

The Quest For “Always-On” Autonomous Mobile Robots

Joydeep Biswas

CICS, University of Massachusetts Amherst, MA, USA

joydeepb@cs.umass.edu

Abstract

Building “always-on” robots to be deployed over extended periods of time in real human environments is challenging for several reasons. Some fundamental questions that arise in the process include: 1) How can the robot reconcile unexpected differences between its observations and its map of the world? 2) How can we scalably test robots for long-term autonomy? 3) Can a robot learn to *predict* its own failures, and their corresponding *causes*? 4) When the robot fails and is unable to recover autonomously, can it utilize partially specified human corrections to overcome its failures? This paper summarizes our research towards addressing all of these questions. We present 1) Episodic non-Markov Localization to maintain the belief of the robot’s location while explicitly reasoning about unmapped observations; 2) a 1,000km Challenge to test for long-term autonomy; 3) feature-based and learning-based approaches to predicting failures; and 4) human-in-the-loop SLAM to overcome robot mapping errors, and SMT-based robot transition repair to overcome state machine failures.

1 Introduction

We seek the ultimate goal of having autonomous service mobile robots working in human environments, performing tasks accurately and robustly on demand. Such robots should be capable of working for not minutes, hours, or days, but for *years* at a stretch. Developing, testing, and overcoming failures for algorithms that can execute robustly and accurately over such protracted time scales involves unique challenges. First, perception algorithms must be robust to changes in the environment. Second, over such long deployments, robots will inevitably make errors. They will have to anticipate errors, and overcome their failures without requiring the assistance of robotics experts. Finally, testing for robustness over such a long duration is a challenge in of itself: we need new approaches to process logs to gain insights into *why* algorithms failed, to support the development of techniques to overcome such limitations. In this paper, we summarize our research towards addressing all of these challenges.

In Section 2, we present Episodic non-Markov Localization (EnML), a localization algorithm for changing environments that reasons about the *nature* of its observations: whether they correspond to permanent objects, movable objects termed, or moving objects. EnML reasons about unmapped objects in real time, at the time of deployments, without relying on up-to-date static maps of the environment.

To demonstrate the robustness and accuracy of localization using these contributions, a few years ago, we proposed the 1,000km Challenge: Demonstrate, on a team of deployed autonomous mobile robots, in multiple real-world human environments, the robustness and accuracy of localization over long-term deployments covering a total distance of more than 1,000km. In Section 3 we summarize our results from the 1,000km Challenge, including accuracy, and robustness in different environments over the 1,000km Challenge.

Over long-term deployments, autonomous robots will inevitably make mistakes. In Section 4, we present our work on predicting such errors. From data collected over the 1,000km Challenge, we demonstrate that the dynamics of the environment correlates with the robustness of localization. We further present our work on *introspective vision* to use supervisory sensing to learn models of failure of stereo vision used for obstacle avoidance. The learned models are capable of not only accurately predicting failures, but also inferring *the number* of distinct classes of failures.

Next, in Section 5 we present our research on overcoming failures of robots by incorporating approximate, partially specified corrections from non-technical users. Specifically, we present our work on human-in-the-loop SLAM [Nashed and Biswas, b] to overcome mapping failures, and SMT-based robot transition repair [Holtz *et al.*,] to overcome action failures arising from incorrect transitions in robot state machines.

We conclude with a discussion of open challenges and promising directions for future work in Section 6.

2 Episodic non-Markov Localization

In support of the goal of long-term autonomy, there have been a number of approaches proposed to address the inevitable changes in human environments that robots must adapt to. Some such approaches include reasoning about the different discrete states of the environment [Meyer-Delius *et al.*, 2012; Stachniss and Burgard,], expanding occupancy grid maps




Figure 1: An example instance of a Varying Graphical Network (VGN) for non-Markov localization. The non-varying nodes and edges are denoted with solid lines, and the varying nodes and edges with dashed lines. The exact structure of the graph will depend on the STFs and DFs present.

with temporal Markov chains [Saarinen *et al.*,], and time-based map-switching [Biber and Duckett, 2005; Krajník *et al.*, 2017]. In our work, we realize that maintaining multiple discrete states of the world, or dynamic models of all the parts of the environment may be infeasible for robots deployed over long periods of time, in large environments. Hence, we introduced Episodic non-Markov Localization (EnML) [Biswas *et al.*, ; Biswas and Veloso, a], an algorithm that 1) explicitly reasons in real-time about its observations that do not match the long-term map, and 2) jointly accounts for the relative pose corrections from such unmapped objects concurrently with global pose corrections from mapped objects, in 3) a single sound probabilistic framework that enables efficient inference of the maximum likelihood estimate of the robot’s trajectory.

At each time step, EnML classifies observations as arising from Long-Term Features (LTFs), Short-Term Features (STFs), or Dynamic Features (DFs). To represent the varying nature of localization in the presence of unmapped observations, we introduce a new graphical model, the *Varying Graphical Network* [Biswas and Veloso, b]. As in the Dynamic Bayesian Network for Markov Localization, a Varying Graphical Network (VGN) includes certain periodically repeating nodes and edges that do not change with the belief. A VGN includes two additional structural elements: varying nodes and varying edges to track the presence and correlations between STFs and DFs. Figure 1 shows an example instance of a VGN. By reasoning about observations from STFs and DFs over extended deployments, we have also shown how we can build model-instance object maps [Biswas and Veloso, 2014] and long-term vector maps [Nashed and Biswas, a].

Given the initial pose of the robot x_0 , observations $s_{1:n}$, odometry $u_{1:n}$, and a static map M , EnML factorizes the belief over the robot’s trajectory $x_{1:n}$, according to the cor-

relations in the VGN. We do not track DFs for localization, hence they assume a constant value based on priors. Exploiting the fact that different classes of observations (m STFs and n LTFs) are independent of each other, the belief for EnML is given by

$$Bel(x_{1:n}) \propto \prod_{i=1}^{i=m} P(s_{1:n}^{STF_i} | x_{1:n}) \prod_{j=1}^{j=n} [P(s_j^{LTF} | x_j, M) P(x_j | x_{j-1}, u_j)].$$

Since the VGN for non-Markov localization has no pre-defined structure, it might seem that computation of the belief would require storing the complete history of all states and observations of the robot. However, in practice this is not necessary. Suppose there exists a time step t_i such that all observations and state estimates made after t_i , given x_i , are independent of all prior observations and state estimates. This conditional independence implies that there are no STF observations after t_i that correspond to STF observations before t_i . In such a case, the history of states and observations prior to t_i , called the “episode” $t_{0:i-1}$, can be discarded when estimating $Bel(x_{i:n})$ over the episode $t_{i:n}$. We thus define an episode $t_{j:k}$ to be a consecutive sequence of time-steps from t_j to t_k such that the observations made between t_j and t_k are independent of all observations made before t_j , given the pose x_j , and the permanent map M .

In order to explicitly account for the effect of the past robot pose estimates $x_{0:n}$ on future observations and pose estimates, we solve for the Belief over the complete history of robot poses, $Bel(x_{0:n})$. Due to the prohibitively large state space, we concentrate on evaluating the Belief in the neighborhood of its maximum likelihood estimate, computed via non-linear least squares optimization of its log-likelihood.

3 Evaluating Long-Term Autonomy in the Real World

One of the challenges of working on long-term autonomy is that real-world evaluation is hard – there is no good surrogate for the rich diversity of scenarios that the real world offers, in order to stress-test algorithms for long-term autonomy. We therefore decided to evaluate our robot localization algorithms over the 1,000km Challenge, which spanned several years of deployments in real human environments. The localization algorithms tested included depth-based localization using plane filtering [Biswas and Veloso, a], LIDAR-based localization using Corrective Gradient Refinement [Biswas *et al.*,] for particle filters, and Episodic non-Markov Localization [Biswas and Veloso, b; Biswas and Veloso, 2017]. The 1,000km Challenge was conducted on a group of autonomous service mobile robots, the CoBots, from May 17, 2011 to November 18, 2014. Over the course of their deployments, four CoBots autonomously performed various tasks [Veloso *et al.*,] for users, including escorting visitors, transporting objects, and engaging in semi-autonomous telepresence.

The robots were deployed in several buildings, including the Gates Hillman Center (GHC) and Newell-Simon Hall

| Map | Samples | Mean error (m) | Median error (m) | Std. Dev. (m) |
|------|---------|----------------|------------------|---------------|
| GHC4 | 7656 | 0.352 | 0.343 | 0.218 |
| GHC5 | 5123 | 0.394 | 0.400 | 0.181 |
| GHC6 | 1235 | 0.296 | 0.314 | 0.122 |
| GHC7 | 17489 | 0.368 | 0.341 | 0.190 |
| GHC8 | 50 | 0.234 | 0.167 | 0.185 |
| GHC9 | 276 | 0.464 | 0.402 | 0.232 |

Table 1: Localization accuracy by sparse ground truth for the different maps over the course of the 1,000km Challenge.

(NSH) at Carnegie Mellon University, and the Center for Urban Science and Progress, 1 Metrotech Center at New York University (NYU). There were 12 floors in total across all the buildings that the CoBots were deployed on.

The 1,000km Challenge resulted in the collection of over 168GB of compressed data logs, which are available online¹. We quantitatively evaluated the accuracy and robustness in localization over the 1,000km Challenge by 1) comparison to scan matching, 2) comparison to sparse ground truth, and 3) tracking operator interventions. The complete results include analysis of the accuracy and robustness as a function of the environment [Biswas and Veloso, 2016], and here we present a short excerpt of accuracy by comparison to ground truth.

CoBot2 and CoBot3 were equipped with Hagisonic StarGazer sensors, which detect StarGazer marker patterns mounted on the ceiling. We had 46 StarGazer markers placed on the ceiling throughout the Gates-Hillman Center to provide sparse ground-truth location estimates. When a CoBot detected a StarGazer marker, the localization estimates were compared to the global location of the observed StarGazer marker to compute the error in localization.

In addition to localization accuracy, we also logged when operator intervention was required to reset robot localization. By calculating the mean distance traversed between such interventions, we were able to quantify the robustness of localization as a function of location and algorithm used. In the next section, we summarize our findings of the primary predictors of robustness over the 1,000km Challenge.

4 Predicting Failures

Robots will inevitably make errors when deployed over extended periods of time. Can we learn to predict the conditions under which the robot is likely to encounter errors? We have investigated this question across two settings: predicting localization failures from the 1,000km Challenge, and learning to predict failures of vision-based obstacle avoidance.

4.1 Analyzing Localization Failures

Over the duration of the 1,000km Challenge, Corrective Gradient Refinement (CGR) [Biswas *et al.*,] was used for localization from September 2011 to January 2014, while Episodic non-Markov Localization (EnML) was used from February 2014 onwards. Table 2 compares the Mean Distance Between Interventions (MDBI) for each map in the Gates-Hillman

¹ <http://www.cs.cmu.edu/~coral/cobot/data.html>

| Map | MDBI | | LTF% | STF% | DF% |
|------|------|-------|------|------|------|
| | CGR | EnML | | | |
| GHC4 | 0.62 | 4.42 | 12.0 | 67.7 | 15.6 |
| GHC5 | 1.23 | 9.49 | 64.2 | 23.8 | 11.9 |
| GHC6 | 8.61 | 9.48 | 69.6 | 22.8 | 6.7 |
| GHC7 | 5.58 | 9.02 | 73.1 | 19.5 | 7.7 |
| GHC8 | 6.04 | 19.36 | 70.5 | 22.0 | 7.0 |
| GHC9 | 5.33 | 20.05 | 67.2 | 24.0 | 7.6 |

Table 2: Mean Distance Between Interventions (MDBI), in km, using CGR and EnML per map over the 1,000km Challenge, along with characteristics of each map: LTF%, the percentage of observations that were LTFs; STF%, the percentage of STFfs; and DF%, the percentage of DFfs.

Center building when using CGR, and when using EnML for localization. The MDBI for EnML is 8.13km, which is significantly higher than the MDBI for CGR, 4.79km, thus demonstrating the higher reliability of EnML for localization in real-world human environments.

The variations in the MDBI across the different floors are due to the differences in the number of movable objects and variations over time between the different floors. Table 2 highlights the different characteristics of each floor by enumerating the median fraction of observations that were LTFs, STFfs, and DFfs. Floors GHC6, GHC7, GHC8, and GHC9 were observed to have the fewest number of unmapped observations (STFfs and DFfs), and correspondingly had longer MDBI for both CGR and EnML. Floors GHC4 and GHC5 had significant human traffic and movable objects – note that only 12% of the observations on GHC4 matched the map. These floors are where EnML had a higher MDBI compared to CGR, thus demonstrating the robustness of EnML while localizing in challenging dynamic environments.

4.2 Introspective Vision for Obstacle Avoidance

Vision, as an inexpensive yet information rich sensor, is commonly used for perception on autonomous mobile robots. Unfortunately, accurate vision-based perception requires a number of assumptions about the environment to hold – some examples of such assumptions, depending on the perception algorithm at hand, include purely Lambertian surfaces, texture-rich scenes, and absence of aliasing features or refractive surfaces. Such assumptions are hard to frame mathematically, hard to detect, and may even be unknown a-priori. We introduce an approach for *introspective vision* for obstacle avoidance (IVOA) to automatically predict when a vision system used for obstacle avoidance is likely to produce erroneous results, stemming from such invalid assumptions. By leveraging a supervisory sensor that is occasionally available, IVOA detects failures of stereo vision by divergence in plans between those generated by vision and by the supervisory sensor. By projecting the 3D coordinates where the plans agree and disagree onto the images used for vision-based perception, IVOA generates a training set of reliable and unreliable image patches for perception. We then use this training dataset to learn a model of which image patches are likely to cause failures of the vision-based perception algorithm. Using this model, IVOA is then able to predict whether the relevant image patches in the observed images are likely to




Figure 2: Prediction of failures by IVOA for two different perception models, JPP-C, and ELAS. On the left column, the obstacle grids produced by each of the two perception models are visualized on the input image. The green and red dots represent the detected obstacle-free and occupied cells respectively. On the right, the output of IVOA is shown as the total probability of failure $\Pr(FP) + \Pr(FN)$ for each perception model. IVOA correctly predicts the reflections to cause false positives for both models, and the textureless wall to cause false negatives for JPP-C.

cause failures due to vision (both false positives and false negatives).

Different vision algorithms that make different assumptions about the world, are likely to have differing failure conditions. We demonstrate the ability of IVOA to predict failures specific to different algorithms. For evaluating this ability of IVOA, we used two vision algorithms: a simplified variant of Joint Perception and Planning [Ghosh and Biswas,] that assumes there are no textureless surfaces (JPP-C), and ELAS [Geiger *et al.*, 2010], a dense stereo matching algorithm. IVOA is trained separately using the supervisory input along with each of these vision algorithms, and the two separate models are then used to predict failures of their respective vision algorithms. Figure 2 shows the failures predicted by IVOA for each vision algorithm.

5 Overcoming Failures

An inevitable challenge for long-term autonomy is that no matter how sophisticated, no perception algorithm is perfect, and robots will inevitably make mistakes. From our own 1,000km Challenge, we showed how EnML [Biswas and Veloso, 2017] significantly improves over Markov Localization on the mean distance traversed between interventions, but there were still a number of inevitable interventions (Section 4.1). To compound the problem of inevitable failures, state of the art robots deployed in the real world need expert roboticists to come to their aid to overcome errors. There are two reasons for this: first, algorithms make *strong assumptions about their models*, and second, they depend on *precise parameter values* for robust execution. Fixing such domain-specific parameter values, or inducing the algorithms to overcome their assumptions, are well beyond what non-roboticists are capable of. However, while non-technical users may not

know about the internal workings of algorithms, when robots make mistakes in terms of perceptual estimates or actions, it is often easy to understand what the *correction* should be, in terms of the correct perceptual estimate, or the correct action to take. With this key insight, we introduced human in the loop SLAM [Nashed and Biswas, b] for correcting errors in mapping, and interactive robot transition repair with SMT [Holtz *et al.*,] for correcting action errors.

5.1 Human in the Loop SLAM

Building large-scale, globally consistent maps is a challenging problem, made more difficult in environments with limited access or sparse features, or when using data collected by novice users. For such scenarios, where state-of-the-art mapping algorithms produce globally inconsistent maps, we introduced Human-in-the-Loop Simultaneous Localization and Mapping (HITL-SLAM) [Nashed and Biswas, b], which incorporates approximate, non-expert corrections into pose graph SLAM, while explicitly reasoning about the potential inaccuracies in the corrections. HITL-SLAM graphically presents the map estimate to the user, and allows the user to paint over sets of observations to specify relative geometric constraints between them – collinear, parallel, coincident, or perpendicular. There are three challenges to incorporating such corrections: 1) the user input will be ambiguous, and HITL-SLAM needs to interpret it; 2) the corrections may be partial, with the user only selecting part of a larger single entity such as a wall; and 3) the corrections need to be applied jointly with the existing correlations between observations in the SLAM problem. HITL-SLAM overcomes the first two problems by running an EM optimization to interpret the user input, and the third problem by using a novel set of *human correction factors* in a factor graph SLAM formulation to jointly solve for the map estimate with user corrections.

5.2 SMT-Based Robot Transition Repair

When a robot makes errors while performing actions instead of during perception, we use a different approach to interactive corrections. Complex robot actions are typically structured as state machines, where each state encapsulates a feedback controller. Even if each state is correct, the transitions between states depend on parameters that are hard to get right, even for experienced roboticists. It is very common for parameter values to work in simulation but fail in the real world, to work in one physical environment but fail in another, or to work on one robot but fail with another.

We introduced SMT-based Robot Transition Repair (SRTR) [Holtz *et al.*,], an approach to adjusting the parameters of robot state machines using *sparse, partial specifications* of corrections from humans. During execution, we log execution traces of the transition function. After execution, the human examines the execution trace and corrects a handful of transitions as needed. SRTR then takes as input the transition function source code, the trace, and the corrections, and produces adjustments to the parameter values in three major steps. 1) It parses the transition function code and converts it to an abstract syntax tree for repair. 2) It uses a lightweight *program analysis* to identify parameters

that cannot be repaired, and for each user-provided correction, it *partially evaluates* the transition function for the inputs and variable values at the time of correction, yielding *residual transition functions*. 3) Finally, it uses the residual transition functions to formulate an optimization problem for an off-the-shelf MaxSMT solver [Bjørner *et al.*, 2015]. The solution to this problem is an adjustment to the parameter values that satisfies as many human corrections as possible while simultaneously minimizing adjustments to the parameters.

The full paper for SRTR [Holtz *et al.*,] includes extensive empirical results demonstrating that SRTR 1) finds parameter adjustments fast; 2) produces adjustments that perform as well as those found by exhaustive search; 3) finds adjustments that generalize to new scenarios; and 4) outperforms a domain-expert at repairing state machines.

6 Discussion and Conclusion

In this paper, we presented our research towards the ultimate goal of “always on” autonomous mobile robots in everyday human environments. While significant progress has been made towards this goal, it still remains elusive. There are many topics of active research in support of this goal. On the topic of long-term mapping and localization, there is growing interest in using *semantic-level* understanding of environments to reason about changes in the world. Simultaneously, *competency-aware* machine learning and adversarial machine learning are promising areas of research with the aim of accounting for, predicting, and overcoming inevitable errors made by autonomous systems, including robots. Finally, there is considerable interest in programming and correcting robots through non-technical means of input, including *interactions with novice users*. We are excited to contribute to, and support the exciting new developments along these directions and others in support of long-term autonomy.

Acknowledgements

I would like to thank my collaborators and mentors, including Manuela Veloso, Reid Simmons, Arjun Guha, and Shlomo Zilberstein – and also my students Samer Nashed, Jarrett Holtz, Sadegh Rabiee, Kyle Vedder, Spencer Lane, Sourish Ghosh, David Balaban, Alyxander Burns, and Emily Pruc. This work was partially supported by AFRL and DARPA under agreement #FA8750-16-2-0042, and by NSF award IIS-1724101.

References

- [Biber and Duckett, 2005] Peter Biber and Tom Duckett. Dynamic maps for long-term operation of mobile service robots. In *RSS 2005*, pages 17–24, 2005.
- [Biswas and Veloso, a] Joydeep Biswas and Manuela Veloso. Depth camera based indoor mobile robot localization and navigation. In *ICRA 2012*, pages 1697–1702.
- [Biswas and Veloso, b] Joydeep Biswas and Manuela Veloso. Episodic non-markov localization: reasoning about short-term and long-term features. In *ICRA 2014*, pages 3969–3974.
- [Biswas and Veloso, 2014] Joydeep Biswas and Manuela Veloso. Model-instance object mapping. In *RoboCup Symposium: Robot World Cup XVIII*, pages 525–536, 2014.
- [Biswas and Veloso, 2016] Joydeep Biswas and Manuela Veloso. The 1,000-km challenge: Insights and quantitative and qualitative results. *IEEE Intelligent Systems*, 31(3):86–96, 2016.
- [Biswas and Veloso, 2017] Joydeep Biswas and Manuela Veloso. Episodic non-markov localization. *Robotics and Autonomous Systems*, 87:162 – 176, 2017.
- [Biswas *et al.*,] Joydeep Biswas, Brian Coltin, and Manuela Veloso. Corrective gradient refinement for mobile robot localization. In *IROS 2011*, pages 73–78.
- [Bjørner *et al.*, 2015] Nikolaj Bjørner, Anh-Dung Phan, and Lars Fleckenstein. νZ -an optimizing SMT solver. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, pages 194–199, 2015.
- [Geiger *et al.*, 2010] Andreas Geiger, Martin Roser, and Raquel Urtasun. Efficient large-scale stereo matching. In *Asian Conference on Computer Vision*, pages 25–38, 2010.
- [Ghosh and Biswas,] Sourish Ghosh and Joydeep Biswas. Joint perception and planning for efficient obstacle avoidance using stereo vision. In *IROS 2017*, pages 1026–1031.
- [Holtz *et al.*,] Jarrett Holtz, Arjun Guha, and Joydeep Biswas. Interactive Robot Transition Repair With SMT. In *IJCAI 2018*, pages 4905–4911.
- [Krajník *et al.*, 2017] Tomáš Krajník, Jaime P Fentanes, Joao M Santos, and Tom Duckett. Fremen: Frequency map enhancement for long-term mobile robot autonomy in changing environments. *IEEE Transactions on Robotics*, 33(4):964–977, 2017.
- [Meyer-Delius *et al.*, 2012] Daniel Meyer-Delius, Jürgen Hess, Giorgio Grisetti, and Wolfram Burgard. Temporary maps for robust localization in semi-static environments. In *IROS 2012*, pages 5750–5755, 2012.
- [Nashed and Biswas, a] Samer Nashed and Joydeep Biswas. Curating long-term vector maps. In *IROS 2016*, pages 4643–4648.
- [Nashed and Biswas, b] Samer Nashed and Joydeep Biswas. Human-in-the-Loop SLAM. In *AAAI 2018*, pages 1503–1510.
- [Saarinen *et al.*,] Jari Saarinen, Henrik Andreasson, and Achim J. Lilienthal. Independent markov chain occupancy grid maps for representation of dynamic environment. In *IROS 2012*, pages 3489 –3495.
- [Stachniss and Burgard,] Cyrill Stachniss and Wolfram Burgard. Mobile robot mapping and localization in non-static environments. In *AAAI 2005*, pages 1324–1329.
- [Veloso *et al.*,] Manuela Veloso, Joydeep Biswas, Brian Coltin, Stephanie Rosenthal, Tom Kollar, Cetin Mericli, Mehdi Samadi, Susana Brandao, and Rodrigo Ventura. CoBots: Collaborative robots servicing multi-floor buildings. In *IROS 2012*, pages 5446–5447.