

Stephane Bazeille*, Emmanuel Battesti and David Filliat

A Light Visual Mapping and Navigation Framework for Low-Cost Robots

Abstract: We address the problems of localization, mapping, and guidance for robots with limited computational resources by combining vision with the metrical information given by the robot odometry. We propose in this article a novel light and robust topometric simultaneous localization and mapping framework using appearance-based visual loop-closure detection enhanced with the odometry. The main advantage of this combination is that the odometry makes the loop-closure detection more accurate and reactive, while the loop-closure detection enables the long-term use of odometry for guidance by correcting the drift. The guidance approach is based on qualitative localization using vision and odometry, and is robust to visual sensor occlusions or changes in the scene. The resulting framework is incremental, real-time, and based on cheap sensors provided on many robots (a camera and odometry encoders). This approach is, moreover, particularly well suited for low-power robots as it is not dependent on the image processing frequency and latency, and thus it can be applied using remote processing. The algorithm has been validated on a Pioneer P3DX mobile robot in indoor environments, and its robustness is demonstrated experimentally for a large range of odometry noise levels.

Keywords: Visual loop-closure detection, topological SLAM, path following.

DOI 10.1515/jisys-2014-0116

Received August 1, 2014; previously published online February 14, 2015.

1 Introduction

To navigate in an unknown environment, a robot requires the ability to build a map and to localize itself using a process named simultaneous localization and mapping (SLAM) [42]. The field of SLAM can be broadly divided into topological and metrical approaches. The topological approach models the environment as a graph of discrete locations and often leads to simple solutions [8, 19]. It is usually light and suitable for many kinds of environments and for human interaction. Building the map requires precise sensor data association to detect when the robot comes back to a previously visited place (a process called loop-closure detection [2, 12]). Its main drawback is the lack of geometric information about the robot surroundings that reduces path planning to searching a path in the corresponding graph. On the contrary, metrical maps are explicitly based on measured distances and positions (e.g., Refs. [13, 42]). The representation of the environment is geometric and clearly corresponds to the real world. With these maps, localization can be done continuously, and planned navigation is easier and more precise. The major problem is to ensure geometrical consistency between position and perception, which is computationally expensive and makes the map more difficult to build.

A number of approaches have attempted to capitalize on the advantages of the two representations. For instance, local metrical maps can be embedded into graphs to enhance scalability [17]. Other graph-based

*Corresponding author: **Stephane Bazeille**, Istituto Italiano di Tecnologia (IIT), Department of Advanced Robotics, Via Morego 30, 16163 Genova, Italy, e-mail: stephane.bazeille@gmail.com

Emmanuel Battesti and David Filliat: ENSTA ParisTech, INRIA FLOWERS Team, 828 boulevard des Marechaux, 91762 Palaiseau Cedex, France

solutions can be used to infer a precise metrical position of the robot, while still allowing for large-scale mapping [29]. In this article, we propose such an approach that builds topometric maps of the environment using a camera and robot odometry (Figure 1). This approach keeps the simplicity of topological approaches by using only simple and fast appearance-based image processing, while making it possible to navigate more robustly in the map, thanks to the metric information provided by the odometry.

Vision sensors are often used for topological mapping. They indeed provide many advantages such as small size and price, light weight, and low energy consumption, and above all, a rich environmental information that is usable as the only information about the environment. Vision sensors could therefore make it possible to perform navigation on small robots; however, the embedded processing power limits the algorithmic complexity of the solutions. A possibility to solve this problem is to execute part of the processing

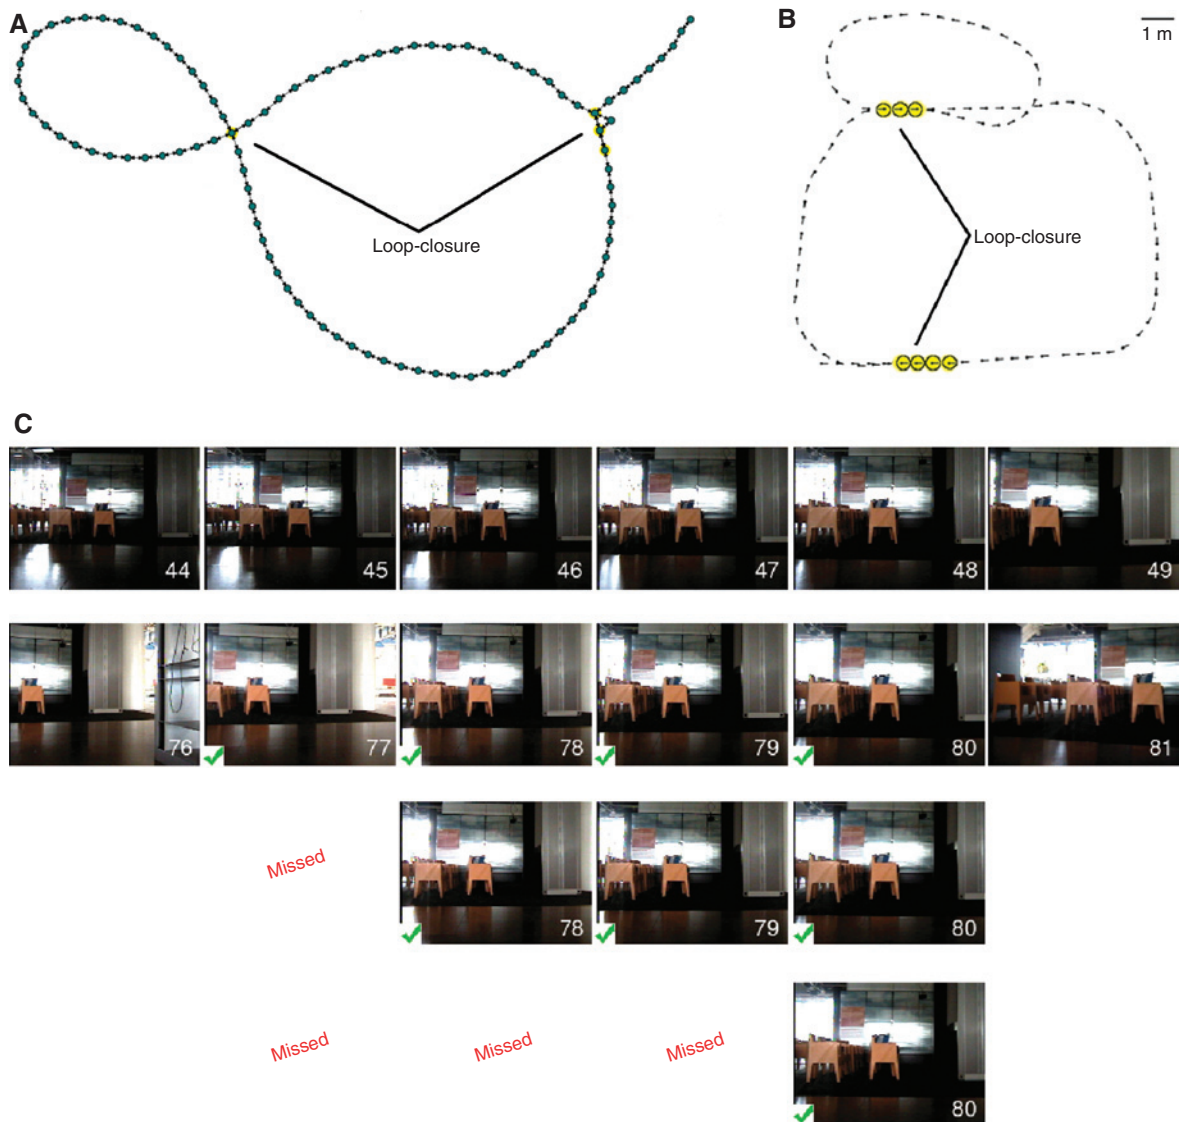


Figure 1. Museum Sequence.

Improvement of the topological map by adding metric information. (A) The topological map built using Ref. [2]. (B) The topometric map built with the approach presented in this article. Yellow circles correspond to the loop-closure locations. (C) Illustration of loop-closure detection in these two maps. The first line shows images taken during the first pass with the robot; the second line shows the second pass with the ground truth loop-closure detections checked in green; the third line shows the loop-closure detected with the proposed approach; and the last line shows the loop-closure detected with Ref. [2], which is not using the odometry.

on remote computers; however, the approach should then deal with communication bandwidth and latency. Contrary to visual odometry (e.g., Ref. [33]) or metric approaches (e.g., Ref. [13]) that require high-frame-rate image processing and cannot deal with remote processing, the topometric approach proposed in this article is cheap, light, and suitable for remote computation and communication latencies, as robot odometry allows to guide the robot between availabilities of the visual information.

1.1 Contributions

This article merges and presents advances on our previously published work on topometric mapping using vision and odometry for loop-closure detection [6] and path-following navigation [5]. The main contributions of this article are (i) the introduction of improvements in the Bayesian filter to increase the precision in loop-closure detection to allow reliable path-following navigation; (ii) the evaluation of the mapping algorithm on large maps; (iii) the thorough evaluation of the path-following behavior depending on robot odometry noise level, occlusions, and changes in the scene; and (iv) a light and fast implementation of our loop-closure framework that can support remote computation.

1.2 Content

In Section 2, we present a review of related work on topological mapping and navigation. In Section 3, we recall our previous work on loop-closure detection and highlight its limitation for topometric mapping and navigation. In Section 4, we detail our visual loop-closure detection framework enhanced with odometry, the mapping, and the navigation methods. In Section 5, we evaluate our system performances. Finally, in Section 6, we discuss the advantages and limitations of our solution.

2 Related Work

Among all the SLAM approaches, we are more particularly interested in vision-based topological SLAM. Many approaches to the topological SLAM problem are based on appearance and rely on omnidirectional vision [8, 23, 43, 44]. A similarity distance between images is defined to set the edges of the map, with similar images considered as originating from the same place and thus as corresponding to the same node. These approaches provide an efficient segmentation of the environment, as omnidirectional images enable to recognize a place independently of the robot orientation. Other approaches [2, 12, 28, 30, 32] use perspective cameras that impose more constraints on location recognition but can be used efficiently as there are some constraints on the robot path such as road or corridors.

During mapping, the problem of recognizing if a new image taken by the robot corresponds to a new place or an already existing place is called loop-closure detection. Several efficient approaches for this problem have been proposed recently [2, 12, 22, 28], many of them relying on the bag of visual words approach [40] that encodes images as an unordered set of local features. It has to be noted that among all these methods, a large part needs camera calibration; however, few methods exist based on uncalibrated cameras [21, 37]. Our method falls in this last category as only the knowledge of the camera field of view is required.

A limitation of topological maps for robot guidance is the lack of information about obstacles and about free space around the robot. This is why navigation methods using topological maps have been limited to following previously learned paths. A first method to follow such a path is visual servoing, which uses the feedback information extracted from images to control the motion of the robot [26]. Those methods generally require camera calibration (homography, fundamental matrix, Jacobian, and removal of lens distortion, e.g., Refs. [9, 38]). Also, some approaches make assumptions on the environment (artificial landmarks, vertical straight lines, parallel walls) or sometimes need more than one camera or specific cameras (e.g.,

omni-directional) [11, 15]. Calibration-free methods without any environment knowledge have been developed [10]; they are based on image features tracking, and use qualitative comparisons of images to control the motion. Such methods are interesting in that they are purely visual but still require heavy computations and are highly dependent of image data quality. Tracking errors, temporary absence of visual information, or changes in the scene lead quickly to system failures.

Another approach to guide robots from a topological map is to complement the map with metrical information so that the robot metrical position can be estimated and used for guidance. This approach leads to topometric maps [6, 7]. A common approach for this is the use of visual odometry. This approach is appealing because it uses the same sensor and processes image sequences taken between nodes to estimate the robot displacement [19, 29, 33, 39]. However, as for visual servoing, this approach is computationally intensive, requires high-frequency image processing, and the results depend strongly on the quality of image matching. A last approach is to use the robot mechanical odometry [18, 36]. While this requires the integration of a second sensor on the robot, it can make the system more robust and lighter from a computational point of view as this relieves the visual system from high-frame-rate computation. Moreover, odometry allows localization for a short time in the absence of visual information due to vision failures (dark or dazzle areas, blurry image, occlusions), important changes in the scene that has been learned (object changes, people, etc.), or network lag in case of remote visual processing.

The main limitation of odometry is the cumulative error due to the integration of noisy measurements, which makes the position estimate more and more inaccurate over time. As a consequence, long-term use of odometry requires complementary information to enable a correction of this cumulative error. Loop-closure detection provides such an information and can be used to estimate a consistent topometric map, an approach also known as pose-graph SLAM [41]. In these approaches, an optimization algorithm is used to correct odometry and estimate the position of nodes that best satisfies loop-closure and odometry constraints. Several relaxation methods have been applied to this problem [16, 25], and efficient recent algorithms can now handle large maps [20, 24, 27, 31, 35] and can also be robust to false loop-closure detection [41].

Our method is different from all the presented ones as it proposes in a single framework, no calibration, low computation with remote processing possibility, and the possibility to build large scale maps and to navigate inside with low-cost platforms.

3 Baseline System and Its Limitations

The approach proposed in the current article relies on visual loop-closure detection [2, 12, 22, 28]. In particular, we started from the approach presented by Angeli et al. [2, 3], who developed a real-time incremental topological SLAM approach without any prior information about the environment, using a monocular calibrated camera. This method was based on visual loop-closure detection using a Bayesian filter and the bag of visual words approach for fast image comparison. It presented many advantages such as its simplicity, speed, lack of learning stage, and efficiency (low false-alarm rate). However, the lack of metrical information (Figure 1A) made the map ill posed for robot navigation.

A straightforward solution to add a metrical information was to extend this algorithm by including the robot's odometry information as the relative position between nodes on each link of the graph and by applying a relaxation algorithm each time a loop closure is detected to enforce map consistency. However, in practice, this approach often leads to maps that are inconsistent with the real world because of several limitations of this purely vision-based loop-closure detection approach.

The first reason was the lack of temporal consistency of the loop-closure detection. We show a typical example in Figure 2 where the matching obtained in Ref. [2] was 4–26, 5–29, 6–27, and 6–28, but the ground truth was 4–26, 5–27, 6–28, and 7–29. While these loop-closure detections are visually correct as the same scene is present in the matched images, the detections are not temporally consistent.

Second, the responsiveness of the algorithm was too low. At least two or three images were needed to validate the first loop closure as the idea was to enforce a very low false-alarm rate (see Figure 1). This is an




Figure 2. Illustration of Loop-Closure Detection Limitations for Topometric Mapping.

A wrong matching is highlighted in red. It shows a system not reactive enough, not temporally consistent, and inaccurate in terms of robot position. In this particular case, we expected the loop-closure: 3–25, 4–26, 5–27, 6–28, and 7–29.

important limitation in indoor environments, where common trajectories are mostly seen when a door is crossed for a very short distance. In these cases, many loop closures are missed and too few constraints are created for map relaxation and odometry correction.

Finally, the loop-closure definition itself was not precise enough. In the original approach [2], as in others such as that in Ref. [12], the loop closures were validated using multiple-view geometry between the matching images. This policy defines loop closure by the fact that the robot sees the same scene but does not enforce that the robot positions are very close. Moreover, the scale ambiguity in multiple-view geometry prevents the easy recovery of the real robot displacement. This strategy prevents false-positive detection but cannot detect well the loop closure corresponding to the smallest robot displacement. To include the graph relaxation algorithm, we therefore had to define a more constrained loop-closure validation stage to only accept loop-closure detections with very close robot positions (see Section 4.1.4).

All these problems have to be solved to allow robust navigation. To solve the responsiveness and consistency problems, in Ref. [6] we replaced the evolution model of the Bayesian filter that applied a diffusion of the probability over the neighboring locations by an odometry-based evolution model (see Figure 5). Through a probabilistic model of odometry, the evolution model can now take into account not only the nodes' topological proximity but also their relative position. We also proposed a new validation method that imposes closer robot positions for loop closure and gives information to navigate, and in Ref. [5] we presented an implementation of a path-following algorithm using the maps previously built in a basic scenario.

In the current article, we summarize in the next section the strategies previously exposed, and present several improvements: an N th-order Bayesian filter, the use of a faster feature detection and a static dictionary in the likelihood computation, and a new fast validation method based on two-dimensional (2D) motion. These modifications improve the whole loop-closure detection framework and allow guidance.

We validate all these improvements on larger maps than previously reported and demonstrate the robustness of our framework by showing navigation in difficult conditions, such as noise on odometry or lack of visual information (people crossing in front of the robot or objects moved in the scene between the learning and the path-following phase). Finally, we describe a light and fast implementation of our loop-closure framework that can be used for mapping when the robot is tele-operated and for autonomous navigation when a map is available. The method is robust to noise on odometry and makes it possible to perform the execution of all image processing on a remote computer, which makes this approach applicable to low-cost robots with wireless connection.

4 Mapping and Navigation Framework

Our approach (summarized in Figure 3) processes images and odometry to build a topometric map of the environment. It is based on a loop-closure detection algorithm [2] using a maximum *a posteriori* scheme that

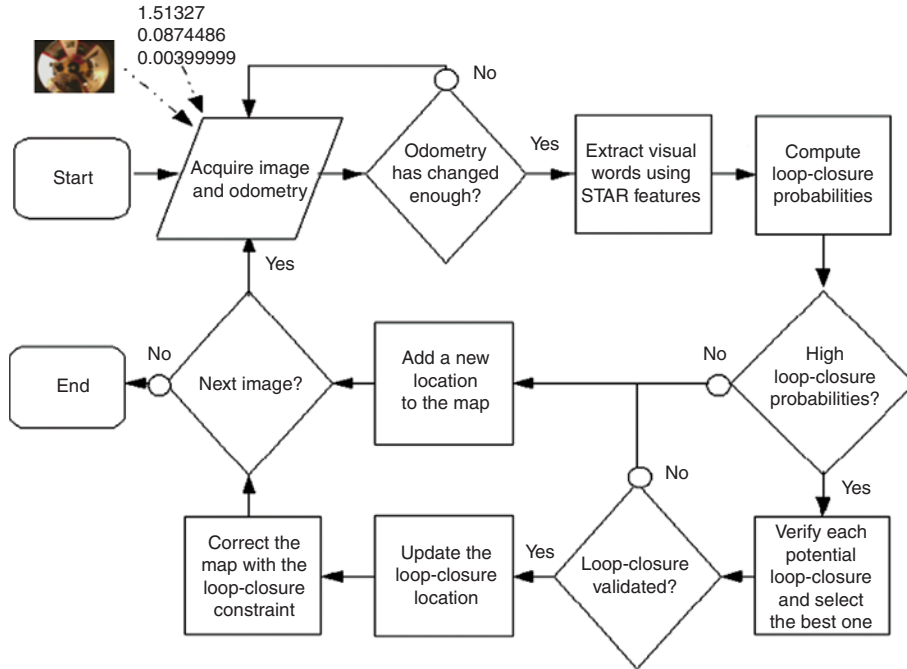


Figure 3. Processing Diagram of the Topometric Map-Building Algorithm.

detects when the robot comes back to a previously visited location. Such detection allows consistent mapping during environment exploration and robust localization during autonomous navigation.

4.1 Odometry-Enhanced Visual Loop-Closure Detection

For each node N_i in the current map, we compute the probability of loop closure with the current image I_t :

$$p(S_t = i | u^t, I^t, M), \quad (1)$$

where $S_t = i$ is event “image I_t comes from node N_i ”; I is the set of images gathered by the robot; u^t is the set of odometry measurements from all the times when images were acquired; and $M = N_0, \dots, N_n$ is the map. Bayes rule, marginalization, and Markov assumption lead to the incremental computation of the *a posteriori* probability as follows (we omit map dependency in all terms for simplicity):

$$p(S_t | u^t, I^t) = \underbrace{\eta \cdot p(I_t | S_t)}_{\text{likelihood model}} \cdot \sum_{j=0}^n \underbrace{p(S_t | u_t, S_{t-1} = j)}_{\text{evolution model}} \underbrace{p(S_{t-1} = j | u^{t-1}, I^{t-1})}_{\text{a priori probability}}. \quad (2)$$

prediction= $p(S_t | u^t, I^{t-1})$

4.1.1 Likelihood Computation

The key point for fast loop-closure detection is the ability to rapidly compute a precise likelihood of the current image for all the nodes of the map. For this, we use a bag of visual words approach [40] that represents images as a set of canonical visual features taken from a codebook. During mapping, we construct an inverted index that links each visual word to the nodes in which it has been seen. Using this inverted index, the current image can be quickly compared with all the map nodes through a voting method whose result is used to estimate the likelihood (Figure 4; see Ref. [2] for details).




Figure 4. Left: Fast Likelihood Computation By Using Bags of Visual Words and an Inverted Index. Right: Computation of the *A Posteriori* Probability Based on the Likelihood and the *A Priori* Probability.

In this article, we optimize the speed of this process by using a static vocabulary tree [34] and the STAR [1] feature detector coupled with the SURF descriptors [4]. A vocabulary containing 10,000 words was created on a set of randomly sampled images in indoor environments for the experiments reported in this article. While even faster feature detectors and descriptors could be used (such as FAST + BRIEF [22]), the proposed approach was accurate enough and already enables to reach a computation time compatible with real-time constraints (see Section 5.1). A proper comparison of the performances of the feature detector and descriptors on our framework will be addressed in our future work.

4.1.2 Evolution Model

By integrating odometry in the evolution model, we predict precisely the evolution of the *a priori* probability and, therefore, we enhance the accuracy of loop-closure detection compared with that in Ref. [2] (see Figure 1 for an example). Starting from a given node, we distribute the probability to each neighboring location in the map depending on the deviation of these nodes' relative positions with the robot displacement since the last update d_u, θ_u, ϕ_u measured by odometry (Figure 5). We used the standard motion model for robot odometry [42], assuming Gaussian noise on the robot displacement measured in polar coordinates:

$$p(S_i | u, S_j) = G_{\mu_d, \sigma_d}(d_{ij} - d_u) \times G_{\mu_\theta, \sigma_\theta}(\theta_{ij} - \theta_u) \times G_{\mu_\phi, \sigma_\phi}(\phi_{ij} - \phi_u), \quad (3)$$

where $u_t = d_u, \theta_u$ gives the odometry displacement in polar coordinates in the frame of the previous robot position and ϕ_u is the variation of robot direction during movement. $G_{\mu, \sigma}(X)$ is the Gaussian distribution of mean μ and variance σ^2 . $d_{ij}, \theta_{ij}, \phi_{ij}$ is the relative position between nodes i and j .

4.1.3 *N*th-Order Evolution Model

To improve the accuracy of the maximum *a posteriori* scheme, we use an *N*th-order Bayesian filter. To do so, we take into account *N* previously computed probabilities, and the evolution model is applied on the corresponding displacements. The *N* predictions are combined to give the final prediction according to

$$p(S_t | u^t, I^{t-1}) = \max_{l=1..N} p(S_t | u_{t-l+1}, S_{t-l}) p(S_{t-l} | u^{t-l}, I^{t-l}), \quad (4)$$

where u_{t-l} is the odometry displacement since time $t-l$. In this article, we used $n = 4$. This modification makes it possible to improve probability prediction when the movement of the robot is different on a second pass in the same area. It also makes it possible to use different image sampling rates during mapping and later reuse of the map.




Figure 5. The Evolution Model Including Odometry and Relative Node Positions.

4.1.4 Validation of the Loop-Closure Candidates

To get an accurate position detection, we do not just verify and validate the more probable loop-closure position that sometimes do not correspond to the closest node in terms of position as it depends on features similarity and on our motion. To improve this, the Bayesian filter presented above is used to extract a small subset of potential loop-closure locations whose probability is above a threshold. Then, to find the previous position, the closest to the current one, we verify all these locations with a 2D image motion computation based on the STAR keypoints, and we select the loop closure that shows the smallest translation, rotation, and scale variation under a threshold. To discard outliers in STAR keypoints matching, the 2D motion (translation, rotation, and zoom in image plane) is computed using RANSAC, accepting the result only if the number of matching points is above a threshold (30 in the experiments reported in this article). As shown in Figure 6, this validation is restrictive to guarantee that the current robot position is very close to the previous one.

4.2 Topometric Mapping

The topometric map is constituted of a set of nodes associated with an image and linked by edges. Each node is associated with an absolute pose in the map (x, y, θ) , where x and y are the 2D position coordinates and θ



Figure 6. Examples of Loop-Closure Acceptance.

First column: match accepted after epipolar geometry validation at a door crossing (Gostai sequence) with >1 m difference between robot positions. Second and third columns: matches accepted using the new validation proposed in this article with 5 cm error.

is an angle representing the direction of the robot when the image was taken. The edges are associated with a relative position between two nodes defined by (d, α, ϕ) , where d and α are the polar coordinates of the second node in the coordinate frame of the first and ϕ is the difference angle between the two nodes' direction.

As shown on Figure 10A, the geometric consistency of the map deteriorates over time owing to the odometry drift. To keep the map consistent, we use a relaxation algorithm each time a loop closure is detected to estimate the nodes' absolute position that best satisfy loop closure and odometry constraints, and correct the odometry drift. In this article, we used the open-source implementation of the TORO algorithm proposed in Ref. [24]. Again, more efficient approaches could be used to improve scalability [27, 31]; however, we preferred TORO, which gives good results with real-time computation for the typical maps of indoor environments.

4.3 Robot Guidance in Topometric Maps

Using the method presented above, it is possible to create a topometric map of an environment while teleoperating the robot. We now present our method to enable autonomous navigation using this topometric map and the same loop-closure framework. The objective is to be able to guide the robot to the position corresponding to any node of the map, while being robust to temporary loop-closure detection failure, noise on odometry, or changes in the scene.

4.3.1 Qualitative Localization

The localization uses the same loop-closure detection method as mapping; however, the incremental part of the system that adds new locations in the graph and the module that relaxes the topometric map are disabled. Robot localization is continuously computed as the position of the robot at the last loop-closure location, to which we add the relative odometry recorded during robot motion since this point in time. This coupling makes it possible to estimate the robot position even in the absence of loop-closure detection during some time, and still corrects the long-term odometry drift each time a loop closure is detected.

As it is difficult to estimate precisely the real robot position given matching images when loop closures are detected, we consider a qualitative localization to navigate in the map. These difficulties are due to the unknown scale factor when computing transformation between two images and to the similarity of image transformation under small camera translation and rotation.

We therefore assume that the image motion is due to a pure rotation of the robot (neglecting translation). This is not true, in general; however, as will be shown in the next section, this leads to corrections in the path-following strategies that guide the robots back to the correct path whatever the real position of the robot is. When a loop closure is detected, the x -axis translation in pixels extracted during the validation of the loop closure is used to estimate the angle between the current direction and the direction recorded during the mapping phase (the conversion factor is deduced from a very simple camera calibration as only the horizontal field of view is required). An example is shown Figure 7.

4.3.2 Path-Following Navigation

We use the Dijkstra algorithm [14] to compute the shortest global path between the node that is the closest to the current robot position and the destination node. This planning takes into account the robot direction in each node to ensure that the robot will travel in the same direction as the mapping phase, thus making it possible to detect loop closure and correct the estimated position.

During the robot motion, a local path is computed joining the current qualitative robot position to the global path. A PID controller applied to the robot direction is then used to compute the velocity of each wheel to follow at best the local path. As stated earlier, the closed-loop operation of this path-following strategy will enforce loop-closure detection that will correct the qualitative position estimate. This strategy is also robust

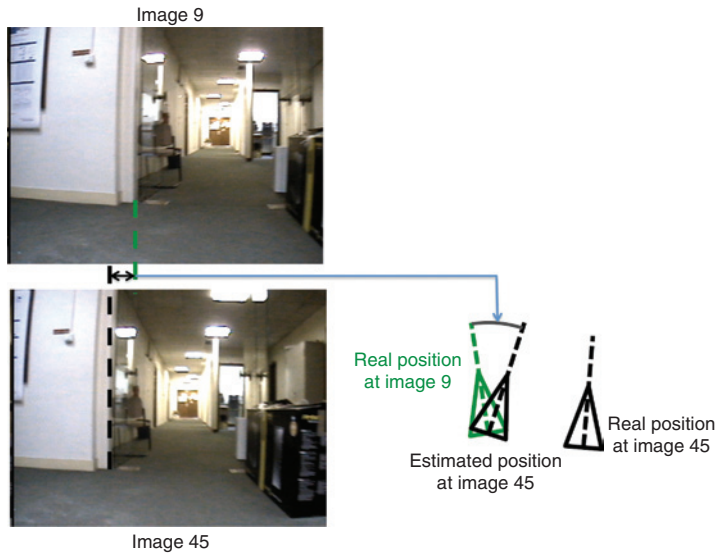


Figure 7. Illustration of the Qualitative Visual Localization.

Left: A loop-closure is detected between two images with 33 pixels of x -axis translation. Given a camera field of view of 45° and 320×240 images, the computed robot rotation is 4.727° . Right: Estimated qualitative robot position and true positions.

to the error made by this qualitative estimate as the localization error always leads to a local path that will guide the robot in the correct direction to reach the global path (Figure 8).

4.3.3 Remote Processing Implementation

An interesting characteristic of this algorithm is that path following can be performed even in the temporary absence of loop-closure detection. Therefore, image processing for loop-closure detection does not need to be performed at a high frame rate, nor in real-time. We took advantage of these facts to implement a remote image processing strategy that makes the approach suitable for robots with very low computational resources.




Figure 8. Path-Following Navigation.

Left: The robot follows the local path but deviates from the true trajectory because of the odometry drift. Middle: The visual loop-closure detection framework gives a qualitative localization of the robot in the graph, taking into account the deviation in direction. As a consequence, in the real world, the local path is modified and the robot corrects its trajectory to stay on the desired path. Right: The robot follows the local path and regains the true trajectory.

In this implementation, only the local path planning, qualitative localization, and PID controller are implemented in real time on board the robot. Low-resolution images (QVGA) are transmitted using wireless connection to a remote computer to detect loop closure. When a loop closure is detected, the loop-closure position and its time stamp are transmitted back to the robot that will use this as the new reference for qualitative localization. Using this strategy, network delays in image transmission and loop-closure detection processing time are not an issue as corrections are retroactively applied.

The system has been implemented using *Urbi*, an open-source software platform to control robots. It includes a C++ component library called *UObject* to interface motors, sensors, and algorithms. We also use the *urbiscript* language to connect the components together using embedded parallel and event-driven semantics. Figure 9 depicts the repartition of the main software modules used for our robotic path-following system. Experiments using this path-following module are shown in the second part of the next section; in the first part, we will present some mapping results.

5 Experimental Results

5.1 Mapping

Experiments have been performed to validate the loop-closure detection framework in indoor environments using a Pioneer P3DX mobile robot equipped with a Canon VC-C50i camera of 45° of horizontal field of view. The robot was guided to perform loops in indoor environments showing strong perceptual aliasing conditions. The images and the odometry information were taken each time the robot moved at least 25 cm or turned of at least 10°. This sampling rate makes it possible to describe the environment without saving too much redundant information; it corresponds to an average acquisition time of one image every 0.7 s, with the average speed of the robot around 0.4 m/s. The computer used for experimentation was based on an Intel Xeon 3G Hz, and we used a small image size: 320 × 240 pixels (QVGA).

Table 1 shows the performances obtained with the approach proposed in this article (new LCD) over the one presented in Ref. [2] for loop-closure detection coupled with map relaxation. We can see that the detection rate is well improved and the inaccurate loop closure is completely removed.

Figure 10 shows some loop-closure details to highlight the improvements in the proposed method. In particular, we can see the increase in the number of loop-closure detection and that the temporal consistency of the detections is improved as there is no more gap between validated loop-closure images.

To illustrate the system performance and show the robustness of the method, we used in parallel the laser SLAM system Karto (<http://www.kartorobotics.com>). Figure 11 shows a map created in an environment

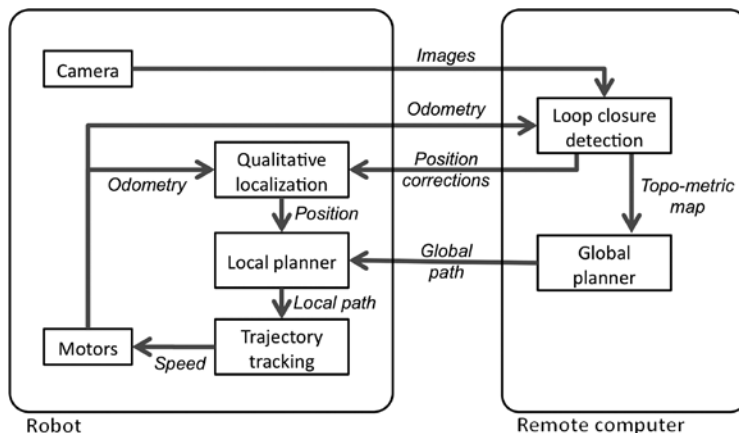


Figure 9. Simplified Diagram of *Urbi* Components Used by the Path-Following System.

Table 1. Compared Results of the Loop-Closure Detection System Presented in Ref. [2] and the Approach Presented in This Article on Four Sequences.

	Science Museum	Gostai Offices	UEI Lab	Six Rooms
Images	112	169	350	1325
Corresponding figure	1	10	11	12
Distance (m)	38	82	98	255
Loop-closure ground truth	9	25	9	≈400
LCD in Ref. [2]				
True positive	4	7	3	–
False positive	1	2	1	–
CPU time/image	370 ms	410 ms	500 ms	–
New LCD				
True positive	7	18	7	352
False positive	0	0	0	0
CPU time/image	13 ms	15 ms	16 ms	22 ms

with a long corridor where the final map is consistent but not perfectly aligned with the ground truth because of the small number of real loop closures. Note that this error is not problematic at all as the map topology is consistent and only relative information between nodes will be used for guidance. As a consequence, the small orientation error on the corridor will not affect navigation. Finally, Figure 12 illustrates a map created in a larger indoor environment with a central hallway and six rooms. This map contains 973 nodes and shows the applicability of our algorithm in an environment of the typical size of standard houses.

5.2 Guidance

Given a map constructed while manually guiding the robot, we have performed autonomous navigation experiments. In all the figures, the reference system gives the real trajectory in a laser map during the learning (green lines) and during the path-following run (magenta lines). The odometry recorded during the learning is shown in yellow. The blue line is the odometry recorded during the path-following run.

Figure 13 shows loop-closure detection and estimation made by the qualitative localization procedure while following a path in an indoor environment. The green circles correspond to the topometric map recorded during the learning phase. The magenta circles correspond to the loop-closure locations detected during the path-following phase. The magenta line in the circles is proportional to the horizontal translation estimated between the matching images that is used by the qualitative localization to estimate robot's direction. This example illustrates the shortest path computation by avoiding the large loop during replay. Moreover, the images were acquired each 5 cm during mapping and each 25 cm during replay. Thanks to the use of the odometry model and the N th-order Markov model, this does not prevent our algorithm to correctly detect the loop closures, and successfully guides the robot to the goal position.

Figure 14 shows another experiment in which the system corrects the odometry drift using the loop closure and guides the robot to the prescribed goal. The image processing rate was roughly 2 Hz, and the average speed of the robot in autonomous navigation was around 0.4 m/s. The loop-closure detection rate during path following was on average 60%, and, as shown, the path following works well. The missed loop closure occurs mainly because while turning or deviating from the trajectory, some images are too close to a wall or an obstacle, i.e., presenting poor features or features different from the learning-stage ones. This makes the loop-closure probability drop below the threshold or prevent the validation stage from accepting the potential loop closure. Another reason for the missed loop closure during path following is that occasionally some images are blurry because of the vibrations of the platforms or fast turns.

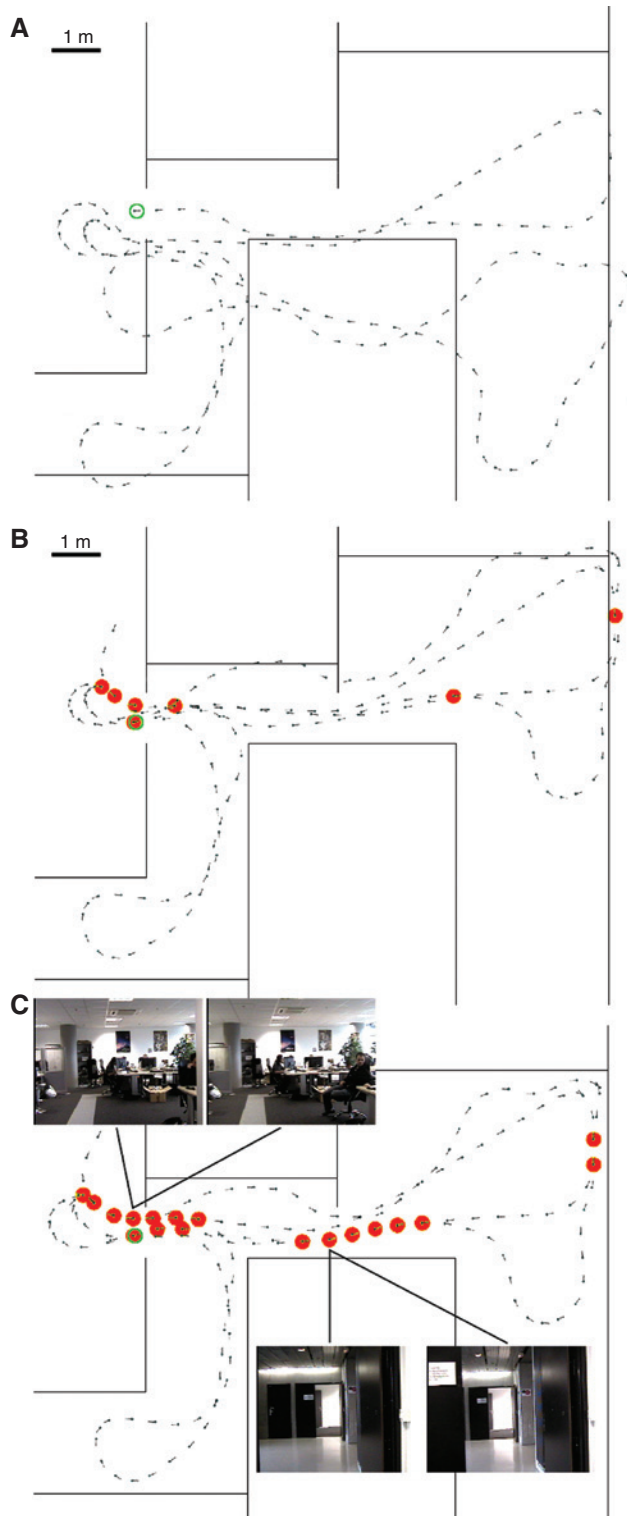


Figure 10. Gostai Sequence: Example of Topometric Mapping Results.

(A) Raw odometry; (B) topometric map using the loop-closure framework from Ref. [2] (7 loop closures); (C) map using the proposed model (18 loop closures).




Figure 11. Correction of the Odometry Drift with Loop-Closure Detections.

The trajectories have been displayed on the 2D laser SLAM map for reference (“UEI Lab” sequence, 350 images, 7 loop-closures). While the global orientation still shows some drift, the local map structure is coherent and the position of the robot at the end of the trajectory is correctly recognized as the starting position.



Figure 12. A Large Topometric Map in a Six-Room Environment Showing the Consistency of the Produced Map (“6 Rooms” Sequence, 1325 Images, 352 Loop Closures).

To better demonstrate the robustness to vision failure or missed detection, on Figure 15 the camera was occluded for short times and people were crossing the robot path in front of the robot. This experiment shows that the system is robust to vision failures up to 20 images, i.e., approximately 5 m. For this experiment with dynamic obstacles and camera occlusions, the average loop-closure detection rate fell down below 40% and the robot was still able to follow the trajectory and reach its goal. It has to be noted that the main odometry drift appears when the robot is turning fast, and that camera occlusions during these turns are much more difficult to compensate. This robustness is, however, sufficient to accommodate reasonable environmental

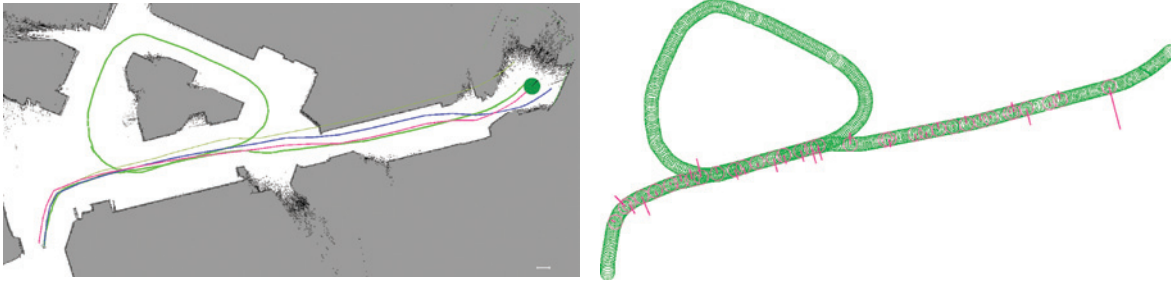


Figure 13. Example of Vision-Based Path Following.

Green circles represent learned places, and magenta circles show the loop-closure locations detected during the navigation. The line in the middle of the magenta circles corresponds to the x-axis translation computed from the matching images.




Figure 14. Large Vision-Based Path Following (≈ 100 m).

changes such as people crossing in front of the camera or large furniture displacement that could prevent the robot from detecting loop closure for a few seconds. It has to be noted that environmental changes should even affect less the loop-closure detections, as it should remove only few features whereas the camera occlusion drastically remove any features.

The Pioneer 3DX platform has a rather efficient odometry compared with other wheeled robots or humanoid robots. To quantify the robustness to odometry drift, we performed navigation experiments while simulating several levels of odometry noise. This was achieved by taking the displacement computed by the laser SLAM and corrupting it every 50 ms with a Gaussian noise with the same mean and variance:

$$d = d + G(\sigma, \sigma); \quad \theta = \theta + G(\sigma, \sigma); \quad \phi = \phi + G(\sigma, \sigma). \quad (5)$$

The variance of these noises were taken as a multiple of a base value of $\sigma = 5 \cdot 10^{-4}$. This noisy displacement value was used as the odometry by our path-following system. In all experiments, the same topometric map was used, created using the laser SLAM displacements. To obtain representative results, we executed the path following five times for each noise level and plotted the average result over these five cases. Figure 16 shows two of these experiments, with the Pioneer robot odometry and with a simulated noise using 6σ variance.




Figure 15. Example of Vision-Based Path Following with Camera Occlusions or People Crossing Just in Front of the Robot.

Figure 17 shows the number of loop closures detected during these experiments. While the number of loop closure decreases with noise, this number remains high enough to ensure position correction for path following. This figure also shows the error in position between the trajectory executed during mapping and the trajectory executed during replay, measured with the laser SLAM position. This position error clearly increases with the odometry noise but remains small enough to ensure path-following success in all trails up to 6σ noise. Figure 18A shows that at 8σ , three of five trials lead to a failure, the robot hitting a corridor wall during path following. In this same figure, we can see that using only the raw odometry, four of five trials lead to failures with the first noise level of 1σ and that all trials fail after the 2σ level. In comparison, our qualitative position estimate makes it possible successfully follow the path in all situations up to 6σ noise. This noise applied is rather important, which demonstrates that our algorithm can be working on many wheeled robot providing odometry. The robustness to odometry noise also lets us predict that the same framework could also be applied on different platforms like legged robots, for example, known to provide poor-quality odometry. This possibility will be tested in our future work.

Finally, Figure 18B shows the mean correction applied by our algorithm during path following. It can be noted that corrections increase with the odometry noise, but that these corrections do not converge to zero with a noise-free odometry. This is caused by the sampling policy of images: as images are not taken at the same position during mapping and path following, each loop closure will produce a direction correction, in particular during rotation, even if the robot estimated position is perfect. However, it can be seen in the

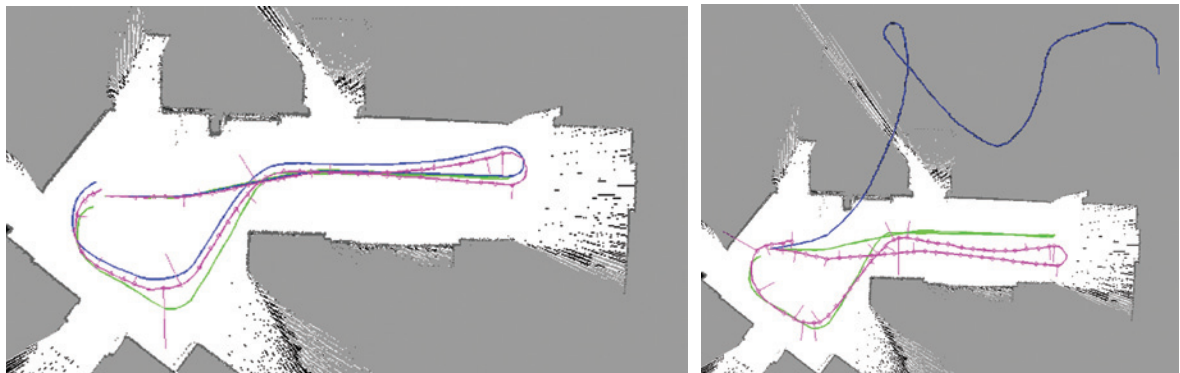


Figure 16. Two Examples of Vision-Based Path Following Using the Same Map (in Green). Raw odometry is in blue; replayed trajectory is in magenta. Left: Path following using the pioneer 3DX odometry. Right: Path following using simulated noise of 6σ .




Figure 17. Number of Loop Closures Detected (Blue) and Sum of Squared Differences between Reference Path and Path Following (Green) for Several Levels of Simulated Odometry Noise.

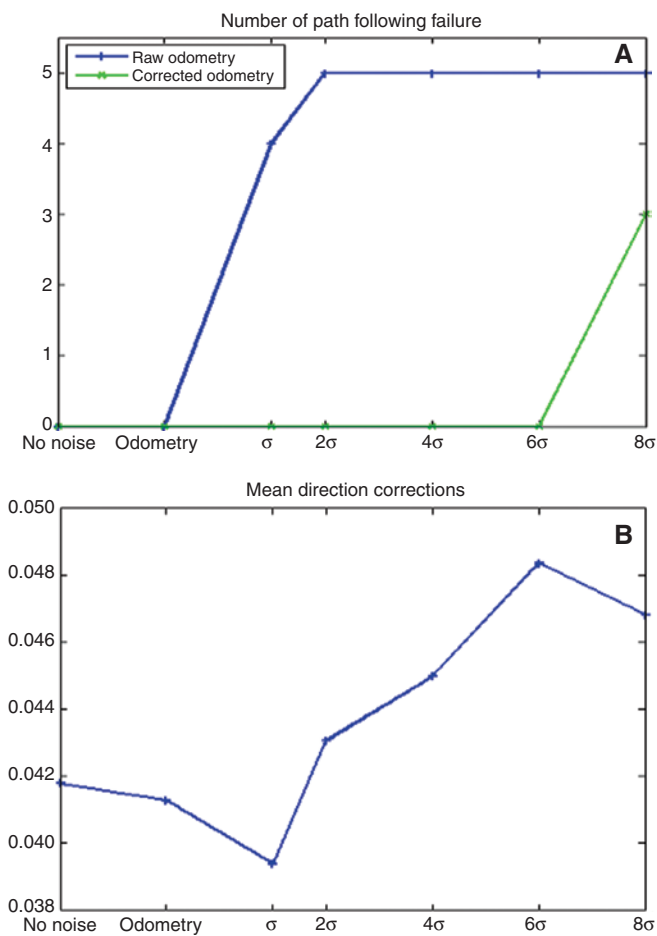


Figure 18. (A) Number of Failed Path Following over Five Trials Using Raw Odometry (Blue) and the Qualitative Position Estimated by Our Algorithm (Green) for Several Levels of Simulated Odometry Noise. (B) Mean Direction Correction Applied during Path Following, in Radians.

path-following error (Figure 17) that these corrections do not prevent efficient path following in the case of a precise odometry, and make path following possible in the case of a noisy odometry.

6 Discussion

The algorithm proposed in this article is well adapted for mapping in indoor environments as the path is more constrained. Indeed, in indoor environments, doors are a required route to go from one room to another and therefore forces the robot to close loops, even with our restricted definition. Outdoor environments have also been tested successfully but usually provide less loop-closure detection as the paths are often less constrained. The qualitative guidance approach is, however, as efficient outdoor as it is indoor as the robot is guided along already executed routes, and is able to close loop frequently and effectively limit the odometry drift.

Our path-following system uses only relative information from the odometry and between current images and past images. This is linked to the fact that the qualitative localization itself is relative to the position of loop-closure detections. As a consequence, an error in the global metrical position of a node such as in Figure 11 where the map diverges from the true corridor direction will not prevent the robot to correctly follow a path inside this corridor. Another consequence is that the use of a relaxation algorithm to obtain a globally consistent map is not required by the path-following system as long as we only follow trajectories really executed during map building. In our application, this relaxation algorithm is, however, important in the prediction step of the Bayesian filter and when following trajectories computed by the Dijkstra algorithm. Indeed, a coherent map is required to correctly predict position probability with the odometry probabilistic model and the global metrical path followed using the qualitative localization must not contain discontinuities.

As shown in Figure 13, another interest of using an odometry model and a qualitative metrical localization is the possibility to use different image sampling policies between mapping and autonomous navigation. In conjunction with the N th-order Markov model, this makes it possible to create high-density maps, sampling images every few centimeters, during supervised navigation, while using a lower sampling frequency to reduce computational burden during autonomous navigation.

Finally, the current model only allows to directly reproduce paths, therefore assuming the absence of dynamic obstacles. In dynamic environments, it would be possible to couple our method with a local obstacle avoidance strategy based on sonars, for example. The qualitative localization based on odometry would make it possible for the robot to come back to the global path after avoiding the obstacle, which would not be possible with a direct visual servoing approach. Adding such an obstacle avoidance strategy could even improve the number of loop-closure detection by forcing the robot to always use similar trajectories (centering the robot along corridors for example), thus improving the precision of the qualitative localization.

7 Conclusion

We have presented a system for topometric mapping and navigation based on an appearance-based loop-closure detection framework combining vision and odometry. In particular, we have proposed an optimized loop-closure detection algorithm that shows high responsiveness, consistent and precise position detections while working with several kinds of uncalibrated cameras, requiring limited computational resources and allowing remote processing. The loop-closure detection framework makes it possible to build incrementally consistent topometric maps in real time while remotely guiding the robot without assumptions about the environment. We also proposed a qualitative localization method based on loop-closure detections and odometry that makes it possible to guide a robot along routes computed from the topometric map. This navigation method is robust to temporary vision failures (people crossing, objects moved), noise on the odometry, and delays in vision processing, thus making it applicable to low-cost platforms.

In future work, we plan to develop a global localization strategy to be able to relocalize the robot inside the map after it has been turned off. This strategy could be based on a standard behavior such as wall or corridor following to provide a first loop-closure detection and initialize localization. We also plan to study the use of a pan-tilt camera to make it possible to detect loop closure with trajectories followed in reverse direction, as in the current framework, the robot can only follow trajectories in the same direction as during mapping. Finally, we plan to add obstacle avoidance capabilities using sonar data to be robust to dynamic obstacles and more important changes in the environment.

Bibliography

- [1] M. Agrawal, K. Konolige and M. R. Blas, CenSurE: Center Surround Extremas for real-time feature detection and matching, in: *European Conference on Computer Vision*, pp. 102–115, Springer, Marseille, 2008.
- [2] A. Angeli, D. Filliat, S. Doncieux and J.-A. Meyer, A fast and incremental method for loop-closure detection using bags of visual words, *IEEE T. Robot., Special Issue on Visual SLAM* **24** (2008), 1027–1037.
- [3] A. Angeli, D. Filliat, S. Doncieux and J.-A. Meyer, Real-time visual loop-closure detection, in: *IEEE International Conference on Robotics and Automation*, pp. 1842–1847, IEEE, Pasadena, 2008.
- [4] H. Bay, T. Tuytelaars and L. V. Gool, SURF: Speeded Up Robust Features, in: *9th European Conference on Computer Vision*, pp. 346–359, Elsevier, 2006.
- [5] S. Bazeille, E. Battesti and D. Filliat, Qualitative localization using vision and odometry for path following in topo-metric maps, in: *European Conference on Mobile Robotics*, pp. 303–308, 2011.
- [6] S. Bazeille and D. Filliat, Incremental topo-metric slam using vision and robot odometry, in: *Proceedings of the International Conference on Robotics and Automation*, IEEE, Shanghai, 2011.
- [7] J. L. Blanco, J. Gonzalez and J.-A. Fernandez-Madriral, Subjective local maps for hybrid metric-topological slam, *Robot. Autonom. Syst.* **57** (2009), 64–74.
- [8] O. Booij, B. Terwijn, Z. Zivkovic and B. Kröse, Navigation using an appearance based topological map, in: *IEEE International Conference on Robotics and Automation*, IEEE, Roma, 2007.
- [9] D. Burschka and G. Hager, Vision based control of mobile robots, in: *Proceedings of the International Conference on Robotics and Automation*, IEEE, Seoul, 2001.
- [10] Z. Chen and T. Birchfield, Qualitative vision-based path following, *IEEE T. Robot.*, (2009), 749–754.
- [11] J. Correa and A. Soto, Active visual perception for mobile robot localization, *J. Intell. Robot. Syst.* **58** (2010), 339–354.
- [12] M. Cummins and P. Newman, Fab-map: probabilistic localization and mapping in the space of appearance, *Int. J. Robot. Res.* **27** (2008), 647–665.
- [13] A. J. Davison, I. D. Reid, N. D. Molton and O. Stasse, Monoslam: real-time single camera slam, *IEEE T. Pattern Anal. Mach. Intell.* **29** (2007), 1052–1067.
- [14] E. W. Dijkstra, A note on two problems in connection with graphs, *Numer. Mathem.* **1** (1959), 269–271.
- [15] A. Diosi, A. Remazeilles, S. Segvic and F. Chaumette, Outdoor visual path following experiments, in: *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 4265–4270, IEEE, San Diego, 2007.
- [16] T. Duckett, S. Marsland and J. Shapiro, Fast, on-line learning of globally consistent maps, *Autonom. Robot.* **12** (2002), 287–300.
- [17] E. Eade and T. Drummond, Monocular SLAM as a graph of coalesced observations, in: *International Conference on Computer Vision*, pp. 1–8, IEEE, Rio, 2007.
- [18] D. Filliat and J. A. Meyer, Global localization and topological map learning for robot navigation, in: *Proceedings of The Seventh International Conference on Simulation of Adaptive Behavior (SAB02) on From Animals to Animats*, pp. 131–140, MIT Press, Cambridge, MA, USA, 2002.
- [19] F. Fraundorfer, C. Engels and D. Nistér, Topological mapping, localization and navigation using image collections, in: *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 3872–3877, IEEE, San Diego, 2007.
- [20] U. Frese, P. Larsson and T. Duckett, A multilevel relaxation algorithm for simultaneous localization and mapping, *IEEE T. Robot. Autom.* **21** (2005), 196–207.
- [21] S. Fu and G. Yang, Uncalibrated monocular based simultaneous localization and mapping for indoor autonomous mobile robot navigation, in: *Networking, Sensing and Control*, pp. 663–668, IEEE, Okayama, 2009.
- [22] D. Galvez-Lopez and J. D. Tardos, Real-time loop detection with bags of binary words, in: *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 51–58, IEEE, San Francisco, 2011.
- [23] T. Goedemé, M. Nuttin, T. Tuytelaars and L. Van Gool, Omnidirectional vision based topological navigation, *Int. J. Comput. Vision* **74** (2007), 219–236.
- [24] G. Grisetti, C. Stachniss, S. Grzonka and W. Burgard, A tree parameterization for efficiently computing maximum likelihood maps using gradient descent, in: *Proceedings of Robotics: Science and Systems*, pp. 65–72, Atlanta, GA, June 2007.

- [25] A. Howard, M. J. Mataric and G. Sukhatme, Relaxation on a mesh: a formalism for generalized localization, in: *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 1055–1060, IEEE, Maui, 2001.
- [26] S. A. Hutchinson, G. D. Hager and P. I. Corke, A tutorial on visual servo control, *IEEE T. Robot. Autom.* **12** (1996), 654–670.
- [27] M. Kaess, H. Johannsson, R. Roberts, V. Ila, J. Leonard and F. Dellaert, ISAM2: Incremental smoothing and mapping with fluid relinearization and incremental variable reordering, in: *IEEE International Conference on Robotics and Automation*, pp. 3281–3288, IEEE, Shanghai, 2011.
- [28] A. Kawewong, N. Tongprasit and O. Hasegawa, Pirf-nav 2.0: fast and online incremental appearance-based loop-closure detection in an indoor environment, *Robot. Autonom. Syst.* **59** (2011), 727–739.
- [29] K. Konolige and M. Agrawal, Frameslam: from bundle adjustment to real-time visual mapping, *IEEE T. Robot.* **24** (2008), 1066–1077.
- [30] J. Kosecká, F. Li and X. Yang, Global localization and relative positioning based on scale-invariant keypoints, *Robot. Autonom. Syst.* **52** (2005), 209–228.
- [31] R. Kummerle, G. Grisetti, H. Strasdat, K. Konolige and W. Burgard, G2o: a general framework for graph optimization, in: *IEEE International Conference on Robotics and Automation*, pp. 3607–3613, IEEE, Shanghai, 2011.
- [32] M. Milford and G. Wyeth, Persistent navigation and mapping using a biologically inspired SLAM system, *Int. J. Robot. Res.* **29** (2009), 1131–1153.
- [33] D. Nistér, O. Naroditsky and J. Bergen, Visual odometry for ground vehicle applications, *J. Field Robot.* **23** (2006), 3–20.
- [34] D. Nister and H. Stewenius, Scalable recognition with a vocabulary tree, in: *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp. 2161–2168, IEEE, Washington, 2006.
- [35] E. Olson, J. Leonard and S. Teller, Fast iterative alignment of pose graphs with poor initial estimates, in: *Proceedings of the IEEE International Conference on Robotics and Automation*, pp. 2262–2269, IEEE, Orlando, 2006.
- [36] P. Rybski, F. Zacharias, J. Lett, O. Masoud, M. Gini and N. Papanikolopoulos, Using visual features to build topological maps of indoor environments, in: *IEEE International Conference on Robotics and Automation*, IEE, Taipei, 2003.
- [37] M. Saedan, L. C. Wang and M. H. Ang, Appearance-based SLAM with map loop closing using an omnidirectional camera, in: *Proceedings of the IEEE/ASME International Conference on Advanced Intelligent Mechatronics*, pp. 1–6, IEEE, Singapore, 2007.
- [38] C. Sagues and J. Guerrero, Visual correction for mobile robot homing, *Robot. Autonom. Syst.* **50** (2005), 41–49.
- [39] G. Sibley, C. Mei, I. Reid and P. Newman, Adaptive relative bundle adjustment, in: *Robotics Science and Systems*, June 2009.
- [40] J. Sivic and A. Zisserman, Video Google: a text retrieval approach to object matching in videos, in: *IEEE International Conference on Computer Vision*, pp. 1470–1477, IEEE, Nice, 2003.
- [41] N. Sünderhauf and P. Protzel, Towards a robust back-end for pose graph slam, in: *Proceedings of the IEEE International Conference on Robotics and Automation*, pp. 1254–1261, IEEE, Saint Paul, 2012.
- [42] S. Thrun, W. Burgard and D. Fox, *Probabilistic robotics*, The MIT Press, Cambridge, MA, 2005.
- [43] C. Valgren, T. Duckett and A. Lilienthal, Incremental spectral clustering and its application to topological mapping, in: *IEEE International Conference on Robotics and Automation*, IEEE, Roma, 2007.
- [44] Z. Zivkovic, B. Bakker and B. Kröse, Hierarchical map building using visual landmarks and geometric constraints, in: *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 2480–2485, IEEE, Edmonton, 2005.