

Towards life-long mapping of dynamic environments using temporal persistence modeling



Georgios Tsamis¹ Ioannis Kostavelis¹ Dimitrios Giakoumis¹ Dimitrios Tzovaras¹

¹Centre for Research and Technology Hellas / Information Technologies Institute, 6th km Harilaou - Thermi, 57001, Thessaloniki, Greece

Abstract

The contemporary SLAM mapping systems assume a static environment and build a map that is then used for mobile robot navigation disregarding the dynamic changes in this environment. The paper at hand presents a novel solution for the problem of lifelong mapping that continually updates a metric map represented as a 2D occupancy grid in large scale indoor environments with movable objects such as people, robots, objects etc. suitable for industrial applications. We formalize each cell's occupancy as a failure analysis problem and contribute temporal persistence modeling (TPM), an algorithm for probabilistic prediction of the time that a cell in an observed location is expected to be "occupied" or "empty" given sparse prior observations from a task specific mobile robot. Our work is evaluated in Gazebo simulation environment against the nominal occupancy of cells and the estimated obstacles persistence. We also show that robot navigation with life-long mapping demands less re-plans and leads to more efficient navigation in highly dynamic environments.

Contribution

One major challenge that robotic agents still face is that of long-term autonomous operation in dynamic environments, such as a factory floor. The robots have to deal with changing conditions where other robots, workers or moving objects such as pallets move around. To tackle this issue, various techniques have been used: multiple maps in various timescales can be retained [1, 2], the future state of the environment can be predicted by identified periodicities [3, 4].

In our work we employ temporal persistence modeling, otherwise used for predicting the location of previously detected objects [5], in order to predict the state of cells in the life-long map by gathering observations from the robot sensors. We implement the life-long observations as persistence probabilities which are integrated into a dedicated plugin for the ROS navigation environment. This is then used during robot global path planning. Our method enables robot navigation by avoiding congested areas reducing the required replans, leading the robot to its target location without unnecessary maneuvering.

Temporal Persistence Modeling

Temporal persistence is defined in this work as the time needed for the state of a map cell to change from "occupied" to "empty". In that sense, temporal persistence modeling is an algorithm for probabilistic modeling the occupancy of a cell based on temporally sparse observations.

Given the current time t_c and the last time t_p a cell c was observed as occupied, the aim is to calculate the probability P that c is still occupied. Exponential distributions can model the time elapsed between events, such as temporal persistence, using one rate parameter, λ and $\lambda = 1/\mu$, where μ is the expected value of the distribution. In the described model, μ_c is the average time between the last time a cell c was observed as occupied, and the time it became free again. Because each cell c cannot always be observed, μ is calculated in a probabilistic manner each time a cell c is observed as free.



Figure 1: Timeline for a cell c, observed at times t_1, t_2, t_3, t_4, t_5

Assuming a cell *c* is observed as occupied at times t_1 , t_2 , t_3 and t_4 , before it is observed as free at t_5 , as shown in Fig. 1, it is probable that *c* stopped being occupied at some time t_f between t_4 and t_5 . The change is assumed to have taken place somewhere in the middle of Δt . This leads to fitting a normal distribution over Δt_i in order to randomly pick a time, needed for calculating μ_c . The mean of the distribution should be centered on the middle of Δt . Also, the estimated time of the occupancy change should lie in Δt , with a probability of 99.7%. The distribution should therefore have a value of $N(\Delta t/2, \Delta t/3)$.

By utilizing the normal distribution, an estimated t_f can be determined that will be later employed for the computation of μ_c , when a cell *c* is observed as free. In order to have a better representation, older values of μ_c are taken into account:

$$\mu_c = \frac{\left(\sum_{i=1}^n (i/n)\mu_{c,i} + t_f\right)}{\left(\sum_{i=1}^n (i/n) + 1\right)}$$

where μ_{ci} are older, similarly produced, mean times for that cell. More recent values of μ_c have larger weights than less recent ones.

Having calculated μ_c , it is possible to compute $\lambda = 1/\mu_c$ for the exponential distribution. To find the probability for a cell *c* to have become free before or on time *t*, the cumulative distribution function is used:

$$df_{exp}(t,\lambda) = \begin{cases} 1 - e^{\lambda t}, & \text{if } t \ge 1\\ 0, & \text{otherwise} \end{cases}$$

So, the probability *P* that a cell c, last observed at time t_p , is occupied at current time t_c is equal to 1 minus the probability c has become free before or on t_c , that is given from the cumulative distribution function. *P* is then calculated as follows:

$$P(c \text{ is occupied}|t_c, t_l) = 1 - cdf_{exp}(t_c - t_l, 1/\mu_c) = e^{-1/\mu_c(t_c - t_l)}$$

Higher *P* indicates that a cell is more likely to be currently occupied. Having calculated the occupancy probability for one cell using temporal persistence modeling (TPM), the same procedure can be followed for each cell of the occupancy grid map of the environment and a new map can be created. This new map constitutes the *life-long map* in our work and can represent the expected occupancy of each cell at a given time, even when some areas are not directly observed by the sensor measurements of the mobile robot. The cells regarding the already known static obstacles can be filtered out of the expected map, thus leaving only the dynamic areas of the map that tend to change.

Navigation with life-long TPM maps

A view of an instance of the simulated factory floor is shown in Fig. 2a. In each of the distinct areas above the conveyor belt a model resembling a human is moving arbitrarily. The mobile robot is parked below the conveyor belt. In Fig. 2b the metric map of the same location the factory floor is shown; black cells on the grid represent the obstacles observed during the initial mapping process. In Fig. 2c the life-long TPM map of the exact area. Brighter colors indicate higher probability of occupancy.



Figure 2: (a) View from a location in factory floor model in Gazebo; (b) Metric map of the same location produced using conventional SLAM techniques; (c) Occupancy probability for cells in the same location obtained from TPM

To incorporate the life-long TPM map into the navigation stack of ROS, an additional costmap layer is introduced, based on the occupancy probabilities of each cell derived from TPM. In case the occupancy probability of a cell is lower than 0.5, the cost for the respective cells is not updated. If the probability is close to 1, a costmap value that indicates that the robot would certainly be in collision, or in close proximity, is assigned. In the rest of the cases, a cost indicating that the robot could be in collision in those cells is assigned. This cost assignment enables the path planning module to confidently circumvent areas with high occupancy probabilities.

Experimental Assessment

In order to assess the effectiveness of the described method, a mobile robot was deployed in the simulated factory floor, where multiple worker models and a robot model were arbitrarily moving. The initial static map of the factory floor was created by a robot using conventional SLAM techniques.

One experiment was comparing the created heatmap with the expected occupancy probabilities to a heatmap of the mean occupancies of the whole factory floor. In Fig. 3 both heatmaps are shown. A cell by cell comparison between the two heatmap grids, after compensating for the thinned lines obtained by Gazebo, showed more than 95% accuracy of the prediction of temporal persistence.

Another experiment was conducted to assess how path planning improved with the proposed method. The global path planner used utilizes the A* algorithm for creating a plan considering the costmap. In Fig. 4a the metric map of the factory floor is shown with the global costmap layered on top of it, without the TPM layer. In Fig. 4b the costmap contains the costs associated with the high occupancy probabilities of persistent objects, calculated by the TPM plugin, resulting in planning around those areas.

The plans produced without the TPM plugin pass through congested areas where the predicted occupancy probability is high. This results in 3 times more replanning actions by the path planner, compared to the paths made when the TPM plugin was used. In these cases, the robot also passes in close proximity of the moving models, which can lead to collisions. On the other hand, paths planned taking the TPM costmap layer into consideration are longer and take longer to be completed. However, this can be mitigated by raising the speed of the robot through low occupancy probability areas.



Figure 3: (a) Heatmap of the mean occupancy of the map cells, derived by the known trajectories of the Gazebo models; (b) Heatmap of the predicted occupancy of each map cell, after employing TPM



Figure 4: Global costmap displayed over part of the metric map of the factory floor (a) without TPM plugin loaded, and (b) with TPM plugin loaded. A planned path to the same goal is shown in blue

Acknowledgement

This work has been supported by the EU Horizon 2020 funded project "Hybrid Human-Robot RECycling plant for electriCal and eLEctRonic Equipment (HR-RECYCLER)" under the grant agreement with no: 820742.

References

[1] Peter Biber, Tom Duckett, et al. Dynamic maps for long-term operation of mobile service robots, Robotics: science and systems, pages 17--24, 2005.

[2] Kurt Konolige and James Bowman. Towards lifelong visual maps, 2009 IEEE/RSJ International Conference on Intelligent Robots and Systems, pages 1156--1163.

[3] Tomas Krajnik, Jaime Pulido Fentanes, Grzegorz Cielniak, Christian Dondrup, and Tom Duckett. Spectral analysis for long-term robotic mapping, 2014 IEEE International Conference on Robotics and Automation (ICRA), pages 3706--3711.

[4] Joao Machado Santos, Tomáš Krajník, Jaime Pulido Fentanes, and Tom Duckett. Lifelong information-driven exploration to complete and refine 4-d spatio-temporal maps, IEEE Robotics and Automation Letters, 1(2):684--691, 2016.

[5] Russell Toris and Sonia Chernova. Temporal persistence modeling for object search, 2017 IEEE international conference on robotics and automation (ICRA), pages 3215--3222.