Predictive Exploration Considering Previously Mapped Environments

Daniel Perea Ström Fabrizio Nenci Cyrill Stachniss

***Abstract*— The ability to explore an unknown environment is an important prerequisite for building truly autonomous robots. The central decision that a robot needs to make when exploring an unknown environment is to select the next view point(s) for gathering observations. In this paper, we consider the problem of how to select view points that support the underlying mapping process. We propose a novel approach that makes predictions about the structure of the environments in the unexplored areas by relying on maps acquired previously. Our approach seeks to find similarities between the current surroundings of the robot and previously acquired maps stored in a database in order to predict how the environment may expand in the unknown areas. This allows us to predict potential future loop closures early. This knowledge is used in the view point selection to actively close loops and in this way reduce the uncertainty in the robot's belief. We implemented and tested the proposed approach. The experiments indicate that our method improves the ability of a robot to explore challenging environments and improves the quality of the resulting maps.**

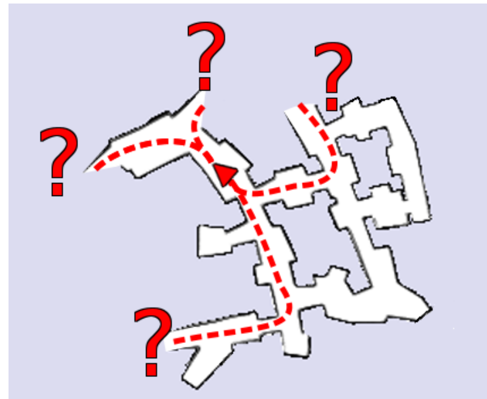

Fig. 1. Mobile robot exploration has to answer the question: "Where to go next?". Our approach exploits previously mapped environments to predict potential future loop closures and thus to select better target locations.

I. INTRODUCTION

Exploration is the task of selecting view points so that a robot can cover the environment with its sensors to build a map. Most exploring robots always start from scratch and do not use any background knowledge about the environment or typical environments. This may be seen as suboptimal as we humans also reason about typical structures even exploring an unknown environment.

While exploring the environment, a robot has to make decisions about where to go and which area to inspect in order to build a model of the environment, see Figure 1 for a small example. The decision of which place to visit can impact the underlying mapping system and can thus be critical for the quality of the resulting map. A typical approach to exploration is the frontier-based approach proposed by Yamauchi [28]. It identifies the frontiers between the free space and unknown areas and guides the robot to the closest one. This strategy typically yields short exploration paths but is generally unaware about the uncertainty in the robot belief, for example, about its current position in the world. Information-theoretic approaches such as [13], [2], [23], [22] consider the expected information gain to evaluate possible target locations. A key challenge, for example in [23], is to reason about the possible environment that the robot may experience when navigating through unknown environments.

Daniel Perea Ström is with Universidad de La Laguna, Departamento de Ingeniería Informática, 38206 Tenerife, Spain. Fabrizio Nenci and Cyrill Stachniss are with the University of Bonn, Inst. of Geodesy and Geoinformation, 53115 Bonn, Germany.

This work has partially been supported by the European Commission under grant agreement No. FP7-ICT-600890-ROVINA, by the DFG under contract number FOR 1505 "Mapping on Demand", and by the Agencia Canaria de Investigación, Innovación y Sociedad de la Información (ACIISI), co-funded by the European Fund for Regional Development (FEDER).

In this paper, we take first steps towards exploiting background knowledge during autonomous exploration. The key idea is to use previously experienced environments to reason about what to find in the unknown parts of the world. Thus, we equip the robot with a database to store all acquired local maps and exploit this knowledge when selecting target locations. Our research is motivated by an autonomous exploration project for autonomously digitizing the Roman catacombs, which are complex underground environments with repetitive structures. To predict possible structures the robot may experience during exploration, we exploit previously visited areas and consider the similarities with the area around the current frontiers. This allows the robot to actively seek for loop-closures and in this way actively reduce its pose uncertainty. Our experiments indicate that this approach is beneficiary for robots when comparing it to a standard frontier-based method.

II. RELATED WORK

The majority of techniques for mobile robot exploration focus on generating motion commands that minimize the time needed to cover the whole terrain. Several techniques also assume that an accurate position estimate is available during exploration [10], [28]. Whaite and Ferrie [26] present an approach that uses the entropy to measure the uncertainty in the geometry of objects that are scanned with a laser range sensor. Similar techniques have been applied to mobile robots [21], [16], but such approaches still assume to know the correct pose of the vehicle. None of the approaches mentioned above takes the pose uncertainty into account when selecting the next vantage point. There are, however, exploration approaches that have been shown to be robust against uncertainties in the pose estimates [5], [9].

Besides the idea of navigating to the next frontier [28], techniques based on stochastic differential equations for goal-directed exploration have been proposed by Shen *et al.* [19]. Similar to that, constrained partial differential equations that provide a scalar field into unknown areas have been presented by Shade et al. [18]. An information-theoretic formulation that seeks to minimize the uncertainty in the belief about the map and the trajectory of the robot has been proposed by Stachniss *et al.* [23]. This approach builds upon the works of Makarenko *et al.* [13] and Bourgault *et al.* [2]. Both extract landmarks out of laser range scans and use an Extended Kalman Filter to solve the underlying SLAM problem. They furthermore introduce an utility function which trades-off the cost of exploring new terrain with the potential reduction of uncertainty by measuring at selected positions. A similar technique has been presented by Sim *et al.* [20], who consider actions to guide the robot back to a known place in order to reduce the pose uncertainty of the vehicle.

In general, the computation of the expected entropy reductions is a complex problem, see Krause and Guestrin [11], and in all real world systems, approximations are needed. Suitable approximations often depend on the environment model, the sensor data, and the application. In some cases, efficient approximations can be found, for example in the context of monitoring lakes using autonomous boats [7].

Other approaches, especially in the context of autonomous micro aerial vehicles (MAVs), seek to estimate the expected feature density in the environment in order to plan a path through areas that support the helicopter localization [17]. This can be seen as related to information-theoretic approaches, although Sadat et al. [17] do not formulate their approach in this framework. A related approach to MAV exploration seeks to select new vantage points during exploration, so that the expected number of visible features is maximized, see Mostegel et al. [14].

An interesting approach by Fox *et al.* [6] aims at incorporating knowledge about *other* environments into a cooperative mapping and exploration system for multiple robots. This allows for predicting simplified laser scans of an unknown environment. This idea was an inspiration for our paper for predicting possible loop closures given the environment structure explored so far. We use this approach for exploring ancient catacombs, which are repetitive underground environments, with a mobile platform, see Figure 2. Chang *et al.* [3] propose an approach for predicting the environment using repetitive structures for SLAM. Other background knowledge about the environment, for example semantic information, can support the exploration process as shown by Wurm et al. [27], Stachniss *et al.* [24] as well as Holz *et al.* [8].

III. Environment Predictive Exploration

The central question in exploration is "Where to go?". Several different cost functions for making the decision of

Fig. 2. A robot for exploring and digitizing Roman catacombs was the motivation for our research.

where to go next can be defined. The most popular one goes back to Yamauchi [28], who guides the robot to the closest unexplored location. Yamauchi introduces the concept of frontiers, which are the cells of an occupancy grid map at the boundary between the free and the unexplored space. In the standard setting, this approach seeks to minimize the time that is needed to cover the environment with the robot's sensors and is a popular choice in mobile robotics.

A. Information-Driven Exploration

Given the fact that most real robots maintain a probabilistic belief about their pose and the map of the environment, an alternative approach is to select the target location that is expected to minimize the uncertainty in the belief of the robot. In this setting, the exploration problem can be formulated as follows. At each time step t, the robot has to decide which action a to execute, i.e., where to move next. During the execution of a, it is assumed that the robot obtains a sequence of observations z (for better readability, we neglect all time indices).

Thus, we can define the expected information gain, also called mutual information, of selecting the action a as the expected change in entropy in the belief about the robot's poses X and the map M:

$$I(X, M; Z^a) = H(M, X) - H(M, X \mid Z^a). \quad (1)$$

The second term in Eq. (1) is the conditional entropy and is defined as

$$H(M, X \mid Z^a) = \int p(z \mid a)\, H(M, X \mid Z^a = z)\, dz. \quad (2)$$

Unfortunately, reasoning about all potential observation sequences z in Eq. (2) is intractable in nearly all real world applications since the number of potential measurements grows exponentially with the dimension of the measurement space and with time. It is therefore crucial to approximate the integral of Eq. (2) so that it can be computed efficiently with sufficient accuracy.

A suitable approximation, however, depends on the environment model, the sensor data, and the application so that no general one-fits-all solution is available. Given our previous work [23], we considered different types of actions: First, *exploration actions* that guide the robot to the closest frontier and this reduces the map uncertainty. As we have no further information about the unseen area, it is difficult

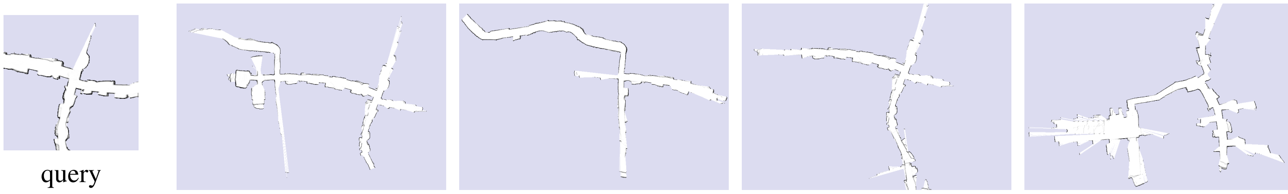

Fig. 3. Example of the submap retrieval using FabMAP2. The left image shows the query map, the other ones the best four matches from the database.

to distinguish two frontiers with respect to the expected uncertainty reduction. Second, *loop-closing and re-localization actions*, which are key to the uncertainty reduction about the robot's pose.

In this work, we aim at combining these types of actions into a single one. We seek to predict what the so far unseen environment beyond a frontier *may* look like, based on background knowledge of previously seen environments, and select the frontier that potentially leads to a loop-closure. In this way, we maximize the expected uncertainty reduction in the belief of the robot about the state of the world.

B. Utility Function for Exploration

Most exploration systems define a utility function to relate the expected gain in information with the cost of obtaining the information. As long as no constraints such as available energy or similar are considered, the distance that the robot has to travel to obtain its measurements is a standard choice. This yields a utility function of the form

$$U(a) = I(M, Z; Z^a) - \mathrm{cost}(a) \tag{3}$$

so that the task of selecting the best action can be formulated as

$$a^* = \operatorname*{argmax}_a I(M, Z; Z^a) - \mathrm{cost}(a). \tag{4}$$

Throughout this work, we define the cost function $\mathrm{cost}(a)$ as the path length corresponding to action a, i.e. the length of the trajectory from the current location of the robot to the designated target location.

As mentioned in the previous section, estimating the expected information gain is challenging and computationally demanding and thus we use the following approximation. We assume that actions can reduce the robot's uncertainty about the map by exploring unseen areas and/or can reduce its uncertainty about the trajectory by closing a loop.

$$a^* = \operatorname*{argmax}_a I_{map}(a) + I_{traj}(a) - \mathrm{cost}(a). \tag{5}$$

As we do not know how large the unknown area and thus the number of unknown grid cells behind a frontier is, we may argue that all frontiers yield the same expected information gain with respect to the map uncertainty. Thus, we can simplify Eq. (5) as long as we consider only exploration actions to frontiers:

$$a^* = \operatorname*{argmax}_a I_{traj}(a) - \mathrm{cost}(a). \tag{6}$$

The expected information gain about the trajectory $I_{traj}(a)$ is mainly influenced by loop closures. The more likely a loop closure can be obtained when executing an exploration action a, the higher its expected gain. Thus, the remainder of this section addresses the problem of predicting possible loop closures.

C. Predictive Exploration

The key contribution of this work is to model the predictive belief describing what the environment may look like in the unexplored areas. To compute this belief, the robot exploits environment structures it has seen in the past—either in the environment explored so far or even from previous mapping runs. Our exploration system uses this predictive belief to evaluate the frontiers as possible target locations for the exploration. This allows us to select the frontiers that are likely to lead to a loop-closure and thus to an active reduction of the uncertainty in the robot's belief. As we will show later during the experimental evaluation, our approach outperforms the traditional frontier-based exploration system.

D. Querying for Similar Environment Structures

The key idea of this belief is to look for similarities between the known areas around a frontier and portions of previously mapped environments. Under the assumption that environments are not random but expose certain structures and that these structures tend to appear more than once, we can use the already mapped areas in order to predict what the environment beyond the frontier may look like.

The first step is to look for portions of the already mapped environment(s) that are similar to the area around the frontier for which the prediction should be performed. To do this, we incrementally build a database storing all local grid maps. To perform a similarity query, we compare our local maps with the maps stored in the database. To avoid a large number of expensive map-to-map comparisons to search for similar submaps, we rely on a bag-of-words inspired approach, a technique that is frequently used in computer vision to search for image similarities. More concretely, we apply FabMAP2 by Cummins and Newman [4], an appearance based approach to efficiently query the database. Although FabMAP2 was originally designed to match camera images, it turns out that we can also use it to effectively search for local grid maps in a large database of maps. As FabMAP2 also provides a likelihood $l(m)$ for each match m, we can obtain a belief about possible environment structures. Figure 3 shows an illustration of this procedure. The image on the left is a query image and the other images are the top 4 matches reported by FabMAP2.

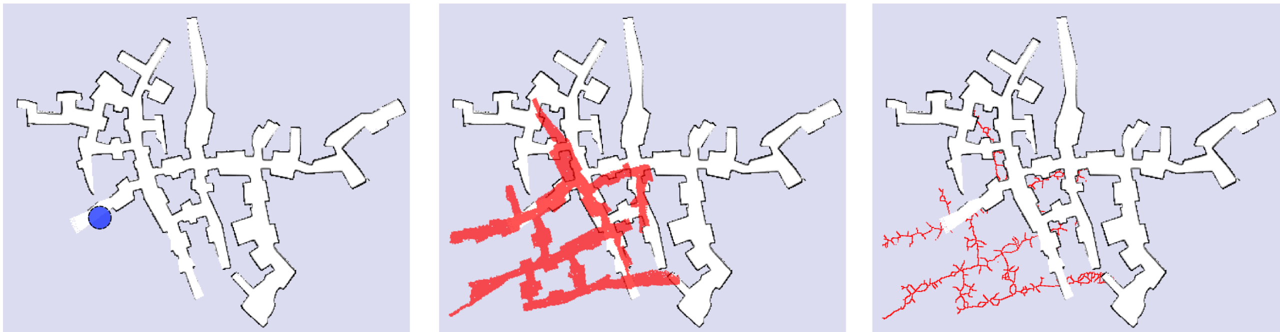

Fig. 4. Illustration of the loop closures prediction. Left: So far explored map with the frontier under consideration (blue circle). Middle: One map from the predictive belief (in red) superimposed on the map explored so far. Right: Voronoi diagram used for the path search.

E. Loop Closures Prediction

As we are mainly interested in the possible paths through the unknown environment in order to find loop closures and not necessarily the exact geometry, we reduce the maps reported by FabMAP2 to extended Voronoi graphs [1] and do all further computations on these graphs.

FabMAP2 provides us with candidates for matching maps but no geometric alignment between the query map and the reported ones. Thus, we align each map reported by FabMAP2 with our query map. This can be done in a robust manner through a RANSAC-based alignment of the Voronoi graphs using its junction points. Figure 4 shows an example of a Voronoi graph aligned with the map explored so far.

The next step, is to search for possible loop closures, for which we use the generalized Voronoi graph. Starting from the frontier point, we traverse the Voronoi graph in a breadth-first manner. During the traversal, we check if the Voronoi graph leads to a position that is close to any other frontier in the map built so far. If this is the case, we regard that as a possible loop closure. Such a situation is illustrated in the left image of Figure 5. This process is executed for each frontier.

F. Estimating the Probability of Closing a Loop

Each map reported by FabMAP2 comes with a likelihood. Thus, we can approximate the probability of closing a loop when executing an exploration action as

$$\mathbf{S}_f = \sum_{m \in \mathcal{M}(f)} l(m) \sum_{c \in \mathcal{C}(f,m)} l(c \mid m) \tag{7}$$

Here, $\mathcal{M}(f)$ is the set of matches returned by FabMAP2 when querying with the frontier f, and $l(m)$ the likelihood of a match m. The term $\mathcal{C}(f, m)$ refers to the set of possible loop closures computed according to the breadth-first traversal explained above and $l(c \mid m)$ is the likelihood that the loop closures can be reached. We assume that $l(c \mid m)$ is proportional to the inverse length of the path of the predicted loop closure. This means that short loop closures are more likely than long ones.

Assuming that every executed loop closure through unknown areas of the map yields the same expected uncertainty reduction, we can approximate the expected information gain I_{traj} of Eq. (6) with the score $\mathbf{S}_f$ according to Eq. (7). This is clearly a strong assumption but we argue that a high score indicates a high expected gain from exploring the frontier.

IV. EXPERIMENTS

The experiments are designed to illustrate the advantages of our predictive exploration approach. We show that our approach selects frontiers that lead to loop closures which in turn result in improved maps of the environment. As a baseline, we use a standard frontier-based exploration approach.

The underlying mapping framework that is used for all experiments is a state-of-the-art graph-based SLAM system that relies on laser range data. The backend is g2o [12] and the frontend uses FLIRT features to search for possible data associations [25], uses correlative scan matching to align scans, and applies single cluster graph partitioning to resolve ambiguities as proposed by Olson [15]. The exploration system is integrated into the mapping framework and has been implemented using ROS.

A. Map Consistency

First, we compare the quality of the maps obtained with frontier-based exploration vs. our predictive exploration. The environments considered here are parts of the Roman catacomb St. Priscilla, a difficult to traverse and large-scale underground environment in Rome. As the robot is equipped with tracks, see Figure 2, its odometry is in general worse than the one of a wheeled robot and it often reveals a bias to one side. This experiment has been conducted in simulation but the environment actually represents the catacomb, with each experiment covering an exploration area of 2,500 m^2. Odometry noise is simulated following a probabilistic motion model, sampling over normal distributions for each motion parameter, with bias for the rotation distribution.

Figure 6 illustrates the obtained results for two environments using exactly the same mapping system and identical parameters for the comparison. The images on the left are the ground truth 2D map used for the simulation. The images in the second column correspond to the results of the frontier-based exploration, while the images on the right show our approach. As can be seen already visually, our approach yielded a consistent model of the environment, while the

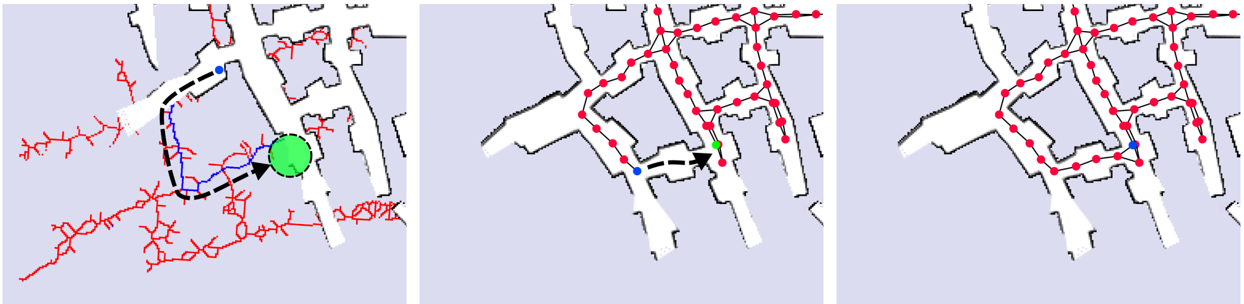

Fig. 5. Illustration of the active loop closing. Left: prediction of the possible path with the loop closure shown in blue. Middle: the robot explores the path along the predicted loop closure and perceives the actual structure of the scene. The graph in the already explored environment shows the pose graph of the SLAM system. Right: successful loop closure Please note that the predicted environment is actually not identical with the real environment but reveals a similar structure. This similarity resulted in the shown loop closure.

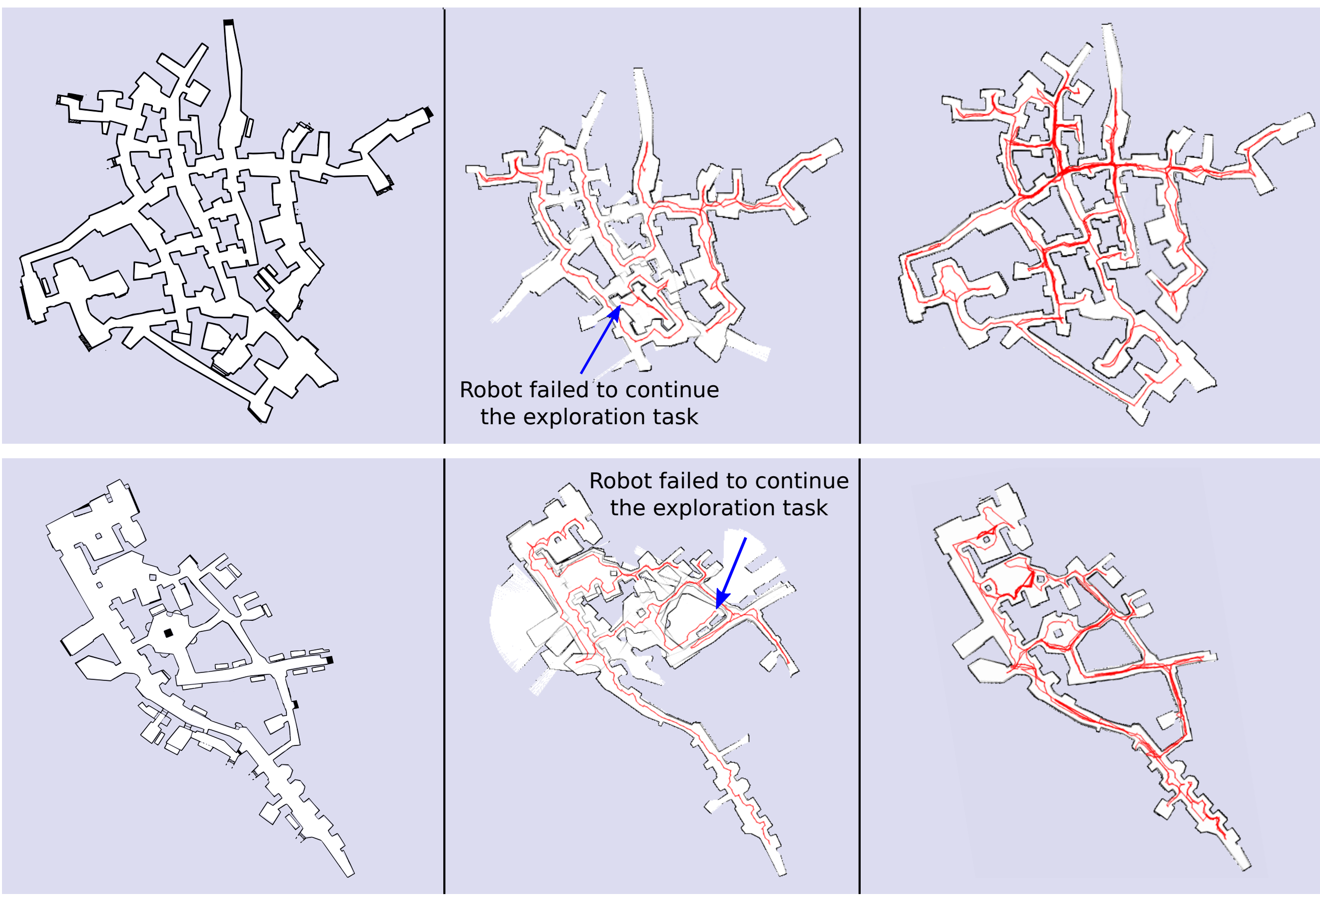

Fig. 6. Two performance comparisons in constant odometry bias scenario. On the left, the original map. In the middle, the closer frontier approach. On the right, our prediction-based approach. Note that the nearest frontier approach produces a map that is non consistent with the original one, so that the robot gets actually lost in it. The map produced by the prediction-based approach is instead consistent with the original one.

frontier-based approach failed. Using the frontier-based approach the robot was unable to continue its exploration task due to an inconsistent map that prevented the computation of further exploration actions. This was the case in all six exploration experiments that we conducted in St. Priscilla environment.

B. Path Length

The advantages of the prediction-based approach come at a cost—the cost of traversing exploration paths that are longer than the ones generated by the frontier-based approach. This experiment is designed to evaluate the increase in path length.

As we are not able to obtain consistent maps for the frontier-based approach under a realistic noise model for the task under consideration, we set the noise to zero in the simulator and repeated the previous experiments. Using a zero noise odometry, also the frontier-based approach is able to build consistent maps. In this settings there is no advantage in using our predictive approach as the pose uncertainty is zero

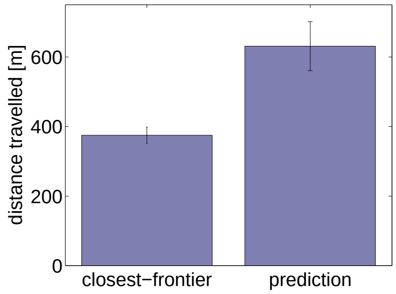

Fig. 7. Mean and standard deviation of the distances travelled in the frontier-based approach and in the proposed approach.

and no uncertainty reduction is gained from closing loops. We compared the distances travelled for the frontier-based and our approach. The distances travelled are summarized in Figure 7. In the worst case scenario, the path generated by our approach was 1.85 times longer than the one of the frontier-based approach. The minimum increase was a factor of 1.5. Generating on average approximatively a 1.7 times longer trajectory is clearly an overhead—for actively closing loops and in this way reducing uncertainty, however, this price must be paid.

V. Conclusions

In this paper, we proposed a novel approach for autonomous exploration of unknown environments. The key contribution of this work is a technique to predict possible environment structures in the unseen parts of the robot's surroundings based on previously explored environments. In our approach, we exploit this belief to predict possible loop closures that the robot may experience when exploring an unknown part of the scene. This allows the robot to actively reduce the uncertainty in its belief through its exploration actions. We implemented and tested our approach. Our experiments illustrate that our technique allows for an effective exploration of difficult to map environments. By actively closing loops, we are able to obtain consistent maps of the environment. In contrast to that, a traditional frontier-based exploration approach is not able to successfully explore the scene.

References

[1] P. Beeson, N.K. Jong, and B. Kuipers. Towards autonomous topological place detection using the extended voronoi graph. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2005.

[2] F. Bourgoult, A.A. Makarenko, S.B. Williams, B. Grocholsky, and F. Durrant-Whyte. Information based adaptive robotic exploration. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, Lausanne, Switzerland, 2002.

[3] H.J. Chang, C.S.G. Lee, Y. Lu, and Y.C. Hu. P-slam: Simultaneous localization and mapping with environmental-structure prediction. *IEEE Transactions on Robotics*, 23(2):281–293, 2007.

[4] M. Cummins and P. Newman. Highly scalable appearance-only slam fab-map 2.0. In *Proc. of Robotics: Science and Systems (RSS)*, 2009.

[5] T. Duckett, S. Marsland, and J. Shapiro. Fast, on-line learning of globally consistent maps. *Journal of Autonomous Robots*, 12(3):287 – 300, 2002.

[6] D. Fox, J. Ko, K. Konolige, and B. Stewart. A hierarchical bayesian approach to the revisiting problem in mobile robot map building. In *Proc. of the Int. Symposium of Robotics Research (ISRR)*, 2003.

[7] G. Hitz, A. Gotovos, F. Pomerleau, M.-E. Garneau, C. Pradalier, A. Krause, and R.Y. Siegwart. Fully autonomous focused exploration for robotic environmental monitoring. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2014.

[8] D. Holz, N. Basilico, F. Amigoni, and S. Behnke. A comparative evaluation of exploration strategies and heuristics to improve them. In *Proc. of the European Conference on Mobile Robots (ECMR)*, Oerebro, Sweden, 2011.

[9] J. Ko, B. Stewart, D. Fox, K. Konolige, and B. Limketkai. A practical, decision-theoretic approach to multi-robot mapping and exploration. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, pages 3232–3238, Las Vegas, NV, USA, 2003.

[10] S. Koenig and C. Tovey. Improved analysis of greedy mapping. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, Las Vegas, NV, USA, 2003.

[11] A. Krause and C. Guestrin. Near-optimal nonmyopic value of information in graphical models. In *Proc. of Uncertainty in Artificial Intelligence (UAI)*, 2005.

[12] R. Kümmerle, G. Grisetti, H. Strasdat, K. Konolige, and W. Burgard. g2o: A general framework for graph optimization. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, pages 3607–3613, 2011.

[13] A.A. Makarenko, S.B. Williams, F. Bourgoult, and F. Durrant-Whyte. An experiment in integrated exploration. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2002.

[14] C. Mostegel, A. Wendel, and H. Bischof. Active monocular localization: Towards autonomous monocular exploration for multirotor mavs. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, Hong Kong, China, 2014.

[15] E. Olson. Recognizing places using spectrally clustered local matches. *Robotics and Autonomous Systems*, 57(12):1157–1172, 2009.

[16] R. Rocha, J. Dias, and A. Carvalho. Exploring information theory for vision-based volumetric mapping. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, pages 2409–2414, Edmonton, Canada, 2005.

[17] S.A. Sadat, K. Chutskoff, D. Jungic, J. Wawerla, and R. Vaughan. Feature-rich path planning for robust navigation of mavs with monoslam. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, Hong Kong, China, 2014.

[18] R. Shade and P. Newman. Choosing where to go: Complete 3d exploration with stereo. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2011.

[19] S. Shen, N. Michael, and V. Kumar. 3d indoor exploration with a computationally constrained mav. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2012.

[20] R. Sim, G. Dudek, and N. Roy. Online control policy optimization for minimizing map uncertainty during exploration. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2004.

[21] C. Stachniss and W. Burgard. Mapping and exploration with mobile robots using coverage maps. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, pages 476–481, 2003.

[22] C. Stachniss and W. Burgard. Particle filters for robot navigation. 3(4):211–282, 2012. Published 2014.

[23] C. Stachniss, G. Grisetti, and W. Burgard. Information gain-based exploration using rao-blackwellized particle filters. In *Proc. of Robotics: Science and Systems (RSS)*, pages 65–72, Cambridge, MA, USA, 2005.

[24] C. Stachniss, O. Martinez Mozos, and W. Burgard. Efficient exploration of unknown indoor environments using a team of mobile robots. *Annals of Mathematics and Artificial Intelligence*, 52:205ff, 2009.

[25] G.D. Tipaldi and K.O. Arras. Flirt–interest regions for 2d range data. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, pages 3616–3622, 2010.

[26] P. Whaite and F. P. Ferrie. Autonomous exploration: Driven by uncertainty. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(3):193–205, 1997.

[27] K.M. Wurm, C. Stachniss, and W. Burgard. Coordinated multi-robot exploration using a segmentation of the environment. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2008.

[28] B. Yamauchi. Frontier-based exploration using multiple robots. In *Proc. of the Second International Conference on Autonomous Agents*, pages 47–53, Minneapolis, MN, USA, 1998.