Long–Term Topological Localisation for Service Robots in Dynamic Environments using Spectral Maps

Tomáš Krajník[1] Jaime P. Fentanes[1] Oscar M. Mozos[1] Tom Duckett[1] Johan Ekekrantz[2] Marc Hanheide[1]

***Abstract*— This paper presents a new approach for topological localisation of service robots in dynamic indoor environments. In contrast to typical localisation approaches that rely mainly on static parts of the environment, our approach makes explicit use of information about changes by learning and modelling the spatio-temporal dynamics of the environment where the robot is acting. The proposed spatio-temporal world model is able to predict environmental changes in time, allowing the robot to improve its localisation capabilities during long-term operations in populated environments. To investigate the proposed approach, we have enabled a mobile robot to autonomously patrol a populated environment over a period of one week while building the proposed model representation. We demonstrate that the experience learned during one week is applicable for topological localization even after a hiatus of three months by showing that the localization error rate is significantly lower compared to static environment representations.**

***Index Terms*— topological localisation, mobile robotics, spatio-temporal representations**

I. Introduction

Self-localisation is a fundamental capability for service robots working in indoor environments. In particular, topological localisation is specifically relevant in the context of large-scale global localization [1], loop-closure detection [2], and for solving the kidnapping [3] problem in mobile robotics.

Typical approaches have made use of static representations of the word to solve the robot localisation problem. However, service robots have to operate in populated environments that are subject to ongoing changes. Moreover, long term operation of mobile robots in changing environments has become a major focus of current robotics research.

In this paper we present a new approach for topological localisation that makes use of information about the dynamics of the environment to improve the localisation process. We assume that many of the changes that occur in populated environments are caused by humans performing their usual daily activities. Several of these activities follow typical patterns and therefore can be exploited by service robots to build more robust representations of their surroundings. Thus, our approach models local states of an environment by means of a probability function of time, which is a superposition of periodical functions that represent recurrent environmental changes [4]. This spectral representation of the time domain allows the robot to identify, analyse and remember regularly occurring environment processes in a computationally efficient way. Although our approach is not applicable in the general case because not all environments exhibit periodic changes, there are cases (e.g. offices or laboratories) where the assumption of periodicity is valid.

[1]Lincoln Centre for Autonomous Systems, University of Lincoln, United Kingdom tkrajnik@lincoln.ac.uk

[2]Computer Vision and Active Perception Lab, The Royal Institute of Technology (KTH), Sweden

Fig. 1. Observations of the same location on different days and times.

In our approach the robot that has to localise itself first makes use of our spectral representation to predict the state of its surroundings at that specific point in time, thus including the changes in the environment that are specific to that time. For example, the predicted representation of an office environment at lunch time will contain less people than usual, because workers are usually in the restaurant at that time. In the same way, fewer people will appear inside the office at night. An example of this situation is shown in Fig. 1. Using this approach, the observations taken by the robot at some particular time will better match the predicted representation of its surroundings at that time, thus improving the match with the model and reducing the error in the localisation.

In this paper we address the problem of topological localisation, in which a service robot has to classify its current position into a set of pre-defined locations inside a specific environment. This process is performed by matching the current robot observation to a previous one that was taken at a different time in (approximately) the same location. In our experiments we select a set of pre-defined locations of a different type (student and staff workplaces, kitchen and meeting area) inside an office environment (see Figure 3) and then let the robot localise itself into one of these locations. The changes in the surroundings of each pre-defined location are learned over one week and modelled using our spectral representation. At the time of localisation, the robot tries to match its current observation to the predicted representations of the surroundings of each location for that specific time. The experiments presented in this paper show significant improvements in the robustness of localisation when using our spectral representations in comparison to static representations of the environment.

II. Related work

There have been different approaches to include changes in the model of the environment, for example, by removing moving objects from the representation of the environment [5], [6] or by tracking these objects and classifying them as moving landmarks [7], [8], [9]. In [10] 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 the current readings. 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 [11], [12]. The paper [13] presents a feature persistence system based on temporal stability in sparse vision-based maps.

For the specific case of topological localisation, algorithms that are based on visual appearance have been shown to be a good choice for image matching and place recognition. [14] shows how SIFT and SURF can be used for robust place identification, [15] showed how using these descriptors along with an epipolar distance it is possible to robustly localise a robot in the long term within an outdoor scenario. [16] combines these image feature techniques to reduce the amount of information to be compared with a CRF algorithm to evaluate the matching rates between different images.

Other authors propose adaptive techniques. [17] presents a method on which the robot adapts its environment model every time it visits a place finding those features that are more stable and "forgetting" those that are less useful. [18] proposes a method in which a dynamic occupancy grid is used that distinguishes between highly dynamical objects, objects that can be moved around and objects that are static.

Other authors use more than one image for place identification. [19] use sequential images to identify places in an across-seasons outdoor scenario and [20] use an "experience" based approach, which uses a set of images acquired at different times of the day to identify the same place. Finally [21] presents a method that predicts appearance changes based on a dataset learned across seasons.

III. Spectral representation of temporal environment domains

Many environment models used in the robotic mapping domain describe the environment by a set of independent components that can assume two distinct states. Examples are occupancy grids with cells that are either occupied of free, topological maps with traversable or non-traversable edges, or landmark maps with occluded or visible features. Typically, the state of each model component is not known exactly due to measurement errors introduced by sensor noise. The uncertainty in the state estimate of the j^{th} world model component is usually expressed by its associated probability p_j. This representation allows to take into account the effect of uncertainty in sensory measurements by employing statistical methods, such as Bayesian filtering. However, most of the world models assume that the p_j of the world components are constant, i.e. they represent the world by a static structure. As a result, the traditional world modelling methods are best suited for representation of slowly changing environments over short periods time.

We propose to consider p_j as a function of time, i.e. to represent the uncertainty of the j^{th} state component as a function $p_j(t)$. However, storing a complete timeline of each $p_j(t)$ is infeasible since complex environments are represented by a very large number of distinct components and such an approach would face memory limitations. Moreover, one of the main reasons to represent the world states as functions of time is to allow for prediction of the environment's future state.

We assume that in mid- and long-term perspectives, most variations of the environment are caused by a finite number of unknown periodic processes. From a mathematical point of view, we propose to represent the $p_j(t)$ as a superposition of sinusoidal functions of unknown frequencies, time offsets and influence. The properties of the periodical processes can be identified by means of frequency transforms, namely the Fourier transform [22].

Although one can argue that the proposed representation is not applicable because many of the world's dynamics are not periodic in nature, we assume that at least some portion of the world's dynamics exhibit certain periodicities and identification of these processes would help to model the environment in a more accurate way. The preliminary experiments presented in [4] suggest that the proposed spectral model reflects the world more faithfully than the traditional static environment models. In particular, [4] shows that representing the world's states as a superposition of 3-15 periodic processes allows to reconstruct (week–long) dynamics of office environments with accuracies of between 90% and 99%. Moreover, the experiments presented in [4] showed that by learning the spectral model parameters on a week-long dataset allows to predict the environment appearance on the following week with ∼90% accuracy. These experiments provide a strong evidence that the mid-term (weeks to months) appearance variations of typical indoor environments are periodic in nature. The main advantage of the representation proposed is its good scalability in terms of time – the memory requirements of the model depend on the number of modelled processes rather than on the represented time period.

A. An Introduction to the Fourier Transform

The Fourier Transform (FT) is a popular mathematical tool with applications mainly in signal and image processing. It transforms a real function of time $f(t)$, into a complex function of frequency $F(\omega)$, which is called a frequency spectrum. The Fourier transform is invertible, and allows to recover the function $f(t)$ from its spectrum $F(\omega)$ and vice versa. The spectrum $F(\omega)$ represents a superposition of sinusoidal functions with amplitudes and phase shifts being equal to $abs(F(\omega))$ and $arg(F(\omega))$ respectively. An important property of the Fourier transform is that the spectrum $F(\omega)$ of a periodic function $f(t)$ is discrete in terms of ω. Considering that in our case, $f(t)$ is also a real discrete function, the spectrum $F(\omega)$ consists of a finite set of complex numbers. For more details on the Fourier transform, refer to [22].

B. The spectral model

The proposed temporal extension applies to world models that represent the environment as a set of independent components, which can be in two distinct states that we denote as 0 and 1. Let us assume that the uncertainty of each state $s_j = \{0, 1\}$ can be represented by its probability of being 1, i.e. $p_j = P(s_j = 1)$. Now assume that s_j is not static, but a function of time $s_j(t)$ that is affected by a set of hidden periodical processes that can be identified by the Fourier Transform. Since the state of individual world components is assumed to be independent, the use of the Fourier transform can be explained on the state $s(t)$ of a single world component.

1) Spectral model representation: Let the temporal sequence of binary states be denoted as $s(t)$. First, we calculate the frequency spectrum of the sequence by means of a Fourier Transform as $S(\omega) = FT(s(t))$. Since we assume that $s(t)$ is periodic, the frequency spectrum $S(\omega)$ is discrete and finite. Therefore, we can select the n most prominent (i.e. of highest absolute value) coefficients S_i of the spectrum $S(\omega)$ and store them along with their frequencies ω_i in the spectral model $\mathcal{P}$. The stored coefficients can be used to recover the smoothed sequence $\tilde{s}(t)$ of $s(t)$ by means of the Inverse Fourier Transform of the model stored in $\mathcal{P}$. Substituting all negative values of $\tilde{s}(t)$ by zeros and limiting the maximal value of $\tilde{s}(t)$ to 1 gives us a function $p(t)$ that can be considered as a probability estimate of $s(t)$. Thus, thresholding the probability $p(t)$ allows us to calculate an estimate $s'(t)$ of the original state $s(t)$. To allow lossless representation of the original signal, the differences between $s'(t)$ and $s(t)$ are stored in an outlier set $\mathcal{O}$ that is Δ-encoded, see Figure 2.

Thus, our model of the state consists of a finite set $\mathcal{P}$ representing the periodic processes and an outlier set $\mathcal{O}$. The set $\mathcal{P}$ consists of n triples $abs(P_i)$, $arg(P_i)$ and ω_i, which describe the amplitudes, phase shifts and frequencies of the model spectrum. Each such triple is an estimate of the importance, time offset and periodicity of one particular periodical process influencing the state $s(t)$. The number of modeled processes n (i.e. the number of triples in $\mathcal{P}$) defines the 'order' of the spectral model. The outlier set $\mathcal{O}$ contains instances when the state $s(t)$ does not match the state $s'(t)$ calculated as $p(t) > 0.5$. The set $\mathcal{O}$ is implemented as a sequence of values, indicating the starts and ends of time intervals when the observed and predicted states do not match, i.e. $s'(t) \neq s(t)$.

2) Spectral model operations: To be able to create, maintain and use this representation, we define four operations: state estimation, state reconstruction, measurement addition and model update.

a) State estimation: The estimation of the state $s'(t)$ from the spectral model allows us to interpolate or even predict the model's state $s(t)$ by the following equation:

$$s'(t) = p(t) > 0.5 = \varsigma(IFT(\mathcal{P})) > 0.5, \qquad (1)$$

where $\varsigma(x)$ is a saturation function that limits the values of x to the interval $< 0, 1 >$. The idea behind this equation is to reconstruct the probability $p(t)$ from the spectrum $\mathcal{P}$ and set the state estimate $s'(t)$ to 1 if $p(t)$ exceeds 0.5.

b) State reconstruction: The function of $s(t)$ is not composed only of periodic processes and $s'(t)$ might not be equal to $s(t)$. To allow lossless representation of the function $s(t)$, we employ the outlier set $\mathcal{O}$ and reconstruct the original state $s(t)$ by means of the following equation:

$$s(t) = s'(t) \oplus (t \notin \mathcal{O}) = (IFT(\mathcal{P}) > 0.5) \oplus (t \notin \mathcal{O}), \quad (2)$$

where $\oplus$ represents a binary XOR operation. The equation first estimates $s'(t)$ from the spectrum $\mathcal{P}$ by means of equation 1 and then inverts the result if t belongs to the set of outliers $\mathcal{O}$.

c) Measurement addition: When a new observation of a real state $s^m(t)$ is obtained at time t, we calculate $s(t)$ by means of Equation (2) and if it differs from $s^m(t)$, the current time t is added to the set $\mathcal{O}$:

$$s^m(t) \neq ((IFT(\mathcal{P}) > 0.5) \oplus (t \notin \mathcal{O})) \rightarrow \mathcal{O} = \mathcal{O} \cup t. \quad (3)$$

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

d) Model update: To update the spectral model, we reconstruct the state $s(t)$ in the desired time interval $< t_{start}, t_{end} >$ and calculate its spectrum $\mathcal{P}$. Again, we select the n coefficients with highest absolute values $|P_i|$ and reconstruct the outlier set $\mathcal{O}$ by means of Equation 3. 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 n can be changed prior to the update step without causing any loss of information. Due to the fact that the Fourier Transform is a well-established mathematical tool and its implementations are optimized for speed, the model update is computationally efficient. Calculating the frequency spectrum of an $s(t)$ with 1000000 samples takes less then 100 ms on an entry level PC.

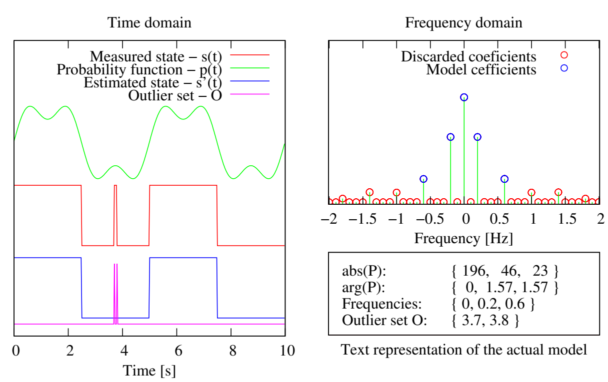

Fig. 2. An example of the measured state and its spectral model. The left part shows the time series of the measured state $s(t)$, probability estimate $p(t)$, predicted state $s'(t)$ and outlier set $\mathcal{O}$. The upper right part shows the absolute values of the frequency spectrum of $s(t)$ and indicates the spectral coefficients, which are included in the model. The last part is a text representation of the model itself.

To further illustrate the spectral representation, Figure 2 provides an example of a third-order spectral model of a quasi-periodic function.

C. Spectrum and periodical processes

In the previous article [4], we have examined the tractability of using the Fourier transform as a core tool for temporal domain representation of traditional world models for mobile robotics. We have shown that storing information about the three most prominent periodic processes allows to represent and predict the environment's appearance with $\sim$90% accuracy over several weeks while the static environment representations achieve accuracies of about $\sim$70%. Since these three processes described by nine real numbers represent several thousand independent measurements of the environment's state, the efficiency of the representation in terms of compression rate exceeds 1:1000. However, the proposed method achieves a lower compression rate than traditional environment models that describe each world's state by a single real number indicating the state's probability. Therefore, there is a tradeoff between the models' accuracy and memory efficiency. Although [4] has shown that the dynamic world model is more accurate, it did not investigate if using the proposed representation provides an advantage in classical problems of mobile robotics like path planning or self-localization. In this paper, we investigate whether introduction of the proposed temporal representation provides an advantage for mobile robot localization in changing environments. We apply the proposed approach to two types of environment models gathered by a mobile robot, which was operating continuously in a human-populated indoor environment for several months. In particular, we build spatiotemporal representations of eight different places in an office environment and let the robot identify its location by comparing the place description to the current observation.

IV. Topological Localisation for Long Term Operation

In this paper we represent the working environment of our service robot using a topological map, which is composed of a number of pre-defined locations. In particular we use eight different locations in our office as depicted in Fig. 3. Each topological place is assigned a set of observations taken at that specific place by the service robot during a certain period of time. In our case, the service robot took an observation every ten minutes during one week of non-stop operations. Each place is then composed of approximately 8000 observations. Our observations are composed of image and point clouds obtained using an RGB-D sensor. Each sensor modality was used to learn a local spectral map that represented the surroundings of each location during one week. In particular, we create a first spectral map represented by 3D occupancy grids obtained by the 3D point clouds; and a second spectral map using visual descriptors from the captured images.

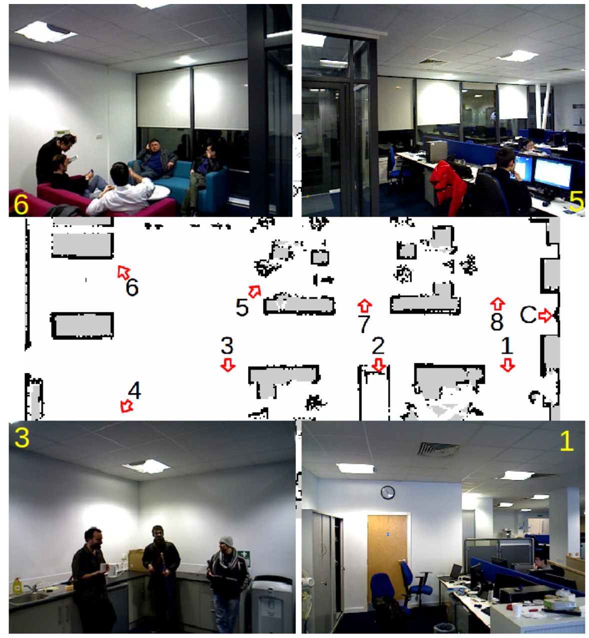

Fig. 3. Topology of the working environment with eight pre-defined topological locations and example observations of four locations.

A. 3D Occupancy Grids

A 3-dimensional occupancy grid of each topological location was built from the gathered point clouds. Since the robot's position information is never absolutely precise, the snapshots of the monitored places were taken from slightly different viewpoints. To deal with variances of robot positions between the different visits, the point clouds gathered by the robot at each topological location are aligned by an adaptive iterative closest keypoint method described in [23]. Since the registered point clouds share a common coordinate frame, we can partition the perceived scene in a 3-dimensional occupancy grid and consider the occupancy of each grid's cell as a function of time. Considering that the dimensions of the grid are given and the occupancy of individual grid cells is considered to be independent, we can represent each location $l \in \mathcal{L}$ by an ordered set of states $\mathcal{S}_l$, where each state $s^l_j \in \mathcal{S}_l$ is associated with a spectral model gathered over the whole week. The spatio-temporal occupancy grid representation of place l thus consists of an ordered set of state evolution functions $\mathcal{S}_l(t)$, where each state representing the occupancy of a particular cell contains its own spectral model.

In the localisation step, we first predict the $\mathcal{S}_l(t)$ of the individual places for the particular time and estimate the occupancy grids. Then, we align the currently perceived point cloud with each grid and use the aligned point cloud to calculate an ordered set $\mathcal{V}$ that represents an occupancy grid of the place the robot is located in. After that, we calculate the similarity of $\mathcal{V}$ to $\mathcal{S}_l(t)$ by means of Hamming distance and assume that the robot is at location k where k = $argmin|\mathcal{V}, \mathcal{S}_l|$.

B. Visual Descriptors

The spectral models for images are created using the BRIEF [24] algorithm, which was evaluated as a best performing image feature extractor in scenarios of long-term localization [25]. The extracted BRIEF features have been used to build a description of a particular place in a similar way as described in [26]. In this approach, the set of visual features $\mathcal{F}_l$ describing a particular place l can be built incrementally from a sequence of sets $\mathcal{V}(t)$ representing the features detected from an image taken at time t. The description of each feature $f_j \in \mathcal{F}_l$ consists of its image coordinates u_j, v_j, descriptor vector $\mathbf{e_j}$ and a binary vector $\mathbf{s_j}(t)$ indicating the presence of the feature f_j in a set $\mathcal{V}(t)$. Each time an image is processed by the BRIEF algorithm, the set $\mathcal{V}(t)$ is populated and tentative correspondences between the sets $\mathcal{V}(t)$ and $\mathcal{F}_l$ are established. These correspondences are filtered by epipolar geometry constraints and the vectors $\mathbf{s_j}(t)$ of the associated features in the set $\mathcal{F}_l$ are updated. The non-associated features in the set $\mathcal{V}(t)$ are added to the set $\mathcal{F}_l$. Once the entire dataset consisting of eight places is processed, we have eight sets $\mathbf{F}_l$ containing the image feature descriptors along with their positions and functions of their occurrence in the particular images. Thus, we can reconstruct eight sets of image features that are likely to be perceived at the eight places at a particular time t.

To determine the location of the robot by means of image feature extraction and matching, we first extract the salient points of the current location's image by the CenSurE detector [27]. Then the descriptor of the point's neighbourhood is calculated by means of the BRIEF [24] algorithm. These positions and descriptors form a set $\mathbf{V}(t)$ that describes the robot's perception of the current location. Then, eight sets $\mathcal{F}_l$ that contain the features likely to be seen at the various locations at the particular time t are created from the spectral environment models. Tentative correspondences between the sets $\mathcal{F}_l$ and the set of currently perceived features $\mathcal{V}$ are established and filtered by means of epipolar geometry. The number of correctly established correspondences is considered as a similarity measure of the l^{th} place to the current view. Again, the robot is considered to be located at a place that is most similar to the captured image, i.e. the place with the highest amount of corresponding features.

V. Experimental Evaluation

To evaluate the proposed spectral map extension, we have applied it to the problem of topological localization in long-term robotic scenarios. Our approach requires that the environment observations are taken frequently and regularly over a long period of time. Since none of the publicly available datasets fulfills these requirements, we have collected the experimental data ourselves. Our experiments have been carried out in the open-plan office of the Lincoln Centre for Autonomous Systems (UK). The experimental platform used was the SCITOS-G5 mobile robot (see Figure 4) equipped with an RGB-D camera and a laser rangefinder.

The autonomous patrolling behaviour was based on combination of the improved ROS navigation stack and the visual localization method proposed in [28]. The robot reported its status regularly using a social network interface, so its occasional failures could be dealt with immediately. While the robot's SICK laser rangefinder was used primarily for autonomous navigation, obstacle avoidance and localization, the primary sensor used to gather the long-term data was the Asus Xtion RGB-D camera, which was placed on the robot's head.

Fig. 4. The SCITOS-G5 robot patrolling the LCAS Witham Wharf office and example observations with extracted BRIEF features.

The robot was programmed to capture 3D point clouds and RGB images of eight designated areas (see Figure 3) of the office every ten minutes. Since it has been continuously patrolling for a week (the second week of November 2013), the entire training dataset consists of approximately 8000 point clouds and 8000 images. Representative examples of the captured images are shown in Figure 3. During the dataset collection, the RGB-D raw data were used to build the different models of the eight monitored topological locations. These models were based on coarse (cell size 1 m) 3D occupancy grids, and visual features as introduced in Sec. IV. The spectral representations for images and point clouds were obtained using 8000 observations corresponding to one week of robot operation during which the robot travelled over 35 km fully autonomously.

A. Results

To test the topological localisation we used a set of $\sim$1000 observations corresponding to 24 hours of a day following the week used to learn the spectral models on November 2013. This would correspond to a situation when a robot attempts to localize itself while having an up-to-date spatio-temporal model of the environment. The second testing dataset consists of approximately $\sim$1000 observations gathered during the 24 hours of a day in early February 2014. This test is aimed at the long-term predictive ability of the spatio-temporal models, because the models learned during the week in November are used to localize the robot after more than two months.

Each of the $\sim$1000 observations corresponding to each day in November and February was matched against the predicted representations obtained for that specific time from the spectral models in the different modalities. The matching errors are presented in Table I. An error occurs when the robot matches its observation with one from a wrong topological location. The experimental results in Table I

TABLE I
OVERALL LOCALIZATION ERROR (%)

Model type	Model order	Image features Nov	Image features Feb	Occupancy grids Nov	Occupancy grids Feb
static	-	35%	45%	21%	17%
spectral	1	25%	26%	14%	13%
spectral	2	22%	27%	14%	8%
spectral	3	18%	24%	14%	17%
spectral	4	17%	29%	7%	17%

indicate that modelling the environment by our approach reduces the localization error. Moreover, the results show that while the predictive capability of high-order temporal models is better in short-term horizons (November dataset), models that include one or two periodic processes perform better in the long term (February dataset). The most important fact is that while the error rate of the static world models increased in the long term, the dynamic models of lower orders perform similarly. One can also see that increasing the number of modelled processes is beneficial only for short-term predictions because only the most significant processes are persistent over long-time periods.

VI. CONCLUSION

A novel approach for mobile robot localization in changing environments has been presented. The approach is based on spatio-temporal mapping in the context of mobile robotics. We assume that in a mid- to long-term perspective, the environment's appearance is affected by a set of hidden, periodic processes. We assume that the dynamics of the environment can be described by the frequency, amplitude and time shift of these processes.

To identify the parameters of these processes and to predict the environment's local state we use the direct and 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 applicability of the method for mobile robot localization in changing environments, we enabled the robot to learn about the environment dynamics by autonomously patrolling an office environment for a period of one week, during which the robot built two types of spatio-temporal models of eight office locations with different dynamics. The applicability of the spatio-temporal models learned has been tested in a topological localization scenario, where the robot had to estimate its location based on its current observation and the spatio-temporal models gathered. Short- and long- term scenarios were considered. In the first one, the robot had to recognize its position during a 24 hour operation on the day following the week the models were created. In the long-term test, we repeated the procedure after three months with the same spatio-temporal models.

The results show that the proposed approach significantly increased the localization success rate compared to the static models, indicating that the knowledge of the assumed periodic processes in the environment helps to explain a significant part of the environment variations observed.

ACKNOWLEDGMENTS

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

REFERENCES

[1] H. Badino, D. Huber, and T. Kanade, "Visual topometric localization," in *IEEE Intelligent Vehicles Symposium (IV)*, 2011.
[2] K. Kosnar, V. Vonasek, M. Kulich, and L. Preucil, "Comparison of shape matching techniques for place recognition," in *ECMR*, 2013.
[3] A. Pronobis *et al.*, "A realistic benchmark for visual indoor place recognition," *Robotics and autonomous systems*, 2010.
[4] T. Krajník, J. P. Fentanes, G. Cielniak, C. Dondrup, and T. Duckett, "Spectral analysis for long-term robotic mapping," in *ICRA*, 2014.
[5] D. Hähnel, D. Schulz, and W. Burgard, "Mobile robot mapping in populated environments," *Advanced Robotics*, 2003.
[6] D. Wolf and G. Sukhatme, "Mobile robot simultaneous localization and mapping in dynamic environments," *Autonomous Robots*, 2005.
[7] L. Montesano, J. Minguez, and L. Montano, "Modeling dynamic scenarios for local sensor-based motion planning," *Aut. Robots*, 2008.
[8] C. C. Wang *et al.*, "Simultaneous localization, mapping and moving object tracking," *International Journal of Robotics Research*, 2007.
[9] D. Migliore *et al.*, "Use a single camera for simultaneous localization and mapping with mobile object tracking in dynamic environments," in *ICRA Workshop on S.N.O.D.E.: App. auton. veh.*, Japan, 2009.
[10] P. Biber and T. Duckett, "Experimental analysis of sample-based maps for long-term SLAM," *The Int. Journal of Robotics Research*, 2009.
[11] M. Milford and G. Wyeth, "Persistent navigation and mapping using a biologically inspired SLAM system," *IJRR*, 2010.
[12] K. Konolige and J. Bowman, "Towards lifelong visual maps," in *Intelligent Robots and Systems, IROS 2009*, St. Louis, October 2009.
[13] 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.
[14] A. Murillo, J. Guerrero, and C. Sagues, "SURF features for efficient robot localization with omnidirectional images," in *ICRA*, 2007.
[15] C. Valgren and A. J. Lilienthal, "SIFT, SURF & seasons: Appearance-based long-term localization in outdoor environments," *RAS*, 2010.
[16] C. Cadena, D. Gálvez-López, J. D. Tardós, and J. Neira, "Robust place recognition with stereo sequences," *IEEE Trans. on Robotics*, 2012.
[17] F. Dayoub and T. Duckett, "An adaptive appearance-based map for long-term topological localization of mobile robots," in *IROS*, 2008.
[18] G. D. Tipaldi, D. Meyer-Delius, and W. Burgard, "Lifelong localization in changing environments," *Int. J. of Robotics Research*, 2013.
[19] M. Milford and W. Fraser, "SeqSLAM: Visual route-based navigation for sunny summer days and stormy winter nights," in *ICRA*, 2012.
[20] W. Churchill and P. Newman, "Practice makes perfect? managing and leveraging visual experiences for lifelong navigation," in *ICRA*, 2012.
[21] P. Neubert, N. Sünderhauf, and P. Protzel, "Appearance change prediction for long-term navigation across seasons," in *ECMR*, 2013.
[22] R. N. Bracewell, *The Fourier transform and its applications*. McGraw-Hill New York, 1986, vol. 31999.
[23] J. Ekekrantz, A. Pronobis, J. Folkesson, and P. Jensfelt, "Adaptive iterative closest keypoint," in *ECMR*, 2013.
[24] M. Calonder, V. Lepetit, C. Strecha, and P. Fua, "BRIEF: binary robust independent elementary features," in *Proceedings of the ICCV*, 2010.
[25] C. Valgren and A. J. Lilienthal, "SIFT, SURF and seasons: Long-term outdoor localization using local features." in *ECMR*, 2007.
[26] T. Krajník, J. Faigl, V. Vonásek, *et al.*, "Simple yet stable bearing-only navigation," *Journal of Field Robotics*, vol. 27, 2010.
[27] M. Agrawal, K. Konolige, and M. Blas, "Censure: Center surround extremas for realtime feature detection and matching," in *ECCV*, 2008.
[28] T. Krajník *et al.*, "A practical multirobot localization system," *Journal of Intelligent and Robotic Systems*, 2014.