Persistent Localization and Life-long Mapping in Changing Environments using the Frequency Map Enhancement

Tomáš Krajník[1] Jaime Pulido Fentanes[1] Marc Hanheide Tom Duckett

***Abstract*— We present a life-long mapping and localisation system that enables long-term autonomous operation of mobile robots in changing environments. The core of the system is a spatio-temporal occupancy grid which explicitly represents persistence and periodicity of the individual cells and can predict the probability of their occupation in the future. During autonomous navigation, our robot builds temporally local maps and integrates then into the global spatio-temporal grid. Through re-observation of the same spatial locations, the spatio-temporal grid obtains information about long-term environment dynamics and gains the ability to predict the future environment states. This predictive ability allows to generate time-specific 2d maps which are used by the robot's localisation and planning modules. By analysing data from a long-term deployment of the robot in a human-populated environment, we show that the proposed spatio-temporal representation improves localisation accuracy and the efficiency of path planning. To allow the use of the method by other roboticists, we show how to integrate it in the ROS navigation stack, which is a de-facto standard in mobile robotics.**

***Index Terms*— mobile robotics, long-term autonomy**

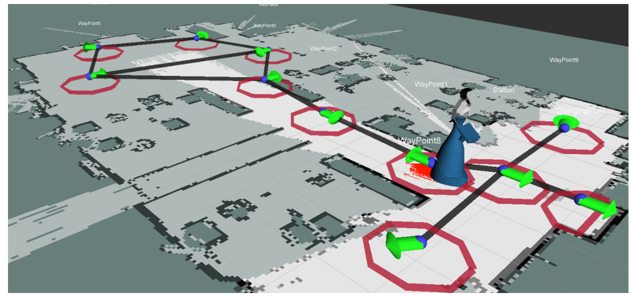

Fig. 1: Predicted and created 2D grids with patrol locations.

Environment change example (cubdoard doors open/closed)

Fig. 2: Example of the regular changes and corresponding map sections predicted for early morning (left) and late afternoon (right).

I. Introduction

One of the many challenges that robots have yet to achieve is the ability of long-term autonomous operation in changing environments. This is particularly difficult because the efficiency of mobile robot operation depends heavily on the quality of the available knowledge about the environment.

Many tasks performed by mobile robots take place in environments where humans perform their usual activities, which causes the environments to change constantly. For example, doors are open and shut, chairs are pushed in and out of tables and furniture is occasionally rearranged.

This means in long-term scenarios, having an environment model that assumes a static world will inevitably lead to robot navigation failures as the knowledge base becomes obsolete over time. Having an obsolete map will lead to many different types of failure such as mislocalisations, navigating to areas that are not accessible anymore or ignoring alternative, shorter paths.

Many authors have tried to deal with changing environments by proposing representations that suppress the effect of the changes [1], [2], keep track of different environment states and choose the one closest to the current state [3], [4], or perform constant remapping, keeping up with the latest environment state [5], [6]. Nevertheless these approaches are usually tailored to specific knowledge representation models and system architectures. We propose an architecture for life-long mapping and persistent localization that is easily integrated within the ROS framework [7]. Through extension of the ROS navigation stack by an additional ROS *gmapping* module, we obtain a system that can create an up-to-date, independent environment map on-the-fly. To enforce the compatibility of the new map with the previous environment models, we propose to use the position estimates of the ROS *AMCL* module as virtual odometry for the *gmapping* node.

The up-to-date maps are integrated into a spatio-temporal occupancy grid where each cell contains a frequency-spectrum of its past states [8] and allows the prediction of the cell's future states at particular times. Previous work has shown that the predictive capabilities of the spectral models improve visual-based mobile robot localisation [9], topological navigation [10] and task planning [11]. However, the previous works were aimed at proof-of-concept verification of the Frequency Map Enhancement (FreMEn) methods using custom modules that were not fully integrated in the ROS navigation framework. In this work, we show that once the ROS navigation stack is extended so that it allows the creation of new maps, integration of the Frequency Map Enhancement is straightforward, as it simply replaces the *mapserver* ROS module.

To validate the proposed method, we set up our robot

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

Preprint submitted to 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems. Received March 1, 2016.

to routinely patrol a human populated office environment, which was subject to frequent changes due to the people's activity. During each patrol run, which would start and end at the robot's charging station, a new map was created and integrated into the FreMEn spatio-temporal models. At the start of each patrol run, the spatio-temporal model predicted a time-specific map that was used by the robot's planning and localization methods. As the model accumulated enough data to infer the long-term environment dynamics, the predicted maps started to differ according to the time of the patrols. For example, the spatio-temporal model predicted that on weekday afternoons, doors of certain cupboards are more likely to be open than during nights, see Figure 2.

To evaluate how the predictive capability of our life-long navigation method affects the efficiency of the robot operation, we provide statistics about the accuracy of the robot localization, navigation performance, planning failures and map quality.

Since the idea of Frequency Map Enhancement [8], which allows to introduce the notion of dynamics into most robotic environment models, was introduced before, the contribution of this work lies in different aspects. In particular, we show that the Frequency Map Enhancement can be easily extended to model not only cyclic-periodic changes, but also their persistence. We propose an architectural modification of the ROS navigation stack that allows for straightforward update of the environment models, which enables ROS integration of the proposed method in an elegant and straightforward way. Last, but not least, we demonstrate that the introduction of the method improves the efficiency of the robot operation in long-term scenarios.

II. Related work

Long-term robot navigation is highly dependant on the precision of the world model, since the robot requires such a model to localise itself, plan its trajectory and find obstacle free paths. Traditionally this model of the environment is called a map and creating it for static environments is a problem that has been widely studied for a long time [12]. However, dynamic scenarios, where the world is constantly changing and uncertainty grows with time, still represent a big problem that is subject to research by the robotics community.

Some authors try to handle these dynamics by finding the most static landmarks and filtering out of the model those that change over time [2]. Another approach involves tracking these "moving" landmarks and labelling them as dynamic [13] In general, these approaches can handle some problems of navigating in a dynamic environment, but they cannot deal with long-term changes to the structure of the environment, which translates into a less robust long-term behaviour.

Other approaches never assume the map to be complete and perform continuous mapping, adding new features to the map with every observation [5], [6], [14]. In these approaches, the key problem is managing map size, especially in long term scenarios where the robot might be making new observations for several weeks or months.

To tackle the long-term challenges of robust navigation in dynamic scenarios, some approaches gather and maintain different temporal representations simultaneously and choose the best one according to its consistency with the current perception of the world (e.g. [3], [4]). However, these approaches are costly in terms of memory and computational efficiency.

Considering this computational cost, many authors have tried taking these approaches to a discrete, topological, level where computational requirements are lower. At this level, most representations use visual appearance for place recognition [15], [16], having shown that visual features can be used for robust place identification. However, these approaches present a decrease in robustness when facing long-term changes [17], as again they are prone to error when features appear and disappear over time. In [9], [10] dynamic models of the topological space that explicitly represent the environment changes and try to identify patterns by means of the Fourier transform are presented, for both localisation (node level) and navigation (transitions between nodes). This ability of pattern identification allows for state prediction, which as shown in [10] can improve navigation performance. However, these approaches are still at the topological level and there is no model that can be used for low-level navigation on node transitions, and the proposed model is limited to higher level planning tasks.

Other authors have looked into metric level representations with state prediction abilities. In [18] the authors propose 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 [19] presents a representation of the environment which models transitions of dynamic objects in the environment, by learning motion patterns from the temporal signal of occupancy in a cell.

The approach presented in this paper presents an occupancy grid map where each cell in the map is enhanced with a spectral model [20] that allows for the prediction of the cell's state at specific times. We show how localisation and navigation are improved using such a representation compared to a static occupancy grid map, which is still the most common approach.

III. System description

A major drawback of the previously proposed dynamic mapping techniques is the fact that they are tailored to particular representations that only work for specific system architectures. Unlike these methods, the objective of this work was to develop an environment representation that is general enough to allow its use with most of the environment models used in robotics. Moreover, we wanted the system to be easy to use by other researchers, and therefore, it was implemented as a module that is compatible with the navigation stack [21] of the Robot Operating System (ROS), which is considered a standard in robotics nowadays. The

Preprint submitted to 2016 IEEE/RSJ International Conference
on Intelligent Robots and Systems. Received March 1, 2016.

software we use is freely available asa component of the STRANDS system [22].

In principle, the framework allows to keep several spatio-temporal models that can be used for localisation simultanneously as proposed by [3], [4]. Building a separate map of the current environment layout allows to postpone the decision on which of the global spatio-temporal models will be updated similarly to the experience-based approach proposed by [4]. This would not have been possible with a classic continuous SLAM approach. In this work, we use only one global spatio-temporal representation and if the curretly build map is detected as anomalous it is simply rejected.

A. *Continuous Mapping with the ROS navigation stack*

The traditional configuration of the navigation stack has four main components, see left part of Figure 3: the *robot*, which apart from being the component that interacts with the world, also provides all the necessary sensory input and coordinate system transformations; the *map server* provides the map that will be used by the localisation and planning systems, where this map is usually created in a previous stage; the *AMCL* localisation system, which provides position estimates in the map coordinate frame using the sensory input from the robot; and a *move base* motion planner that uses the map, position estimate and sensor information to plan the robot's motion.

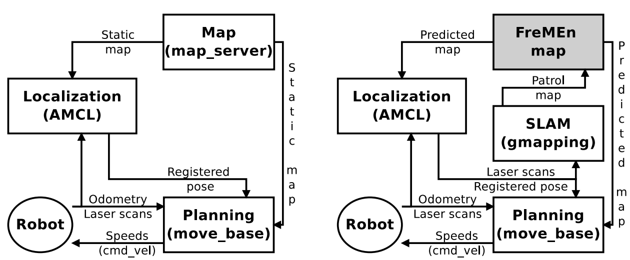

Fig. 3: Classic and proposed navigation stack.

One of the key requirements of long-term operation is the ability to keep the environment models up-to-date. This does not necessarily mean that the robot has to be able to make long-term predictions, but that it should be able to update its map whenever a change occurs. In our architecture, a traditional SLAM-based method is used to create a completely new map every time the robot performs a patrol run and integrate this single map into the proposed spatio-temporal representation.

The ability to update the environment map is achieved through a minor modification of the ROS navigation stack, see right part of Figure 3. The main changes are the continuous mapping component and the spatio-temporal representation, called the *FreMEn map*, which provides time-specific maps to the other system modules. Continuous mapping is achieved by setting up the *gmapping* module to work in parallel with the rest of the *navigation stack* during robot operation and creating a new 'patrol' map, see Figure 3. However, the SLAM-based *gmapping* implementation is subject to a slight localisation drift and running it as a completely separate process would cause its map to diverge from the global 'FreMEn' one, which would make its integration into the FreMEn map impossible. To prevent this drift, we simply inject the output of the *AMCL* position estimation in the *odometry* input of the *gmapping* module. This position injection ensures that the individual cells of the 'patrol' and FreMEn maps correspond closely to each other and that the cell differences are caused by environment changes and not by localisation drift. This makes integration of the 'patrol' map into the global one a straightforward process.

B. *FreMEn 2D Grid*

The idea of the *FreMEn map* is based on the observation that most of the environment states are not changing chaotically and the nature of their dynamics can be learned from their repeated observations. Thus, the *FreMEn map* is a 2d occupancy grid that does not represent the uncertainty of the individual cells by a constant probability, but as a function of time, which is estimated from re-observations of the cell occupancy over long time periods. The *FreMEn map* can integrate local 2d grids created at different times ('patrol maps') into a global spatio-temporal representation that captures not only the spatial environment layout, but also persistence and cyclic behaviour of its changing states.

From an architectural point of view the *FreMEn map* components provides the 2d environment maps in the same way as the original ROS *map server*. However, unlike the *map server*, which can load, save and transmit static maps only, the *FreMEn map* generates time-dependent maps that reflect the expected environment state at the time of robot operation.

Although the model update step can be performed at any time, but the mapping process should be long enough to filter out fluctuations in the environment, and the map update should only happen when the robot is not moving to avoid anomalous maps.

C. *Anomaly detection*

However, continuous mapping is exposed to two significant threats. First, the *gmapping* method can fail and produce an incorrect map, which, when integrated into the *FreMEn map*, might corrupt the entire spatio-temporal representation. The second threat is more subtle: due to the environment changes, sensor noise and localisation inaccuracies, exact registration of the recently-gathered 'patrol' maps with the FreMEn grid is not absolutely precise. This introduces a certain amount of noise every time a new map is integrated into the global one. As the noise accumulates, the global map might become less and less accurate over time, which might lead to its destruction.

Both effects exhibit themselves at the moment when a new map is being integrated into the FreMEn grid. Therefore, the *FreMEn map* checks how probable is the new map, i.e. how much it conforms to the predicted maps for that

Preprint submitted to 2016 IEEE/RSJ International Conference
on Intelligent Robots and Systems. Received March 1, 2016.

particular time. This allows to identify maps that differ significantly from the predicted representations. We assume that the number of changes in our environment is low and reject these maps as outliers. However, the maps could also be kept as alternative representations of the environment as mentioned at the start of this section.

IV. Frequency Map Enhanced occupancy grid

The spatio-temporal representation that forms the core of our approach is a 2d occupancy grid that models the occupancy of each cell by a probabilistic function of time. This probabilistic function consists of two distinct components: persistency and periodicity. The persistency component acts as a short-term memory that represents the expectation that the cell state did not change since the last observation if the observation was performed recently. The periodicity component is related to the fact that from a long-term perspective (days to months), some of the environment states might be influenced by hidden processes that could exhibit certain periodicities.

The idea of identifying periodic patterns of binary environment states via the Fourier Transform and using them for future predictions was originally presented in [20]. Later on, the authors demonstrated that the predictive capability of the proposed representation improves mobile robot localization, planning and exploration [8].

In our version of FreMEn, we represent the periodic behaviour of each cell by its sparse frequency spectrum, which is a set $\mathcal{A}$ of complex numbers α_k. These correspond to the set Ω of modeled periodicities ω_k that might be present in the environment. The persistence of each state is represented by the mean time between state transitions τ and the time and value of the last observation t_l and $s(t_l)$. Moreover, we store each cell's number of observations n and the cell's mean occupancy μ. Each time a cell occupancy $s(t)$ is observed at time t, the aforementioned representation is updated as follows:

$$\begin{array}{lcll} \mu & \leftarrow & \frac{1}{n+1}\left(n\mu + s(t)\right), & \\ \alpha_k & \leftarrow & \frac{1}{n+1}\left(n\alpha_k + s(t)e^{-jt\omega_k}\right) & \forall\, \omega_k \in \Omega, \\ \tau & \leftarrow & \frac{1}{n+1}\left(n\tau + \frac{|s(t)-s(t_l)|}{t-t_l}\right), & \\ s(t_l) & \leftarrow & s(t), & \\ t_l & \leftarrow & t, & \\ n & \leftarrow & n+1. & \end{array} \tag{1}$$

The proposed update step is analogous to incremental averaging – the absolute values of $|\alpha_k|$ actually correspond to the average influence of a periodic process (with a frequency of ω_k) on the values of $s(t)$. Note that the size of the representation of the state (i.e. the number of elements in $\mathcal{A}$) is independent of the number of observations, which means that the memory requirements of the proposed representation do not grow over time.

To predict the value of state $s(t)$ for a future time t, we first sort the set $\mathcal{A}$ descendingly according to the absolute values $|\alpha_k|$. Then, we extract the first m elements α_l along with their corresponding frequencies ω_l and calculate the state's probability over time as

$$p(t) = s(t_l)e^{\frac{t_l-t}{\tau}} + f(t)(1-e^{\frac{t_l-t}{\tau}}), \tag{2}$$

where

$$f(t) = \varsigma(\mu + \sum_{l=1}^{m} |\alpha_l| cos(\omega_k t + \arg(\alpha_l))), \tag{3}$$

where $\varsigma(.)$ ensures that $p(t) \in [0,1]$. Note that for preditions which immediatelly follow the last observation, i.e. $|t-t_l| << \tau$, the expression $e^{\frac{t_l-t}{\tau}}$ is close to 1, which means that the expected occupancy would be the same as the one recently observed. If we use Equation 2 to predict further into the future, i.e. $|t-t_l| >> \tau$, the expression $e^{\frac{t_l-t}{\tau}}$ is close to 0, which supresses the effect of the last observation on $p(t)$ and emphasizes $f(t)$, which represents the behaviour of the modelled cell from a long-term perspective.

Note that the choice of m, which determines how many periodic processes are considered for prediction, and Ω, which determines the periods of the potential cyclic processes, are crucial for the prediction performance. Omitting some periods from the set Ω would prevent the system from capturing the processes with these periodicities, while including too many elements in Ω would cause the model to consume too much memory. Setting m too low results in omitting some environment processes that actually influence the state, while setting m too high includes components of $\mathcal{A}$ that are caused by sensor noise. The discussion about the optimal choice of m and Ω is beyond the scope of this paper. In our case, m was set to 2 and the Ω was selected as in the paper [23], where the choice of m and Ω is explained in detail.

V. Experiments

To evaluate the utility of the proposed dynamic map for long-term deployment of mobile robots, we used data that were gathered during several days of routine autonomous operation of a mobile robot at the Lincoln Centre for Autonomous Systems. The SCITOS-G5 mobile robot (see Figure 4) regularly patrolled a large open-plan office every ten minutes while recording data from its odometry, RGB-D and laser range-finding sensors. Its autonomous navigation was based on the ROS navigation stack, which used our FreMEn 2D grid instead of the traditional map server. To achieve autonomous operation, the robot uses a precise visual servoing method for reliable docking to its charging station [24]. Our evaluation was based on three criteria: localization accuracy, navigation efficiency and map quality. To evaluate the localization accuracy, we covered part of the environment with an external localization system, which provided us with a ground truth of the robot position with millimetre precision. To quantify the efficiency of the robot navigation, we measured the time it took to perform a patrol where the robot had to visit five different locations. We also measured the times it took for the robot to navigate through a narrow area that exhibited regular changes. To assess the

Fig. 4: Overview and 2D occupancy grid of the Witham Wharf office.

Fig. 5: Example image captured by the ceiling-mounted camera of the external localisation system. The position of the circular marker on top of the robot is used as a ground truth to determine the accuracy of the robot's localization system.

quality of the built maps, we quantified the amount of noise in the maps.

A. Localisation accuracy

To evaluate the accuracy of the robot self-localization, we installed an independent localisation infrastructure at the Witham Wharf office. The infrastructure consisted of two ceiling-mounted fish eye Kodak PixPro SP360 cameras, a large circular marker on top of the robot and another set of markers close to the robot's charging station. While the marker on the robot's top was used to determine its x and y position, the markers positioned at the charging station area allowed for precise, independent estimation of the robot heading. Detection and position estimation of the markers, localisation system calibration and coordinate system setup was based on a freely-available, open-source method presented in [24]. To ensure millimetre accuracy of the localisation system, we had to use rather large markers as suggested by the mathematical model of the system [24], see Figure 5. We selected approximately 2000 images in 20 different image sequences captured by the overhead cameras and established the positions of the robot. To avoid potential accuracy drop-off caused by the use of the wide-angle lens cameras, the selected images have the robot position close to the center of the image.

The individual sequences captured the movement of the robot through a 1.5 m wide corridor outlined by eight storage cupboards. These cupboards are used by the research staff of the office and some of the cupboard doors are typically open during the day and closed at night. The cupboards are 0.5 m deep, so when a cupboard door is left open, the corridor appears to be 2 m instead of 1.5 m wide and its center appears to be offset by 0.25 m aside. Thus, when moving through this corridor, the discrepancy of the 2D map with the perceived environment state might negatively affect the accuracy of robot self-localization.

In our case, the 20 m range of the robot laser rangefinder ensures that it will almost always perceive areas that did not change, which should keep the position estimate accurate. However, if the range of the laser sensor was shorter, e.g. when using a Hokyo URG04, then the localization accuracy would be affected severely.

To estimate the impact of the environment change and sensor range on the localization precision, we processed laser, odometry and ground truth data from 20 different passes of the robot through the monitored corridor. To emulate the limited range of the laser rangefinder, we trimmed the laser data at different lengths. Using the trimmed data from 20 different runs, we performed standard ROS-based AMCL localisation on the 'static', 'averaged' and 'predicted' 2d maps and compared the robot positions to the ground truth obtained by the overhead cameras. The results shown in

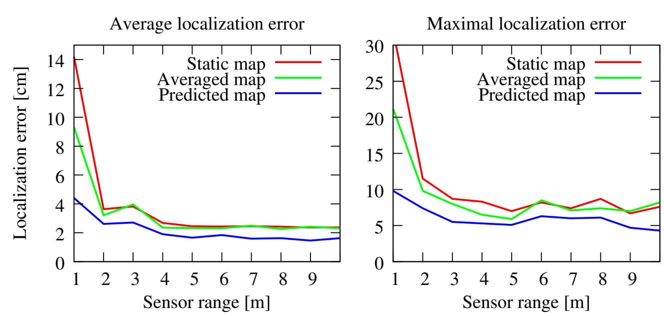

Fig. 6: Localization error for different ranges of the laser scanner and different types of the maps. Predicting a map for a particular time improves localization accuracy, although the improvement is only marginal for long-range sensors.

Figure 6 indicate that use of the time-specific, predicted maps results improves the localization precision in a significant way if the range of the laser rangefinder is lower than the overall map size. If the rangefinder provides a complete overview of the operational environment, the reduction of the position estimation error is only marginal. However, a small difference in localization precision can have a significant impact on the efficiency of the robot navigation and quality of the constructed maps.

B. Navigation efficiency

To evaluate the navigation efficiency, we processed navigation statistics of 60 different patrol runs. During each patrol, the robot undocked from its charging station, visited several different locations in the office (see Figure 2) and returned back to recharge. The data from each patrol run contains the robot's average speed and the number of events where standard navigation behaviour failed and the robot had to perform custom recovery behaviours in order to proceed with its patrol. The gathered navigation statistics were divided into three groups of 20 patrols each. The first group contained patrols that were happening during weekends, where the amount of environment changes in the office is more likely to be low. The second group contained patrols from weekday afternoons, where the robot was using an 'averaged' map, which slowly adapts to the observed change. The third group contained patrols from weekday afternoons, where the robot was using a 'predicted', time specific map. Table I indicates

TABLE I: Navigation statistics

Environment		Static	Changing	
Map		Static	Static	Predicted
Average speed	[$\frac{m}{s}$]	0.21	0.15	0.18
Recovery events	[-]	1	21	12

that in a static environment, the robot could navigate efficiently even when using a static map, but as soon as the environment began to change, the navigation efficiency was affected in a negative way. However, the negative effect of the changes was slightly lowered through the use of the proposed map, which represents the environment changes in an explicit way.

C. Map quality

This experiment evaluates the effect of an anomalous map detection mechanism. This mechanism verifies whether a newly created map conforms with the representation that was gathered so far, which allows to reject corrupt or otherwise incorrect maps. To verify the utility of the anomaly detection mechanism, we replayed laser and odometry data from 100 consecutive patrols with the anomaly detection component being deactivated and compared the resulting spatio-temporal representation with the one built while the anomaly detection was used to rejectd potentially corrupted maps. Figure 7 shows the amount of changes detected in the consecutively created maps. The figure shows that at a certain point (run 36), integration of an incorrect map corrupts the FreMEn grid, which breaks the map update process. However, the

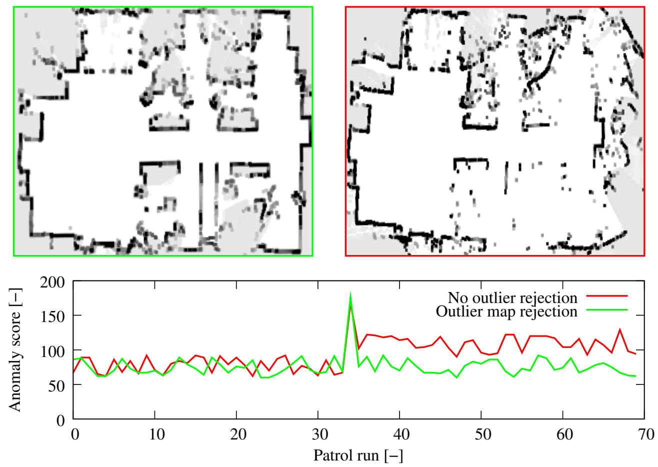

Fig. 7: Effect of the anomalous map detection and rejection. The peak in the graph indicates mapping system failure and consequent discrepancy of the individual maps. The two maps show the resulting models with and without anomaly rejection.

anomalous map rejection mechanism prevents this situation and the map update process continues to produce a faithful 2D environment model.

VI. Conclusion

We presented an approach for mobile robot life-long mapping and persistent localization in changing environments. First, we show that the ability to update the environment model does not require introduction of custom modules to the ROS navigation stack. Instead, the navigation stack can be augmented by the *gmapping* module that builds a new map every time the robot navigates around its operational environment. To ensure that the new map is consistent with the previously built model, we propose to use the *AMCL* module position estimation as virtual odometry for *gmapping*. Second, we demonstrate that maps of the individual navigation runs can be integrated into a spatio-temporal model that captures the persistency and periodicity of the environment changes. This spatio-temporal environment representation, which explicitly models the environment dynamics, is used to predict time-specific maps, which serve our robot both for localization, path-planning and navigation.

Our experimental evaluation, based on data gathered over the course of several weeks, shows that using the model's predictive capabilities improves the accuracy of robot localization and increases the efficiency of the robot navigation. The tests indicate that the proposed environment model is especially beneficial for mobile robots that do not have a complete overview of their environment, e.g. due to the limited sensor range such as when operating outdoors or in large warehouses.

While encouraging, the experiments were too short to demonstrate that the proposed method enables life-long autonomous operation in changing environments. Therefore, as part of our project goals [22], we plan the deployment of the method on a mobile robot that will operate at a large care home for a period of four months.

Moreover, we plan to extend the anomaly detection mechanism so that an anomalous map would not be rejected, but stored as an alternative map. This alternative map would represent a hypothesis that that the map change was caused by an actual environment change rather than mapping malfunction. This could result in additional robustness of the system to significant environment changes.

Acknowledgments

The work was funded by the EU ICT project 600623 'STRANDS'.

References

[1] D. Hähnel, D. Schulz, and W. Burgard, "Mobile robot mapping in populated environments," *Advanced Robotics*, 2003.

[2] D. Wolf and G. Sukhatme, "Mobile robot simultaneous localization and mapping in dynamic environments," *Autonomous Robots*, 2005.

[3] P. Biber and T. Duckett, "Experimental analysis of sample-based maps for long-term SLAM," *International Journal of Robotics Research*, 2009.

[4] W. Churchill and P. Newman, "Practice makes perfect? managing and leveraging visual experiences for lifelong navigation," in *International Conference on Robotics and Automation (ICRA)*, 2012.

[5] 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.

[6] K. Konolige and J. Bowman, "Towards lifelong visual maps," in *International Conference on Intelligent Robots and Systems*, 2009.

[7] M. Quigley, K. Conley, B. P. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Y. Ng, "ROS: an open-source robot operating system," in *ICRA Workshop on Open Source Software*, 2009.

[8] T. Krajnik, J. Pulido Fentanes, J. Santos, K. Kusumam, and T. Duckett, "Fremen: Frequency map enhancement for long-term mobile robot autonomy in changing environments," in *Workshop on Visual Place Recognition in Changing Environments (VPRiCE) ICRA2015*. IEEE, 2015.

[9] 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.

[10] J. Pulido Fentanes, B. Lacerda, T. Krajník, N. Hawes, and M. Hanheide, "Now or later? predicting and maximising success of navigation actions from long-term experience," in *International Conference on Robotics and Automation (ICRA)*, 2015.

[11] 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 *International Conference on Robotics and Automation (ICRA)*, 2015.

[12] S. Thrun, "Robotic mapping: A survey," *Exploring artificial intelligence in the new millennium*, pp. 1–35, 2002.

[13] L. Montesano, J. Minguez, and L. Montano, "Modeling dynamic scenarios for local sensor-based motion planning," *Autonomous Robots*, vol. 25, no. 3, pp. 231–251, 2008.

[14] 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.

[15] T. Krajník, S. Pedre, and L. Přeučil, "Monocular navigation system for long-term autonomy," in *International Conference on Advanced Robotics (ICAR)*, 2013.

[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 *International Conference on Intelligent Robots and Systems (IROS)*, 2008.

[18] M. Yguel, O. Aycard, and C. Laugier, "Wavelet occupancy grids: a method for compact map building," in *Field and Service Robotics*. Springer, 2006, pp. 219–230.

[19] T. Kucner *et al.*, "Conditional transition maps: Learning motion patterns in dynamic environments," in *Int. Robots and Systems*, 2013.

[20] T. Krajník, J. P. Fentanes, G. Cielniak, C. Dondrup, and T. Duckett, "Spectral analysis for long-term robotic mapping," in *International Conference on Robotics and Automation (ICRA)*, 2014.

[21] E. Marder-Eppstein, E. Berger, T. Foote, B. Gerkey, and K. Konolige, "The office marathon: Robust navigation in an indoor office environment," in *International Conference on Robotics and Automation (ICRA)*, May 2010, pp. 300–307.

[22] "Spatio-temporal representations and activities for cognitive control in long-term scenarios," [Cit: 2016-02-29]. [Online]. Available: http://www.strands-project.eu

[23] J. M. Santos, T. Krajnik, J. Pulido Fentanes, and T. Duckett, "Lifelong information-driven exploration to complete and refine 4d spatio-temporal maps," *Robotics and Automation Letters*, 2016.

[24] T. Krajník, M. Nitsche, J. Faigl, P. Vaněk, M. Saska, L. Přeučil, T. Duckett, and M. Mejail, "A practical multirobot localization system," *Journal of Intelligent & Robotic Systems*, vol. 76, no. 3-4, pp. 539–562, 2014.