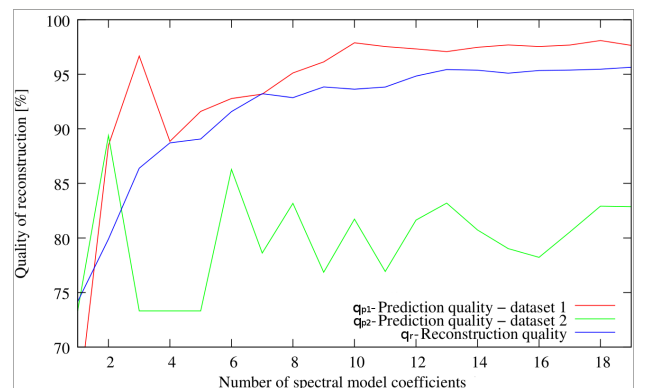


Fig. 4. Model accuracy vs. model complexity - office dataset.

of 15 parameters achieves 95% reconstruction accuracy. As expected, the reconstruction accuracy q_r increases monotonically with the number of spectral model coefficients j , but the prediction quality does not. The local maxima of q_{p1}

and q_{p2} at $l = 2$ and $l = 3$ suggest that for the purpose of prediction, one should use a spectral model of order 3.

The test indicates that the spectral model allows to represent millions of measurements with only a few complex numbers. Thus, the spectral representation \mathcal{P} without the outlier set \mathcal{O} achieves compression ratios in the order of millions while losing less than 5% of information. The full model is composed of 15 triples of spectrum \mathcal{P} and 160 values in the set \mathcal{O} that represent 80 time intervals where the spectral model did not match the measured sequence. Thus, the proposed model achieves lossless compression of the temporal data with a compression ratio reaching $\sim 10^5$.

Since the model that achieved the most accurate signal reconstruction included 15 spectral components, we used this model order for the Lab- n D datasets as well. The reconstruction quality of the ‘Office’ dataset was 0.95 and the corresponding values for the ‘Lab’ datasets are shown in Table I.

TABLE I
RECONSTRUCTION QUALITY FOR DIFFERENT DATASETS

Dataset	Lab-1D	Lab-2D	Lab-3D
Location 1	0.95	0.99	0.99
Location 2	0.99	0.97	1.00
Location 3	0.96	0.98	0.99

The data indicate that the model size remains more or less constant regardless of the time span it covers. Therefore, higher compression rates are achieved simply by representing larger datasets. On average, the proposed model represented the environment state with 98% accuracy while stationary models achieved 95% accuracy.

B. Anomaly detection

An anomalous situation can be defined as a local state of the world which deviates from the internal world model of the robot. Since our model can predict the state $s(t)$ with a given confidence value by Equation 3, we can assume that a measurement $s^m(t)$ is anomalous with confidence level c if

$$s^m(t) \neq (\mathcal{F}'(P) > c). \quad (6)$$

While setting c too high would lower the algorithm’s sensitivity to anomalies, a low value of c would result in an increased number of false positives. An optimal confidence level c can be calculated from the statistical properties of $p(t)$ and the requirements for the number of false positives and failed detections.

Figure 5 shows the results of anomaly detection for the Lab-1D dataset. The anomaly detection was performed “post-hoc”, i.e. we first build the model from the entire dataset and then applied Equation 6. The confidence level c was set to 90% and the anomalous situations correspond to the room being accessed at night or a sudden absence of all people just before and after a meeting, see Figures 6 and 2.

These examples demonstrate how the model adapts its inner dynamics to represent the observed environment. The

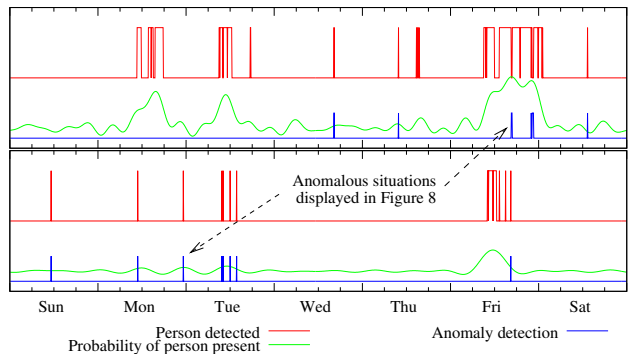


Fig. 5. Anomaly detection for two locations of the Lab-1D dataset.



Fig. 6. Anomalous situations in the ‘Lab’ environment. Left: Workplace empty on Friday early afternoon. Right: Person entering the room at night.

differences between the inner dynamics and real observations allow for natural detection of anomalous situations with arbitrary confidence levels. Figure 5 shows how the spectral models for each laboratory location adapt to the particular location’s dynamics. This results in a different anomaly detector for each location. For example, the present of a person on Monday morning is considered normal for location 1, but triggers the anomaly detector on locations 2 and 3.

C. Building spatiotemporal maps for mobile robotics

The quality of the aforementioned datasets was not significantly affected by the sensor position, because either the sensor was static or the measurement did not require perfect sensor localization. However, metric-map-building algorithms require precise localization of the sensor. While localization can be solved by means of known SLAM methods, the position estimate is never absolutely error-proof or perfectly accurate. Therefore, localization glitches and inherent sensor noise cause degradation of metric maps over time. This degradation can be countered by means of Bayesian filtering methods. Since our approach is not explicitly designed to suppress sensor noise, one might assume that it would not deal with the aforementioned issues.

However, from the theoretical point of view, taking into account only the most prominent spectral components in the update step (see section II-B) is equivalent to sensor data filtering. Thus, map errors caused by noisy measurements should fade out similarly as with the classical mapping approaches. To emphasize new measurements, the update step introduced in Section II-B might be performed for a limited slice of the data gathered causing the model to “forget”

observations that were made too long time ago. However, such a forgetting scheme is beyond the scope of this paper. To













	RGB Image	Point Cloud	Measured Grid	Estimated Grid
Day 1				
Day 2				
Day 6				

Fig. 7. Reconstructed and real environment state at different times.

demonstrate that our representation can deal with the noisy measurements of the environment states $s_j(t)$, we used the 3D and 2D measurements to build 3D and 2D occupancy grids, see Section III. To suppress the effects of imprecise localization, the point clouds gathered by the range sensors were aligned by means of the ICP algorithm prior to the occupancy grid calculation. The individual grid maps were aligned over each other. Therefore we were able to track how the occupancy of the individual grid cells changes over time. Since the cells are considered independent, each cell of the aforementioned grids maintains its own spectral model built from approximately 2000 measurements gathered regularly during one week. Figure 7 provides a detail of the 3D occupancy grid at location 2 of the ‘Lab’ dataset, showing the estimated probabilities and measured states of individual cells at different times.

The accuracy of reconstruction of each cell’s spectral model was calculated, see Table I. The results indicate that the accuracy and compression ratio of the spectral model was not significantly affected by the imperfections of the robot localization.

V. CONCLUSION

A novel approach for spatiotemporal mapping in the context of mobile robotics has been presented. The approach is based on an assumption that in a mid-term perspective, the environment is influenced by several processes which might be periodical and that the evolution of the environment 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 the direct, respectively inverse Fourier transform. The core of the proposed temporal representation is composed of the most prominent frequency components of the Fourier spectrum – these relate to the most important periodical processes influencing the environment.

To evaluate the performance of the proposed method in a real scenario, we have applied it to data gathered by a mobile robot continuously patrolling in an indoor environment for a

period of one week while detecting people and building 2- and 3-D occupancy grids of designated locations.

The results indicate that the proposed method allows to represent arbitrary timescales with constant (and low) memory requirements, achieving compression rates up to 10^6 . Moreover, we demonstrate that the representation allows prediction of future states of the environment with accuracies ranging from 88% to 99%, which enables easy detection of unexpected (anomalous) environment states.

In the future, we will study the utility of the spectral model for mobile robot localization, path planning, information-driven spatio-temporal exploration and semantic segmentation. We will examine the method’s accuracy in long-term scenarios and extend the spectral representation by modelling the uncertainty of its long-term predictions.

ACKNOWLEDGMENTS

The work has been supported by the EU ICT project 600623 ‘STRANDS’. We thank Miroslav Kulich and Pablo Urcola for valuable remarks regarding probability theory.

REFERENCES

- [1] S. Thrun, “Robotic mapping: A survey,” *Exploring artificial intelligence in the new millennium*, pp. 1–35, 2002.
- [2] D. Hähnel, D. Schulz, and W. Burgard, “Mobile robot mapping in populated environments,” *Advanced Robotics*, 2003.
- [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] C. Wang *et al.*, “Simultaneous localization, mapping and moving object tracking,” *International Journal of Robotics Research*, 2007.
- [6] D. Migliore *et al.*, “Use a single camera for simultaneous localization and mapping with mobile object tracking in dynamic environments,” in *ICRA Workshop on Safe navigation in dynamic environments*, 2009.
- [7] P. Biber and T. Duckett, “Experimental analysis of sample-based maps for long-term slam,” *International Journal of Robotics Research*, 2009.
- [8] M. Milford and G. Wyeth, “Persistent navigation and mapping using a biologically inspired SLAM system,” *The International Journal of Robotics Research*, vol. 29, no. 9, pp. 1131–1153, 2010.
- [9] K. Konolige and J. Bowman, “Towards lifelong visual maps,” in *International Conference on Intelligent Robots and Systems*, 2009.
- [10] S. Hochdorfer and C. Schlegel, “Towards a robust visual SLAM approach,” in *International Conference on Advanced Robotics*, 2009.
- [11] 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.
- [12] C. Stachniss and W. Burgard, “Mobile robot mapping and localization in non-static environments,” in *Conf. on Artificial Intelligence*, 2005.
- [13] N. Mitsou and C. Tzafestas, “Temporal occupancy grid for mobile robot dynamic environment mapping,” in *Mediterranean Conference on Control Automation*, 2007.
- [14] M. Yguel *et al.*, “Wavelet occupancy grids: a method for compact map building,” in *Field and Service Robotics*, 2006.
- [15] T. Kucner *et al.*, “Conditional transition maps: Learning motion patterns in dynamic environments,” in *Int. Robots and Systems*, 2013.
- [16] M. Pfingsthorn *et al.*, “Maximum likelihood mapping with spectral image registration,” in *Int. Conf. on Robotics and Automation*, 2010.
- [17] A. Makadia *et al.*, “Fully automatic registration of 3d point clouds,” in *Computer Vision and Pattern Recognition*, 2006.
- [18] P. Neubert *et al.*, “Appearance change prediction for long-term navigation across seasons,” in *European Conf. on Mobile Robotics*, 2013.
- [19] U. Frese and L. Schroder, “Closing a million-landmarks loop,” in *Intelligent Robots and Systems*, 2006.
- [20] R. N. Bracewell and R. Bracewell, *The Fourier transform and its applications*. McGraw-Hill New York, 1986, vol. 31999.
- [21] T. Krajník, M. Nitsche *et al.*, “Visual Localization System for Mobile Robotics,” in *International Conference on Advanced Robotics*, 2013.
- [22] “Linda’s status.” [Online]. Available: <http://twitter.com/LindaStrands>