



Queensland University of Technology
Brisbane Australia

This is the author's version of a work that was submitted/accepted for publication in the following source:

[Morris, Timothy](#), Dayoub, Feras, [Corke, Peter](#), & [Upcroft, Ben](#)
(2014)

Simultaneous localization and planning on multiple map hypotheses. In *Proceedings of the 2014 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2014)*, IEEE, Chicago, The United States of America, pp. 4531-4536.

This file was downloaded from: <https://eprints.qut.edu.au/75193/>

© ©2014 IEEE

Notice: *Changes introduced as a result of publishing processes such as copy-editing and formatting may not be reflected in this document. For a definitive version of this work, please refer to the published source:*

<https://doi.org/10.1109/IROS.2014.6943204>

Simultaneous Localization and Planning on Multiple Map Hypotheses

Timothy Morris, Feras Dayoub, Peter Corke and Ben Uprocft¹

Abstract—This paper presents a novel method to rank map hypotheses by the quality of localization they afford. The highest ranked hypothesis at any moment becomes the active representation that is used to guide the robot to its goal location. A single static representation is insufficient for navigation in dynamic environments where paths can be blocked periodically, a common scenario which poses significant challenges for typical planners. In our approach we simultaneously rank multiple map hypotheses by the influence that localization in each of them has on locally accurate odometry. This is done online for the current locally accurate window by formulating a factor graph of odometry relaxed by localization constraints. Comparison of the resulting perturbed odometry of each hypothesis with the original odometry yields a score that can be used to rank map hypotheses by their utility. We deploy the proposed approach on a real robot navigating a structurally noisy office environment. The configuration of the environment is physically altered outside the robots sensory horizon during navigation tasks to demonstrate the proposed approach of hypothesis selection.

I. INTRODUCTION

When operating in non-stationary environments the assumption of a static world representation is not valid. Structural changes affect the geometry perceived by the robot which in turn influences navigation and planing performance. In such an environment it is not possible to have an up to date map due the nature of the real world. Some examples of structural changes are in areas of construction or delivery. To deal with these dynamics we consider multiple representations extracted from an efficient memory architecture [1] to account for this change, see Figure 1. However in one instance of time a robot needs to select a single representation to act in the environment, this can be done by selecting the most likely representation based on the current experience of the world. Our research presented herein will focus on the selection of actionable representations during navigation given that multiple representations are already available.

A multi hypothesis representation supersedes the classical approach of a single hypothesis representation in the following areas:

- **Localization:** The ability to provide localization and alternate routes should large portions of the represented world change before / during navigation.
- **Planning:** Over time representations will be ranked by their successful use, enabling the robot to learn from mistakes and ultimately provide for temporal planning.

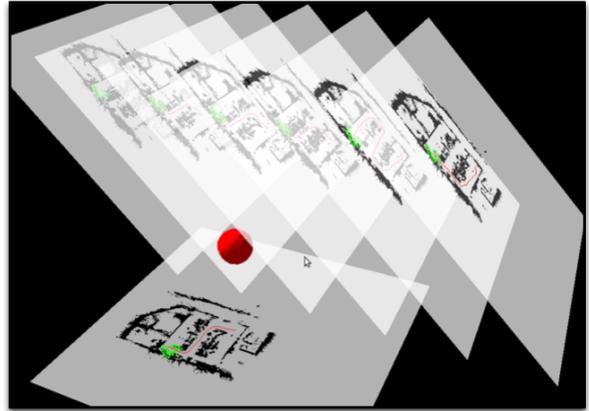


Fig. 1. This visualization taken in *rviz* demonstrates the hypotheses being simultaneously tracked. The highest ranked hypothesis at present marked by a red sphere is actioned for guiding the robot during navigation. Should an alternative map configuration provide greater localization performance and contain a valid plan it can replace the current map and plan as the new primary configuration.

- **Mapping:** Identification of which configurations represented have been the most useful and which might require updating.

The contribution of this paper is a novel approach to continually ranking alternate hypotheses during navigation. The stability in this ranking approach is a result of incorporating information from localization and odometry. A robust factor graph is used to consider only localization which conforms to the locally accurate window of integrated odometry. The sensitivity of excluding a localization estimate is adjusted by the distance to the expected location along a plan made on the map hypothesis itself. This measures how well a hypothesis predicted the route taken by the robot and excludes information provided by localization when using unexpected regions of the map to localize.

The rest of the paper is structured as follows. After an overview of related work Section II, we will introduce our proposed approach III. Subsections in the approach will describe key components to the approach and how they relate to the benefits of a multi hypothesis representation. Section IV details experiments and results supporting the approach. Finally we conclude in Section V and discuss future research.

II. RELATED WORK

Localization in dynamic environments is an established area of research in robotics. Estimating the robot pose in

¹ The authors are with the CyPhy Lab, Queensland University of Technology, Brisbane, Australia `firstname.lastname@qut.edu.au`

some known space by way of matching observation to an internal representation. We adopt the classification of dynamics as provided in [2] of static, semi-static and dynamic.

- **Static:** Unchanging location and classification of observation.
- **Semi-static:** Change occurring outside the sensory horizon.
- **Dynamic:** Change occurring within the sensory horizon.

These definitions are provided to aid in the discussion of problems relating to long-term mapping and navigation, not to restrict the membership of objects.

Research towards dealing with dynamics in this area largely fall into two broad categories of modelling or detection. For detection, some attempt to identify and remove dynamics or semi-static information from sensory input [3]. While others suggest accumulating dynamics in case they later prove to be useful for localization [4]. Both assume the most relevant landmarks of the world are known. The alternative approach to dealing with dynamics involves factoring in some expected dynamics by maintaining a margin of uncertainty of their pose estimates [5]. Research in this field demonstrates a factor which is intuitively true when navigating in the real world under changing conditions: At some point a decision must be actioned and mistakes will be made.

Structural change occurs most commonly from objects that occupy various states in the world. If the assumption is made that static can be defined as a period of time left unchanged, then observations can be temporally filtered into a single static representation [2], [6], [7]. Other methods include using multiple timescales [8] and choosing the single representation that most closely matches the current observation.

A single representation that is updated based on recurring observation can provide good localization performance [2] over relatively short periods. This is because semi-static and dynamic observations are considered noise, this noise can overcome the available landmarks used for localization and the robot can become lost. An alternative is to maintain a representation that contains both static and semi-static observations [9]. This can be achieved by tracking semi-static objects and their states [10], [3] or by compiling similar episodes of observation and determining which are useful Short Term Features [11].

Many of these approaches deal well with filtering dynamics from the stored representation, unfortunately details of semi-static configurations are also lost. To capture and use semi-static change it was proposed [12] to attach semi-static configurations to a static representation and recall them when localization in the static map is poor. This approach requires an initial static map that does not evolve over time as semi-static configurations are referenced to it. In addition, localization is tracked only in a single representation at any one time, thereby limiting recovery as the only hypotheses available are those near the potentially corrupted localized position.

A solution to poor / aliased localization and repeated planning mistakes is to maintain and use a set of local map hypotheses. Simultaneously localizing and planning within a set of hypotheses enables recovery to previously observed world configurations during navigation tasks, should the utility of the actioned hypotheses fall. Switching during a navigation task based on a metric of map utility was introduced in [1]. This is a mode of recovery that is not available to approaches that track only a single hypotheses of maximum likelihood.

Map hypotheses can be built using traditional SLAM techniques such as [13], constructed at key intervals to capture various states of the working world. Although such an approach would require a trigger to determine key intervals and full coverage each time to produce a complete map. An alternative is to use an incremental mapping approach where observations accumulated during short experiences determined by the navigation task. In our case, these observations pass through a two stage filter (Short Term Memory (STM)- Long Term Memory (LTM) [14]) as in [1]. Structure that is observed sufficiently will migrate from the STM to the LTM, this structure is provided the label of a configurable 'piece' of the environment. In this fashion map hypotheses are built from static and semi-static configurations which are free to evolve and change membership or be forgotten over time.

In previous work [1], we considered uncertainty provided by a Markovian localizer as an indication of map performance. While this metric is a broad indication of the accuracy in localization, we found a ranking system solely reliant on the Markovian process can be fooled by aliased localization. This aliasing effect occurs in regions where the map is similar and as a result localization uncertainty is underestimated. In addition, localization performance is only half the reason for maintaining a map of the environment. We therefore need to consider how well a particular map hypotheses is able to 'predict' the environment outside the robots Field of View in the way of planning a path though it.

III. APPROACH

In this work we propose a novel approach to computing a metric for comparing map hypotheses during the navigation task. A useful map supports accurate localization which can be used to constrain odometry drift over longer periods. The key to this work is to instead use the fact that odometry can be locally accurate over short distances. With this knowledge, localization that complements or refines local odometry estimates indicate that useful information was provided by the map. We can then take the influence of localization from each map hypothesis as a metric for selecting the most relevant hypothesis at any moment during the navigation task. Finally, to consider the combined performance of many locally accurate trajectories considered during a navigation task we should also consider how closely the trajectories tracked the planned route.

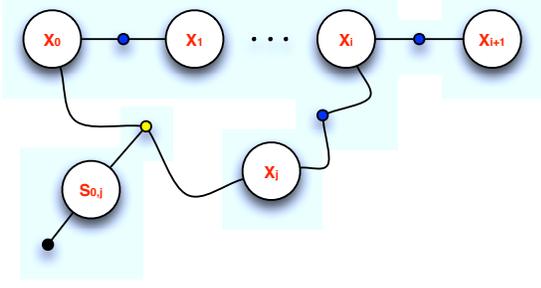


Fig. 2. Factor Graph representation of a locally accurate trajectory when using localization estimates as loop closures on odometry in SE2. Position estimates of X_j are considered to be localization estimates sampled at the same place and time as odometry X_i , i.e. $u_{ij} = \{0, 0, 0\}$. The variable S_{0j} governs the sensitivity of inclusion when considering differing position estimates X_i (integrated odometry) and X_j (fixed frame localization).

We propose the ranking of hypotheses can be made during navigation by measuring information gain / loss to locally accurate odometry given localization estimates. To find its influence we relax the trajectory provided by locally accurate odometry X_i by localization constraints. This is intuitive as localization estimates can be paired with odometry measurements taken at the same physical place and time. The relaxed odometry trajectory of each hypothesis, referred to as the perturbed odometry X_i^* , represents the influence of localization on odometry. The information provided by localization can be determined by comparison of X_i and X_i^* .

In the remainder of this section we will briefly review the source of map hypotheses. We then discuss how odometry can be perturbed by localization in two parts. Firstly how we incorporate localization from hypotheses into *Perturbed Odometry* trajectories and how to bias what localization should contribute using *Plan Deviation*. Finally we will show how the original integrated odometry trajectory can be compared to the new perturbed odometry trajectories to compute an information gain / loss metric for each hypothesis used for online ranking during navigation.

A. Gathering Hypotheses

The source of map hypotheses is inconsequential here as this work focuses simply on how to measure utility of any set of map hypotheses for navigation tasks online. In principle this work could be used to compare multiple mapping approaches and their utility for actual navigation. For this work we use a memory architecture as in [1] to provide configurable regions of the environment. Configurations are maintained by the memory architecture and the hypotheses used are subject to availability and may change at any time between requests. However, the choice to localize against the structure of a particular configuration will influence when the memory will next observe the region. Therefore regions of memory frequently used because they facilitate accurate planning and navigation will be rehearsed to a greater degree than those that do not.

B. Perturbed Odometry

In this work we address the stability of selecting an active representation by considering how the locally accurate odometry trajectory can be perturbed by localization. Here we effectively say over a short trajectory, good localization should complement odometry and its integrated uncertainty.

To consider localization estimates as part of a continuous trajectory we match each localization estimate with its temporally co-occurring integrated odometry pose. This effectively provides a series of position estimates which correspond to the single trajectory actually traversed.

We find the perturbation of odometry given localization by formulating the problem as a factor graph in SE2 (x, y, yaw), *Figure 2*. Here an integrated odometry estimate x_i is made at the same *Time* as the localization estimate x_j . We therefore know with complete certainty that the constraint u_{ij} that connects these factors is a zero transform. The constraint u_{0j} is between factors representing the latest locally accurate origin and localization estimate. The additional variable S_{0j} governs the sensitivity of the localizations inclusion in the final information comparison between odometry and odometry with localization.

Let x_i be the poses provided by integration of odometry constraints u_i . Let x_j be the poses provided by localization in a fixed frame relative to x_0 given constraint u_{0j} . Finally we can say that localization estimates can be paired with odometry estimates sampled at the same place and time by a perfect zero constraint $u_{ij} = \{0, 0, 0\}$. Therefore similar to [15] we can find the conditional probability over all variables $X = \{x_i \cup x_j\}$ and constraints $U = \{u_i \cup u_{ij} \cup u_{0j}\}$

$$P(X | U) \propto \prod_i P(x_{i+1} | x_i, u_i) \quad (1)$$

$$\cdot \prod_{ij} P(x_j | x_i, u_{ij})$$

$$\cdot \prod_j P(x_j | x_0, u_{0j})$$

The maximum a posteriori configuration of poses, X^* , can be found by maximising the joint probability:

$$X^* = \underset{X}{\operatorname{argmax}} P(X | U) = \underset{X}{\operatorname{argmin}} -\log P(X | U) \quad (2)$$

$$= \underset{X}{\operatorname{argmin}} \sum_i \|f(x_i, u_i) - x_{i+1}\|_{\Sigma_i}^2$$

$$+ \sum_{ij} \|f(x_i, u_{ij}) - x_j\|_{\Phi_{ij}}^2$$

$$+ \sum_j \|f(x_0, u_{0j}) - x_j\|_{\Lambda_{0j}}^2$$

which is a nonlinear least squares problem with $\|a - b\|_{\Sigma}^2$ as a squared Mahalanobis distance with covariance Σ .

Figure 3 demonstrates how hypotheses are evaluated at each time step. Localization estimates with position uncertainty are used to perturb an integrated odometry trajectory and its accumulated uncertainty. The original integrated

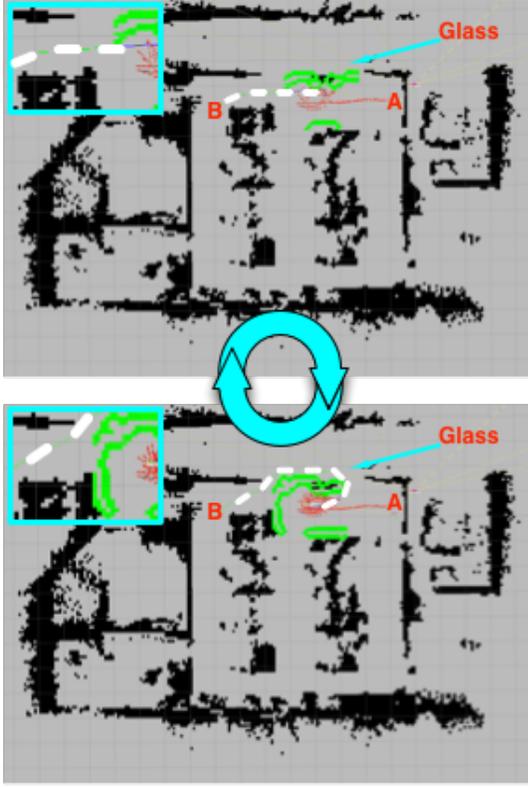


Fig. 4. A to B where measuring the uncertainty of Markovian localization itself is insufficient to promote an alternative hypothesis. An example where there is no sensible solution to the navigation task, due to partially visible structure (Glass). In this case a glass wall extends horizontally above ‘A’ and ‘B’ but only portions have made the projection to 2D. The top left zoomed images show the plan attempting to pass through the semi-static obstacle, then enter through the top glass wall when blocked, then back to the semi-static obstacle when it is forgotten. Black represents occupied space of the hypothesis. Green represents the observed obstacle boundaries with a safety margin (inflated obstacles).

odometry is then compared to the resulting perturbed odometry for each hypothesis. Providing each hypothesis with a measure of information gain / loss for odometry given localization.

We use similar algorithms for the maintenance of ranked hypotheses and the consolidation of weights accumulated during a navigation experience as in prior research [1] and details will not be replicated here.

C. Plan Deviation

The utility of a hypothesis is also in its ability to predict the path to be taken. Preliminary experimentation demonstrated that good localization can be ‘faked’ by hypotheses in the real world due to aliasing caused by noisy sensors and maps. This can lead to a similar situation as the single hypothesis plan oscillation (Figure 4) if a ‘bad’ hypothesis benefits from aliased localization confidence. The scenario occurs when hypotheses share common areas or are initially localized in different locations structurally similar. To overcome this we needed to increase the sensitivity of excluding localization estimates for individual hypotheses

at times when the robot deviates from their expected plan. To achieve this we introduce a switching prior as in [15], although instead of switching ‘loop-closures’ from multiple passes we are switching ‘localization estimates’ relative to a fixed locally accurate odometry frame.

$$\begin{aligned}
 X^*, S^* = \underset{X, S}{\operatorname{argmin}} & \sum_i \|f(x_i, u_i) - x_{i+1}\|_{\Sigma_i}^2 & (3) \\
 & + \sum_{ij} \|f(x_i, u_{ij}) - x_j\|_{\Phi_{ij}}^2 \\
 & + \sum_j \|sig(s_{0j}) \cdot f(x_i, u_{0j}) - x_j\|_{\Lambda_{0j}}^2 \\
 & + \sum_j \|\gamma_{0j} - s_{0j}\|_{\Xi_{0j}}^2
 \end{aligned}$$

The switching prior γ_{0j} can vary by the scaled plan deviation of the constraint’s localised pose. This alters the sensitivity to localization error by the sigmoid scaling function $sig(s_{0j})$, in turn varying the error at which the constraint will be ignored ($sig(s_{0j}) = 0$) in the relaxation process. Using a variance of 1 the scaled value used for the prior ranges from 0 to 2, being fixed ‘off’ to fixed ‘on’ at the extremes. We calculate the prior as follows:

$$\gamma_{0j} = \left(1 - \frac{\text{DistanceToPlan}_{0j}}{\text{MaxPlanDeviation}}\right) \cdot 2.0 \quad (4)$$

D. Perturbed Odometry Comparison

If we now consider two pose sets, the original integrated odometry X_i and the new perturbed by localization X_i^* . We can find the divergence of X_i^* from X_i by the sum of Kullback-Leibler divergence over corresponding poses:

$$\begin{aligned}
 D_{KL}(X \| X^*) &= \frac{1}{2} (\operatorname{tr}(\Sigma^{*-1} \Sigma) & (5) \\
 & + (\mu^* - \mu)^T \Sigma^{*-1} (\mu^* - \mu) \\
 & - n - \ln(\frac{\det \Sigma}{\det \Sigma^*}))
 \end{aligned}$$

Where $n = 3$ degrees of freedom and means μ, μ^* and covariance matrices Σ, Σ^* .

Now let the number of hypotheses be k , most recent pose of hypothesis X^{h^*} be p and the iterative reverse search limit m .

$$f(h_i, m) \Big|_{i=1}^k = \sum_{j=0}^m D_{KL}(X_{p-j}^h \| X_{p-j}^{h^*}) \quad (6)$$

We now select the hypothesis h_i who’s perturbed poses provide the least divergence over the largest sequence of poses beginning at the most recent.

$$\underset{m}{\operatorname{argmin}} f(h_i, m) \Big|_{i=1}^k \quad (7)$$

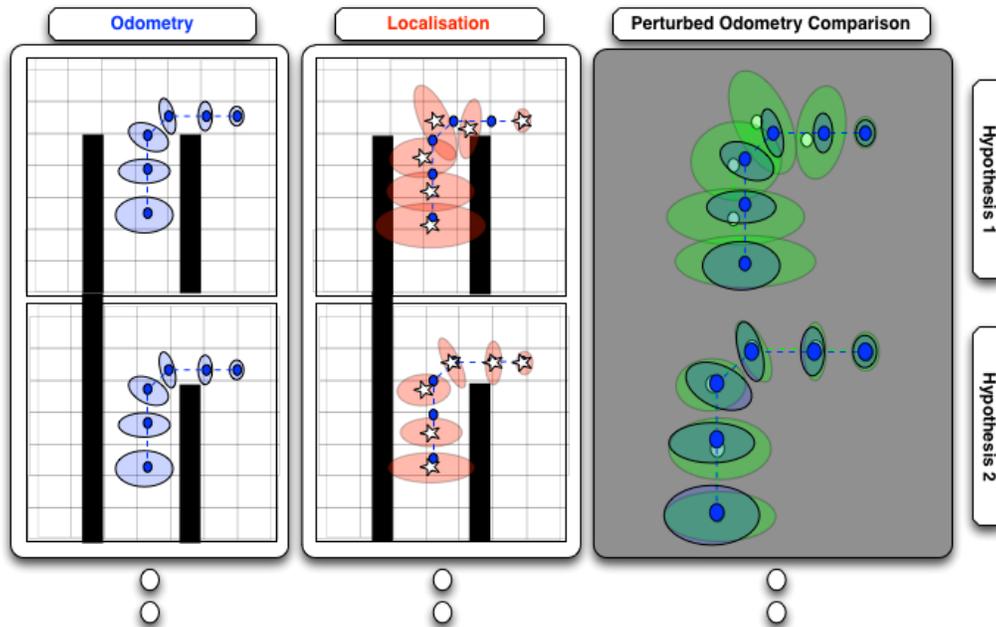


Fig. 3. This illustration demonstrates the hypothesis ranking process. An integrated trajectory of odometry (*Blue*) common to all hypotheses is relaxed with pose estimates from localization (*Red*). The perturbed odometry (*Green*) obtained for each hypothesis is then compared to the initial odometry trajectory. The comparison is made using the Kullback-Leibler divergence of odometry given localization, providing an information gain / loss at each pose pair based on their mean and covariance. The net information change over the hypothesis is used to rank alternative hypotheses.

IV. EXPERIMENTAL EVALUATION

Here we demonstrate our approach through experimental evaluation on the MobileRobots’ Research GuiaBot platform. The robot is equipped with a Kinect RGB-D sensor that provides 3D point clouds to the memory. Configurations taken from memory are projected onto 2D Occupancy grids as in our case we are navigating with 2 degrees of freedom.

The office environment where the experiment took place is of size $16m \times 20m$. In the following experiments we have frozen the state of our memory architecture to provide 16 map hypotheses that contain both major and mini structural change. This is done as we are investigating the ranking of hypotheses rather than their source. It should be noted here that noisy point clouds from the RGB-D sensor are the input to both the memory and the localizer. Therefore the noisy visual appearance of the map hypotheses is to be expected, especially in cluttered environments with walls of glass.

The maximum plan deviation was set to 2 Meters for the experiment shown in Figure 5.

A brief example of navigation on a single hypothesis is provided in Figure 4. In this example the single hypothesis does not quite match the physical world. In this case the robot oscillates between two plans indefinitely as both are short and can not be observed blocked within a single sensory horizon window. It is common practice to select the shortest path when comparing plans, the only escape from this oscillation is to use an arbitrary set of heuristics. Unfortunately such a heuristic might reduce the robots ability to deal with short term blockages (humans) and plan deviations. Instead our approach suggests an alternative hypothesis (learnt by

memory) that restricts the planer into taking the alternate longer route to achieve the goal, see Figure 5.

We repeated the experiment, navigating from ‘A’ to ‘B’ with the direct path blocked. We then removed the blockage while the robot was in transit along the longer route and out of sensor range. Once the initial goal of ‘B’ was reached the robot was tasked with returning to ‘A’. To demonstrate the benefit of including the plan deviation as a switching prior we also tried increasing the max plan deviation to a level effectively forcing all localization estimates to be included in the evaluation process. Details in Table I show that using plan deviation results in the minimum number of switches and the most direct route. When plan deviation is effectively ignored the navigation is much less efficient as localization performance was similar in the centre of the map. This would cause the robot to drop its current plan being traversed and re-plan to the shorter plan between ‘A’ and ‘B’.

MaxPlanDeviation (M)	Switches	Time (S)	Distance (M)
2	2	127	27.47
100	6	370	50.90

TABLE I
NAVIGATION PERFORMANCE WHEN USING PLAN DEVIATION.

In such an approach there is a fair assumption that odometry uncertainty is expected to account for potential errors in odometry measurements. Failing this assumption does not entirely invalidate the approach as any errors falling outside the expected confidence interval will impact



Fig. 5. A to B using the proposed approach. Each panel representing the robots navigation state at stages during the same run. Black represents occupied space of the hypothesis. Green represents the observed obstacle boundaries with a safety margin (inflated obstacles). Enchantments added for the reader include: Letters 'A' and 'B' for starting and goal locations. Blue dash's for plan completed and White plan under current action.

all hypotheses with the same corruption. The effect may be felt if there is significant deviation from the intended plan. But this effect is determined by the maximum plan deviation aloud, 2 meters in our case, which would require a significant odometry failure. Perhaps this is an avenue for future research. Although it is likely that the best solution would be to trigger a new locally accurate segment when such a corruption could be detected. Possible detection could be provided by comparing the delta pose estimate with the commands sent to the robot.

V. CONCLUSIONS AND FUTURE WORK

This paper introduced a novel method of measuring the navigation performance of alternate map hypotheses. This approach incorporates both localization and planing performance and is generally applicable to ranking the utility of maps given a source of odometry and localization. We have demonstrated the approach implemented on a real robot for active navigation given alternative hypotheses accumulated through our memory filter introduced in prior work. It is our belief that for life-long learning and navigation we need to hypothesise based on past observation and that is equally important to determine the utility of these hypotheses for the

navigation task itself.

In future we would like to use the ranking approach described in this paper to measure the utility of maps provided by various mapping approaches. This comparison will provide greater insight into the requirements of internal representations for the task of real world navigation. Future experimentation will look at a long term deployment in a bookshop and food court area where structural change such as deliveries, lunch crowds, book displays is common place. We expect to learn from this deployment temporal patterns of structural change that impact navigation performance and how to use varying performance to influence the rate of decay in memory.

ACKNOWLEDGMENT

This research was supported under Australian Research Council's Discovery Projects funding scheme (project number DP110103006).

REFERENCES

- [1] T. Morris, F. Dayoub, P. Corke, G. Wyeth, and B. Upcroft, "Multiple map hypotheses for planning and navigating in non-stationary environments," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2014.
- [2] M. Hentschel and B. Wagner, "An adaptive memory model for long-term navigation of autonomous mobile robots," *Journal of Robotics*, 2012.
- [3] D. Anguelov, R. Biswas, D. Koller, B. Limketkai, and S. Thrun, "Learning hierarchical object maps of non-stationary environments with mobile robots," in *International Conference on Uncertainty in artificial intelligence*, 2002.
- [4] D. F. Wolf and G. S. Sukhatme, "Mobile robot simultaneous localization and mapping in dynamic environments," *Journal of Autonomous Robots*, vol. 19, no. 1, pp. 53–65, 2005.
- [5] D. Fox, "KLD-sampling: Adaptive particle filters," in *In Advances in Neural Information Processing Systems 14*. MIT Press, 2001.
- [6] G. D. Tipaldi, D. Meyer-Delius, M. Beinhofer, and W. Burgard, "Lifelong localization and dynamic map estimation in changing environments," in *RSS Workshop on Robots in Clutter*, 2012.
- [7] J. Saarinen, H. Andreasson, and A. J. Lilienthal, "Independent markov chain occupancy grid maps for representation of dynamic environment," in *IEEE International Conference on Intelligent Robots and Systems (IROS)*, 2012.
- [8] P. Biber, T. Duckett *et al.*, "Dynamic maps for long-term operation of mobile service robots," in *Robotics Science and Systems (RSS)*, 2005.
- [9] C. Stachniss and W. Burgard, "Mobile robot mapping and localization in non-static environments," in *National Conference on Artificial Intelligence*. AAAI Press; MIT Press, 2005.
- [10] C.-C. Wang and C. Thorpe, "Simultaneous localization and mapping with detection and tracking of moving objects," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2002.
- [11] J. Biswas and M. Veloso, "Episodic non-markov localization: Reasoning about short-term and long-term features," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2014.
- [12] D. Meyer-Delius, J. rgen Hess, G. Grisetti, and W. Burgard, "Temporary maps for robust localization in semi-static environments," in *IEEE International Conference on Intelligent Robots and Systems (IROS)*, 2010.
- [13] G. Grisetti, C. Stachniss, and W. Burgard, "Improving grid-based slam with rao-blackwellized particle filters by adaptive proposals and selective resampling," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2005.
- [14] F. Dayoub, G. Cielniak, and T. Duckett, "Long-term experiments with an adaptive spherical view representation for navigation in changing environments," *Journal of Robotics and Autonomous Systems*, vol. 59, no. 5, pp. 285–295, 2011.
- [15] N. Sunderhauf and P. Protzel, "Towards a robust back-end for pose graph slam," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2012.