Copyright Notice

This is the authors version of a work that was accepted for publication in Pattern Recognition Letters.

Cristoforis, Nitsche, Krajnik, Pire, Mejail: *Hybrid vision*based navigation for mobile robots in mixed indoor/outdoor environments, Pattern Recognition Letters, 2014.

Changes resulting from the publishing process, such as peer review, editing, corrections, structural formatting, and other quality control mechanisms may not be reflected in this document. Changes may have been made to this work since it was submitted for publication. The full version of the article is available on Elsevier's sites with DOI: 10.1016/j.patrec.2014.10.010.



Pattern Recognition Letters journal homepage: www.elsevier.com

Hybrid vision-based navigation for mobile robots in mixed indoor/outdoor environments

Pablo De Cristóforis^{a,**}, Matias Nitsche^a, Tomáš Krajník^{b,c}, Taihú Pire^a, Marta Mejail^a

^aLaboratory of Robotics and Embedded Systems, Faculty of Exact and Natural Sciences, University of Buenos Aires, Argentina ^bLincoln Centre for Autonomous Systems, School of Computer Science, University of Lincoln, United Kingdom

^cDepartment of Cybernetics, Faculty of Electrical Engineering, Czech Technical University in Prague, Czech republic

ARTICLE INFO

ABSTRACT

In this paper we present a vision-based navigation system for mobile robots Article history: equipped with a single, off-the-shelf camera in mixed indoor/outdoor environments. A hybrid approach is proposed, based on the teach-and-replay technique, which combines a path-following and a feature-based navigation al-2000 MSC: 41A05 gorithm. We describe the navigation algorithms and show that both of them 41A10 correct the robot's lateral displacement from the intended path. After that, we 65D05 claim that even though neither of the methods explicitly estimates the robot po-65D17 sition, the heading corrections themselves keep the robot position error bound. We show that combination of the methods outperforms the pure feature-based Keywords. Vison-based navigation approach in terms of localization precision and that this combination reduces Mixed indoor/outdoor environmets map size and simplifies the learning phase. Experiments in mixed indoor/out-Mobile robotics door environments were carried out with a wheeled and a tracked mobile robots in order to demonstrate the validity and the benefits of the hybrid approach. © 2014 Elsevier Ltd. All rights reserved.

1. Introduction

Autonomous navigation can be roughly described as the process of moving safely along a path between a starting and a final point. In mobile robotics, different sensors have been used for this purpose, which has led to a varied spectrum of solutions. Active sensors such as sonars (Borenstein and Koren, 1988), laser range finders (Surmann et al., 2001) and radars (Kaliyaperumal et al., 2001) have been used in autonomous navigation methods. These sensors are inherently suited for the task of obstacle detection and can be used easily because they directly measure the distances from obstacles to the robot.

Other sensors that are becoming a standard part in mobile robotic systems, particularly in field robotics, are visual sensors. Quality cameras have become increasingly affordable, they are small and can provide high resolution data in real time. They are passive and therefore do not interfere with other sensors. Unlike range sensors, they can not only be used to detect obstacles, but also to identify forbidden areas and navigate mobile robots using human-defined rules (i.e. keep off the grass). Such forbidden areas are not obstacles, since they are in the same plane as the path, but should be considered as nontraversable. Moreover, and probably more important, nowadays computational power required by image processing techniques is readily available in consumer hardware. For these reasons, in recent years vision-based navigation for mobile robots has been a widely studied topic.

Many visual navigation methods generally rely on the extraction of salient visual features, using algorithms such as well-known SIFT (Scale Invariant Feature Transform) or SURF (Speeded Up Robust Features), among others. These are then used as external references from which information about the structure of the surrounding environment and/or the ego-motion of the robot is estimated.

A classical approach to visual navigation is known as teachand-replay. This technique is closely related to visual servoing (Chaumette and Hutchinson, 2006; F. Chaumette and S. Hutchinson, 2007), where the task is to reach a desired pose using vision for control feedback. Several teach-and-replay works employ visual features as landmarks to guide the autonomous

^{**}Corresponding author:

e-mail: pdecris@dc.uba.ar (Pablo De Cristóforis)

navigation. A particularly succesful work has been presented by Furgale and Barfoot (2010). This method uses a stereocamera to build a manifold map of overlapping submaps as the robot is piloted along a route. Each submap represent a metric reconstruction of a portion of the environment. The use of local submaps allows the rover to faithfully repeat long routes without the need for an accurate global environment reconstruction.

In contrast to metric reconstruction of the environment, the authors of Krajník et al. (2010), and Chen and Birchfield (2009) present a simplistic monocular navigation algorithm that performs autonomous navigation without building metric maps, so-called appearance-based or qualitative navigation. In their approach, the robot steering is calculated directly from horizontal displacement of the currently perceived and previously mapped image features. The article of Krajník et al. (2010) contains a mathematical proof of the stability of their feature-based navigation.

In recent years, teach-and-replay methods have been studied for both terrestrial and aerial robots. On the one hand, many works use metric reconstruction techniques, such as stereo-based triangulation proposed by Ostafew et al. (2013) and Pfrunder et al. (2014). On the other hand, several works still employ solely appearance-based navigation (Ok et al., 2012; Augustine et al., 2012; Krajnik et al., 2012; Nitsche et al., 2014).

The main drawback of the aforementioned methods lies in the fact that the robot workspace is limited only to the regions mapped during the training step. The user has to guide the robot all around the entire path before performing autonomous navigation, which may represent a very tedious process, especially in large outdoor environments. We argue that the mathematical proof presented in Krajník et al. (2010) is not limited only to feature-based map-and-replay navigation methods. This fact should allow to design hybrid navigation systems that use both map-and-replay and map-less reactive navigation methods.

An example of such a method is a real-time monocular path following algorithm that allows a mobile robot navigate through semi-structured or unstructured paths De Cristóforis et al. (2013). This method uses a probabilistic approach similar to the proposed by Dahlkamp et al. (2006), learning a visual model of the nearby road using an area in front of the vehicle as a reference. The traversable path appearance is modeled as a mixture of Gaussians and the classification between navigable or not navigable area is done of segments rather than pixel level. From the previously determined navigable area, a path center trajectory is estimated and using a motion control law the robot is guided to the center of the detected path and maintain it away from path edges during navigation.

Though different solutions have been proposed to address the indoor or outdoor vision-based navigation, there are not many works that solve the problem when it comes to mixed indoor and outdoor environments. The main contribution of this work is to propose a hybrid navigation method than can work successfully both in indoor and outdoor environments while overcoming the aforementioned drawbacks of existing methods, which are generally oriented towards a specific scenario. The proposed method combines a teach-and-replay feature-based method (Krajník et al., 2010) and a the segmentation-based approach for following paths (De Cristóforis et al., 2013). Both methods have the same principle behind: to correct the heading of the robot to guide the autonomous navigation.

When transitioning to environment contains semi-structured paths, the visual path-following method is used. When there is no path to follow, a map is built using feature-based approach. By recording odometric information for establishing the length of each segment in the map, a large mixed indoor/outdoor environment can be traversed autonomously. We claim that even though neither of the two methods explicitly estimates the robot position in the environment, the heading corrections themselves keep the robot position error bound. In other words, by combining two navigation methods based on the same principle, i.e. heading correction, a hybrid navigation approach which takes the best of each can be obtained, without loosing accuracy in autonomous navigation. In this paper this is proved in both theoretical and practical point of views.

The remainder of this paper is organized as follows: the hybrid navigation approach is presented in Section 2, description of the feature-based method is introduced in Section 3, description of the segmentation-based method for following path is introduced in Section 4, the convergence analysis is performed in Section 5, evaluation criteria and experimental results are discussed in 6 and the paper is concluded in Section 7.

2. Hybrid visual navigation

The main idea behind the proposed hybrid navigation method is to combine two separate approaches, feature-based and segmentation-based, by alternating between one and the other depending on the part of the environment the robot is at.

Both of the aforementioned methods correct the robot's heading so it is steered towards the intended path. However, neither of the methods can compute the localization of the robot within the environment. Odometry is only used to determine the relative position of the robot within a segment, and therefore, when the robot has to change between the two navigation methods. One might argue that the cumulative nature of the odometric error would eventually cause the navigation method to fail. However, Krajník et al. (2010) indicate that the heading correction itself can compensate for odometric errors and keep the robot position error bound. In this article, we will extend this idea of heading compensation to a broader set of monocular navigation systems based on robot's heading correction. In particular, we will show that stability of bearing-only navigation is not limited to landmark-based localization only, but can be applied to reactive path-following methods as well.

The advantages of the proposed approach compared to the one presented in Krajník et al. (2010) can be summarized as follows:

- 1. It is not necessary to teach the robot the entire trajectory in areas where a detectable path is present. This significantly simplifies and shortens the learning phase.
- 2. Since the path-following method is adaptive, it is more robust to environment appearance variations compared to the

feature-based method. This increases the overall reliability of the navigation.

- 3. The path following method reduces robot positioning error faster than the feature-based approach. Thus, combination of both methods allows for a more precise navigation.
- 4. The path following method does not need to store information about the entire edge that is to be traversed. Rather than that, it just need information about the path texture, which dramatically reduces spatial requirements for storage of the topological map especially when traversing long outdoor paths.

2.1. Method overview

The topological map that the method uses is a graph where nodes represent places in the environment and edges represent the navigation methods that can be used to move between these places. While the nodes are associated with azimuths indicating directions of the outgoing edges, the edges contain information that allows the robot to traverse between the individual places. In other words, an edge of the hybrid topological map represents an action that the robot should take during its length in different parts of the route: i.e. either the path-following or the feature-based navigation.

The environment map is created during a training phase where the user builds a topological map by driving a semiautonomously moving robot through the environment. During the training, the robot either moves forwards while creating an image-feature based map or it follows a pathway of the environment. Whenever the user terminates the current behaviour, the robot records the length of the edge and creates a topological node. Then, the user can manually set the robot orientation and start path following or feature mapping again. It is important to note that if the length of the path-following edge is known (e.g. from an overhead map), there is no need to traverse the edge during the learning phase.



Fig. 1. Schematic of an example map of the hybrid method for a given indoor/outdoor scenario. Lines marked in blue consist of the portions of the map which are traversed using the path-following method, while lines marked in red, portions which are traversed using the feature-based approach.

In the autonomous navigation or "replay" phase, the robot traverses a sequence of edges by either one of the two movement primitives. Note that neither or the methods performs localization of the robot. Both methods simply keep the robot close to the trained paths while using odometric reading to decide whether the particular edge has been traversed or not. In Section 5 we will show that despite of simplicity of the approach, the robot position error does not diverge over time.

In the next sections we will describe each navigation movement primitive: feature-based for open areas (indoor or outdoor) where there is not bounded path to go through, or cross intersection of two or more paths, and image segmentation based method conceived to follow structured or semi-structured paths.

3. Feature-based visual navigation

The feature-based navigation method used as a motion primitive of the proposed approach is based on Krajník et al. (2010). This method uses the map-and-replay technique which is described in this section. This method allows a robot to follow a given path by performing an initial learning phase where features extracted from the camera images are stored. In a second replay phase, the robot autonomously navigates through the learned path while correcting its heading by comparing currently seen and previously stored features.

3.1. Image Features

Since image features are considered as visual landmarks, the feature detector and descriptor algorithms are a critical component of the navigation method. Image features must provide enough information to steer the robot in a correct direction. Furthermore, they should be robust to real world conditions, i.e. changing illumination, scale, viewpoint and partial occlusions and of course its performance must allow real-time operation. In the original version of the method, Krajník et al. used SURF (Speeded Up Robust Feature) algorithm proposed by Bay et al. (2008) to identify visual landmarks in the image. The SURF method is reported to perform better than most SIFT (Lowe, 2004) implementations in terms of speed and robustness to viewpoint and illumination change.

For this work, further evaluation of a variety of extractor and descriptor algorithms, in the context of visual navigation, was performed. The goal of these off-line experiments was to determine which algorithm combination is the best choice for visual navigation in terms of performance, robustness and repeatability (Krajník et al., 2013). As a result of this evaluation, we conclude that the combination of STAR based on CenSurE (Center Surround Extremas) algorithm (Agrawal et al., 2008) to detect image features and BRIEF (Binary Robust Independent Elementary Features) algorithm (Calonder et al., 2010) to build the descriptor of the feature outperforms the other image feature extraction algorithms. The STAR extractor gives better results to detect landmarks than SURF, because keypoints extracted with STAR are more salient and thus, also more stable than keypoints extracted with SURF. The BRIEF descriptor uses a simpler coding scheme to describe a feature, this reduces storage requirements and increases matching speed. The SURF use 64 (or 128 in the extended version) double-precision floating-point values to describe a keypoint, this results in 512

bytes (or 1024) of data for each descriptor. In contrast, BRIEF uses a string of bits that can be packed in 32 bytes. Taking into account the high number of features present in the map, this can amount to hundreds of megabytes which are saved in large-scale maps. Furthermore, in terms of speed, instead of using the Euclidean distance to compare descriptors as it is done with SURF, BRIEF descriptors can be compared using Hamming distance, which can be implemented extremely fast using SSE instructions.

3.2. Learning Phase

In the initial learning phase, the robot is manually guided through the environment while the map is built. The operator can either let the robot go forward or stop it and turn it in a desired direction. A map can therefore be built, which will have the shape of a series of segments of certain length as computed by the robots odometry. During forward movement, the robot tracks features extracted from the robot's camera images. For each feature, its position in the image is saved along its descriptor. Also, initial and final distances (as measured by the robot's odometry) along the segment where the feature was first and last seen are also saved. Tracking is performed by extracting features for the present frame and comparing these to others that were continously seen in previous frames. While features remain in the set of tracked landmarks, their information can be updated. When features are not seen anymore, they are added to the map and removed from this set. Matching is performed by comparing descriptors.

The procedure which creates a map of one segment is as follows in the listing 1. Each landmark has an associated descriptor l_{desc} , pixel position l_{pos_0} , l_{pos_1} and robot relative distance l_{d_0} , l_{d_1} , for when the landmark was first and last seen, respectively.

The map obtained by this approach is thus composed of a series of straight segments, each described by its length *s*, azimuth ϕ_0 and a set of detected landmarks *L*. A landmark $l \in L$ is described by a tuple ($\mathbf{e}, k, \mathbf{u}, \mathbf{v}, f, g$), where \mathbf{e} is its BRIEF descriptor and *k* indicates the number of images, in which the feature was detected. Vectors \mathbf{u} and \mathbf{v} denote positions of the feature in the captured image at the moment of its first and last detection and *f* and *g* are distances of the robot from the segment start at these moments.

3.3. Autonomous Navigation Phase

In order to navigate the environment using this method, the robot is initially placed near the starting position of a known segment. Then, the robot moves forward at constant speed while correcting its heading. This is performed by retrieving the set of relevant landmarks from the map at each moment and to estimate their position in the current camera image. These landmarks are paired with the currently detected ones and the differences between the estimated and real positions are calculated. As these differences in image coordinates are related to a displacement of the robot in the world, it is possible to actively minimize these by moving the robot accordingly, and thus the previously learnt path is replayed. Algorithm 1: Learning Phase **Input**: φ : initial robot orientation **Output**: L: landmarks learned for the current segment, s: segment length, φ_0 : segment orientation $\varphi_0 \leftarrow \varphi$ read robot orientation from gyro or compass /* clear the learned landmark set */ $L \leftarrow \emptyset$: $T \leftarrow \emptyset;$ /* clear the tracked landmark set */ repeat $F \leftarrow$ features extracted from the current camera image foreach $l \in T$ do $f \leftarrow \operatorname{find}_\operatorname{match}(l,F)$ if no match then $T \leftarrow T - \{l\};$ /* stop tracking */ $L \leftarrow L \cup \{l\};$ /* add to segment */ else $F \leftarrow F - \{f\}$ $l \leftarrow (l_{desc}, l_{pos_0}, f_{pos}, l_{d_0}, d) / *$ update image coordinates & robot position */ foreach $f \in F$ do $T \leftarrow T \cup \{ (f_{desc}, f_{pos_0}, f_{pos_0}, d, d) \};$ /* start tracking new landmark */ until operator terminates learning mode

```
Algorithm 2: Navigation Phase
 Input: L: landmarks of the current segment s: current
          segment length
 Output: \omega: angular robot speed v: forward robot speed
                              /* start moving forwards */
 v \leftarrow v_0;
 while d < s do
      H \leftarrow \emptyset:
                      /* pixel-position differences */
      T \leftarrow \emptyset;
                                    /* tracked landmarks */
      foreach l \in L do
          if l_{d_0} < d < l_{d_1} then
               T \leftarrow T \cup \{l\};
                                                /* get expected
               landmarks according to d * /
      F \leftarrow features extracted from the current camera image
      while T not empty do
          f \leftarrow \text{find}_\text{match}(l,F)
          if matched then
               /* compare feature position to
                    estimated current landmark
                    position by interpolation
                                                                     */
              h \leftarrow f_{pos} - \left( (l_{pos_1} - l_{pos_0}) \frac{d - l_{d_0}}{l_{d_1} - l_{d_0}} + l_{pos_0} \right)H \leftarrow H \cup \{h\}
          T \leftarrow T - \{l\}
     \omega \leftarrow \gamma \operatorname{mode}(H)
 v, \omega \leftarrow 0; /* segment completed, stop moving */
```

More formally, the replay phase is as follows. Each time a new image is captured, the robot updates the estimated traveled distance d from the start of the segment and builds the set T by retrieving landmarks in L that were visible at distance d during learning. Then, features are extracted from the current camera image and put in the set F. A pairing between the sets F and T is established as in the learning phase. After establishing these correspondences, for each pair of matching features their image position is compared. The expected image position of landmarks in T is obtained by linearly interpolating according to d, since only initial and last image positions of landmarks are stored during mapping to reduce map size. By taking the horizontal difference in image positions a heading correction for the robot can be established. To obtain a single correction value, the most likely horizontal position needs to be obtained since the matching process may produce outliers. Thus, this single deviation value is obtained by sampling all horizontal differences in a histogram and taking its mode.

While the robot moves forward at constant speed, this heading correction is performed for each frame until d equals the segment length stored in the map. The listing 2 presents the algorithm used to traverse or 'replay' one segment.

4. Segmentation-based path-following

The path-following method used as motion primitive of the proposed approach is based on a previous work by De Cristóforis et al. (2013). Here, a mobile robot is autonomously steered in order to remain inside a semi-structured path by means of image processing alone.

The method works by first dividing the image into regions or segments of similar appearance and then classifying these as belonging to traversable or non-traversable areas. It does not require a learning phase since it uses the area immediately in front of the robot in order to obtain a sample of traversable terrain. A gaussian-mixture model is used to describe this sample (using the Hue, Saturation and Value components) and compare to the rest of the regions in the image. After classification, the most likely group of interconnected image segments is used to obtain a contour of traversable path. With this contour a simple control law which infers the path middle line and steers the robot to follow it is used.

Since the method only requires the area of the image below the horizon level (i.e. the ground) to steer the robot, an initial horizon detection method can be used to infer this area and restrict processing to the meaninful portion of the image. Figure 2 shows the image processing pipeline.

4.1. Image Segmentation

In outdoor semi-structured or unstructured environments, the traversable area is often cluttered by small objects with a color of the forbidden area, for example grass, tree leaves, water or even snow in the middle of the road. In this case, most classification methods working at a pixel level, like the one proposed by Dahlkamp et al. (2006), would perform worse than methods which first segment the image to several areas with similar texture or color. Thus, the initial step in the path-detection process



(j) Super-pixel clas- (k) Filtered binary sification as binary mask mask

(1) Path contour

Fig. 2. Pipeline description: (a) input image as acquired from the camera, (b) cropped image below automatically detected horizon, (c) segment map. For each segment, mean and covariance are computed:(e), (i) and (g) show segment hue, saturation and value means respectively (d), (h) and(f) show pixel hue, saturation and value for reference. (j) binary mask obtained from classification using ground models from ROI. (k) morphological opening filter is applied in order to 'close' small gaps in the binary mask. (l) path contour extracted from processed binary mask and middle points computed, on top-left linear and angular speeds values obtained by control law.

is to segment the image into a set of super-pixels. After the image is cropped to contain the region below the horizon (either automatically or manually), the result is initially filtered to reduce noise (using a median-blur, which preserves edge contrast) and then segmented. The segmentation algorithm used is the graph-based approach proposed by Felzenszwalb and Huttenlocher (2004). The algorithm first constructs a fully connected graph where each node corresponds to a pixel in the image. Pixel intensities between neighbors are analyzed and edges are broken whenever a threshold is exceeded. The resulting unconnected sub-graphs define the segments (or super-pixels).

4.2. Classification

To achieve a robust classification of the segments, a probabilistic approach is used, based on Dahlkamp et al. (2006). This step determines if a population sample (a segment), modeled by a Gaussian probability distribution $\mathcal{N}(\mu, \Sigma)$, with mean vector μ and covariance matrix Σ , represents an instance of a more general model of the 'navigable path' class or not. This navigable path class is represented in turn by a Mixture-of-Gaussians (MoG) model. The steps involved in the classification task are as follows:

- 1. Segment model computation: each segment is modeled by its mean μ and covariance Σ of HSV color values.
- 2. Navigable class computation: a rectangular ROI (Region Of Interest) in the lower part of the image is used as a reference of navigable path (see Fig. 2). All models contained in this ROI are merged by similarity (using Mahalanobis distace). The resulting models which cover the ROI area by more than a defined percentage are used as references for navigable regions in the image.
- Classification of all segments: All segments in the image are compared to the reference models for navigable path using the Mahalanobis distance. A binary mask for all pixels is created indicating the membership to this class.
- 4. Contour extraction: the binary mask is post-processed by a morphological opening operation to reduce artifacts and from all regions classified as navigable path, the area intersecting the ROI region is chosen. The contour of this region is computed.

4.3. Motion control

The goal of the motion control is to correct the heading of the robot to keep it in the middle of the road. From the previously determined contour, a path center trajectory is estimated in order to guide the robot to the center of the detected path and maintain it away from path edges. First, by going rowby-row in the image, the middle point for the current row is obtained from the leftmost and rightmost pixel's horizontal positions of the path contour. The list of middle points is then used to compute angular and linear speeds with a simple yet effective control-law as follows. From the list of *n* horizontal values u_i of the *i*-th middle point of the detected path region, angular speed ω and forward speed *v* are computed as follows:

$$\omega = \alpha \sum_{i}^{n} \left(x_{i} - \frac{w}{2} \right), \qquad v = \beta n - |\omega| \tag{1}$$

where w is the width of the image and α and β are constants specific to the robot's speed controllers.

The effects of this control law are such that the robot will turn in the direction where there is the highest deviation, in average, from middle points with respect to a vertical line passing through the image center. This line can be assumed to be the position of the camera, which is mounted centered on the robot. Therefore, whenever a turn in the path or an obstacle is present, the robot will turn to remain inside the path and avoid the obstacle. The linear speed of the robot is reduced accordingly to the angular speed determined by the previous computation. This has the effect of slowing down the robot whenever it has to take a sharp turn.

5. Robot position error analysis

In this section we show that despite the fact that the robot localization is based on odometry, the robot position error does not diverge over time due to the heading corrections caused by the two movement primitives. Our analysis assumes that a mobile robot moves through a sequence of *n* waypoints by using reactive navigation algorithms (or motion primitives) that steer the robot towards the waypoint's connecting lines. We describe robot position at each waypoint in the waypoint's coordinate frame and show that the robot's relative position to a particular waypoint can be calculated by a recursive formula $\mathbf{x}_{k+1} = M_k x_k$, where the matrix \mathbf{M}_k eigenvalues are 1 and m_k . First, we show that m_k is lower than 1 because of the heading corrections introduced by the motion primitive used to navigate between k^{th} and $(k + 1)^{th}$ waypoint. Finally, we show that if all the path waypoints do not lie on a straight line, $\mathbf{M} = \prod k = 0^{n-1} M_k$ dominant eigenvalue is lower than one, meaning that the robot position error is reduced as it travels the entire path.

5.1. Robot movement model

Let us assume that the robot moves on a plane and its position and heading is x, y, φ respectively. Assume that the edge to be traversed starts at coordinate origin and leads in the direction of the *x*-axis.

The path following method presented in Section 4.1 calculates the robot's steering speed ω directly from image coordinates of the segmented path center. These coordinates are affected by both robot heading ϕ and lateral displacement y. Larger robot displacement and larger heading deviation cause larger values of the steering speed, i.e.

$$\omega \sim -c_0 y - c_1 \varphi. \tag{2}$$

Given that the robot's speed v is much lower compared to the maximal rotation speed ω , the robot can change its heading much faster than its horizontal displacement y. Thus, a robot steered by the method presented in Section 4 will quickly stabilize its heading and reach a state when its rotation speed ω will be close to zero. This allows to rewrite Equation 2 to

$$0 \approx -c_0 y - c_1 \varphi. \tag{3}$$

Thus, heading φ of a robot following the path can be expressed as

$$\varphi \approx -k_0 y, \tag{4}$$

where k_0 is a positive constant.

The case of the feature-based method is similar. If the robot is displaced in a lateral direction from the path, i.e. in the direction of the *y*-axis, the perceived features will not be visible at the expected positions, but will be shifted to the right for y > 0 and to the left y < 0.

Since the feature-based method steers the robot in a direction that minimizes the difference (in horizontal coordinates) of the expected and detected features, the robot heading will be stabilized at a value that is proportional to its lateral displacement. In other words, the robot heading φ satisfies a similar constrain as in the previous case:

$$\varphi \approx -k_1 y,$$
 (5)

where k_1 is a positive constant.

Thus, both of the motion primitives of the hybrid method adjust the robot heading φ proportionally to its lateral displacement from the path, i.e. $\varphi \approx -ky$, where k > 0. Both methods move the robot along the edge until the robots odometric counter reports that the distance travelled exceeded the length of the edge stored in the map. Assuming that the robot heading φ is small, we can state that

$$\frac{dy}{dx} = \varphi = -ky. \tag{6}$$

Let the initial position of the robot be (a_x, a_y) , where the values of (a_x, a_y) are significantly smaller than the segment length. Solving this differential equation with boundary conditions $a_y = f(a_x)$ allows us to compute the robot position:

$$y = a_y e^{-kx}.$$
 (7)

Considering a segment of length *s*, we can calculate the robot position (b_x, b_y) after it traverses the entire segment as:

$$b_x = a_x + s$$

$$b_y = a_y e^{-ks}.$$
(8)

Equation (8) would hold for an error-less odometry and noiseless camera. Considering the camera and odometry noise, Equation (8) will be rewritten to

$$\begin{pmatrix} b_x \\ b_y \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & e^{-sk} \end{pmatrix} \begin{pmatrix} a_x \\ a_y \end{pmatrix} + \begin{pmatrix} s\nu \\ \xi \end{pmatrix},$$
(9)

where v and ξ are normally distributed random variables, $\mu(v) = 1, \sigma(v) = \epsilon, \mu(\xi) = 0$ and $\sigma(\xi) = \tau$. A compact form of (9) is

$$\mathbf{b} = \mathbf{M}\mathbf{a} + \mathbf{s}.\tag{10}$$

For a segment with an arbitrary azimuth, one can rotate the coordinate system by the rotation matrix **R**, apply (10) and rotate the result back by \mathbf{R}^{T} . Thus, Equation (10) would become

$$\mathbf{b} = \mathbf{R}^{\mathrm{T}} (\mathbf{M}\mathbf{R}\mathbf{a} + \mathbf{s}) = \mathbf{R}^{\mathrm{T}}\mathbf{M}\mathbf{R}\mathbf{a} + \mathbf{R}^{\mathrm{T}}\mathbf{s}.$$
 (11)

Using (11), the robot position **b** at the end of the segment can be computed from its starting position **a**. Equations (11) allow to calculate robot position as it traverses along the intended path.

5.2. Position error

Let the robot initial position **a** be a random variable drawn from a two-dimensional normal distribution with the mean \hat{a} . Since Equation (11) has only linear and absolute terms, the final robot position **b** will have a normal distribution as well. Let $\mathbf{a} = \hat{\mathbf{a}} + \tilde{\mathbf{a}}$, where $\tilde{\mathbf{a}}$ is a random normal variable with zero mean. Assuming the same notation for **b** and **s**, Equation (11) can be rewritten as

$$\tilde{\mathbf{b}} = \mathbf{R}^{\mathrm{T}} \mathbf{M} \mathbf{R} (\hat{\mathbf{a}} + \tilde{\mathbf{a}}) + \mathbf{R}^{\mathrm{T}} (\hat{\mathbf{s}} + \tilde{\mathbf{s}}) - \hat{\mathbf{b}}.$$
 (12)

Substituting $\mathbf{R}^{T}\mathbf{M}\mathbf{R}\hat{\mathbf{a}} + \mathbf{R}^{T}\hat{\mathbf{s}}$ for $\hat{\mathbf{b}}$, Equation (12) becomes

$$\tilde{\mathbf{b}} = \mathbf{R}^{\mathrm{T}} \mathbf{M} \mathbf{R} \tilde{\mathbf{a}} + \mathbf{R}^{\mathrm{T}} \tilde{\mathbf{s}}, \tag{13}$$

where $\tilde{a}, \tilde{b}, \tilde{s}$ are Gaussian random variables with zero mean. The \tilde{a} and \tilde{b} represent the robot position error at the start and end of the traversed segment.

5.3. Traversing multiple segments

Let the robot path is a closed chain of *n* linear segments denoted by numbers from 0 to n - 1. Let a segment *i* be oriented in the direction α_i and its length be s_i . Let the robot positions at the start and end of the *i*th segment are \mathbf{a}_i and \mathbf{b}_i respectively. Since the segments are joined, $\mathbf{b}_i = \mathbf{a}_{i+1}$ and Equation (13) for the *i*th traveled segment is

$$\tilde{\mathbf{a}}_{i+1} = \mathbf{N}_i \tilde{\mathbf{a}}_i + \mathbf{R}_i^{\mathrm{T}} \tilde{\mathbf{s}}_i, \tag{14}$$

where $N_i = R_i^T M_i R_i$. Thus, the robot position after traversing the entire path consisting of the *n* segments will be

$$\tilde{\mathbf{a}}_{\mathbf{n}} = \check{\mathbf{N}}\tilde{\mathbf{a}}_{\mathbf{0}} + \check{\mathbf{t}},\tag{15}$$

where

$$\breve{\mathbf{N}} = \prod_{j=n-1}^{0} \mathbf{N}_{j} \text{ and } \breve{\mathbf{t}} = \sum_{\mathbf{i}=0}^{n-1} \left(\prod_{j=\mathbf{i}+1}^{n-1} \mathbf{N}_{j} \right) \mathbf{R}_{\mathbf{i}}^{\mathrm{T}} \widetilde{\mathbf{s}}_{\mathbf{i}}.$$
 (16)

If the robot traverses the entire path *k*-times, its position can be calculated in a recursive way by

$$\tilde{\mathbf{a}}_{(\mathbf{k}+1)\mathbf{n}} = \tilde{\mathbf{N}}\tilde{\mathbf{a}}_{\mathbf{k}\mathbf{n}} + \tilde{\mathbf{t}}.$$
(17)

Since the mean of every \tilde{s}_i is equal to zero and \tilde{s}_i have normal distribution, \check{t} , that is a linear combination of \tilde{s}_i , has a normal distribution with zero mean as well. Therefore, Equation (17) describes a linear discrete stochastic system, where the vector \check{t} represents a disturbance vector. If the system characterized by (17) is stable, then the position deviation of the robot from the intended path \tilde{a}_i does not diverge for $k \to +\infty$. The system is stable if all eigenvalues of \check{N} lie within a unit circle.

As every N_i equals to $R_i^T M_i R_i$, its eigenvalues lie on the diagonal of M_i and its eigenvectors constitute columns of R_i . Therefore, each matrix N_i is positive-definite and symmetric. Since the dominant eigenvalue of every Ni equals to one, eigenvalues of \check{N} are smaller or equal to one. Since the dominant eigenvalue of \check{N} is equal to one if and only if the dominant eigenvalues of products $N_{i+1}N_i$ are equal 1 for all *i*. However, dominant eigenvalue of a product $N_{i+1}N_i$ equals 1 only if the dominant eigenvectors of both Ni and Ni+1 are linearly dependent, which corresponds to collinearity of i^{th} and $(i + 1)^{th}$ segment. Thus, a dominant eigenvalue \breve{n} of the matrix \breve{N} equals 1 if and only if all path segments are collinear, i.e. the entire path is just a straight line. For any other case, the spectral radius of \tilde{N} is lower than one, meaning that the system described by (17) is stable. This means that if the robot travels the trajectory repeatably, its position error $\tilde{\mathbf{a}}_{\mathbf{kn}}$ does not diverge.

6. Experimental results

The proposed hybrid vision-based navigation system was implemented in C/C++ within the ROS (Robot Operating System) framework as a set of separate ROS modules¹. The performance of the navigation system has been evaluated in an indoor/outdoor environment. The experiments were performed

¹We will plan to release an open version of the software if the paper is accepted.

outside and inside of the Pabellón 1 building, Ciudad Universitaria, Buenos Aires, Argentina. Two sets of experiments were performed with two different robots. Two robotic platforms were used for experiments: a small differential caterpillartracked robot called ExaBot (Pedre et al., 2014) and a fourwheeled P3-AT from MobileRobots. Both platforms were equipped with a single off-the-shelf camera and a Core i5 laptop as the sensing and processing elements, respectively.

Additionally, the motion model correctness was also demonstrated in practice, by teaching and repeating a single straight segment in an outdoor scenario, using both methods. Finally, the computational efficiency of the method was measured.

6.1. Motion model correctness

The first scenario was aimed at verification of the assumptions given in Section 5.1. In particular, we wanted to calculate if the robot motion model established by equations 7 conforms to real robot motion along one straight segment. To verify the motion model, we have taught the robot to traverse a 5.6 m long straight path. Then, we executed the autonomous navigation system in order to repeat the path 5 times, with an initial lateral displacement of about 0.5 m from the original starting point. The robot motion was tracked using an external localization system Krajník et al. (2014). We used both segmentationand feature-based methods for this test. The recovered trajectories, compared with an expected exponential shape given by 7, are presented on Figure 3, where the stability of the system is appreciated. The trajectories described by both methods correspond to the exponential model presented in equation 7.

As can be seen, both navigation methods converge. However, it can be easily noted that in this outdoor scenario, the segmentation-based path following method converges much faster and with less error. Since the scene is comprised mostly of distant elements, the feature-based method aligns its view and converges much slower. On the other hand, on indoor scenarios, the convergence of the feature-based method is much faster. Thus, it makes much more sense to use a faster converging method such as path-following when the robot has to navigate outdoors through a structured or semi-structured path.

6.2. Algorithm computational efficiency

During the aforementioned experiment, we have also measured the real-time performance of both navigation algorithms. Typically, one cycle of the feature-based navigation algorithm takes approximately 30 ms, out of which one half is spent with the image features extraction and the other with the frameto-map matching. Segmentation of one image into superpixels typically takes 150 ms of a total of 215 ms that takes the whole algorithm. Thus, the feature- and segmentation-based algorithms issue about 30 and 4 steering commands per second respectively, which allows to quickly stabilize the robot at the heading desired.

6.3. Large-scale experiments

The purpose of the large-scale experiments was to test the robot in realistic outdoor conditions. However, it was not technically possible to cover the entire operation area by the



Fig. 3. Robot trajectory along a straight segment path. 3(a) shows the results using the feature-based method and 3(b) using the path following method. Five replays were performed for each case. The data points of the recorded trajectories are fitted using the equation $y = a_y e^{-kx}$ to verify the motion model used in the navigation method.

localization system. Therefore, the algorithm precision has been assessed in the same way as Chen and Birchfield (2009) and Krajník et al. (2010), i.e. the relative robot position has been measured each time the robot traversed an entire path.

In both cases, a closed path was first taught during a teleoperated run, traversing an open area around the entrance of the building and a square path inside the building hall. On the outdoor portion, when a path was available, the path-detection method was used. Otherwise, the feature-based navigation method was selected. During this training phase, a hybrid map was built, which was later used during autonomous navigation. For the P3-AT, the outdoor portion was considerably longer, of a total of 150 m (see Figure 4), while for the ExaBot, the path length was of 68 m. Two sets of experiments were performed with the P3-AT on different days using different mappings of the environment.

In order to test the robustness of the approach, a series of autonomous replays of the previously taught path were performed. While this type of navigation has been theoretically proved to reduce odometric errors by means of processing visual information, during the experiments the aim was to test this aspect in practice. To this end, the position of the robot at the end of each replay was measured and compared to the ending position. Furthermore, in order to better appreciate the previous error reduction effect, during the experiments performed with the ExaBot, the robot was laterally displaced by 0.6 m before the first repetition while still remaining inside the bounds of the path. This was possible due to the reduced size of the ExaBot. On the other hand, due to its larger size, the P3-AT was only displaced about 0.3 m. However, given such a long path, the expected errors of odometry alone (generally around of 10%) would be higher than this amount and what is of importance is that after several repetitions the robot always reaches the expected position without significant error.

For the first set of experiments with the ExaBot, the positions differences were measured by marking the final robot positions on the floor and measuring by hand. In contrast, for the second set of experiments performed with the P3-AT, an external visual localization system, WhyCon (Krajník et al., 2014), was used as ground-truth data of the robot position. The localization system allows obtaining the pose of the robot within its 2D plane of motion, with respect to a user set coordinate system, by detecting a circular pattern attached to the robot by means of an external fixed camera. This localization system has been shown to have around 1% relative error (in relation to the measurement area).

The results for both sets can be seen in Table 1, corresponding to five repeats for the case of the P3-AT and three for the ExaBot (due to limited autonomy). As can be seen, the second set of repeats with the P3-AT (second column) gave better results than the first one (first column). This may be because in the first case the learning phase was made in one day and the autonomous navigation phase in another, while the weather changed from cloudy/light rain to partly cloudy/sunny, affecting the learned map of the environment. In the second set of repeats with the P3-AT the whole experiment (both learning and replay phase) was done in the same day. In total, using the presented hybrid method, the P3-AT robot was able to autonomously navigate 1.5 km. Example images corresponding to the robot views during navigation can be seen in Figure 5.

	P3-AT				ExaBot		
replay	abs. err. [<i>m</i>]		rel. err. [%]		abs. err. [<i>m</i>]	rel. err. [%]	
1	0.40	0.09	0.27	0.06	0.12	0.12	
2	0.58	0.19	0.39	0.12	0.07	0.10	
3	0.90	0.20	0.60	0.13	0.03	0.05	
4	0.86	0.21	0.57	0.13	-	-	
5	0.80	0.25	0.53	0.16	-	-	
avg.	0.71	0.19	0.47	0.12	0.07	0.09	

Table 1. Position errors of the indoor/outdoor experiment with the hybrid method, for each robot. Errors were measured after each repetition with respect to the ending positions reached after training.

6.4. Performace comparison

To compare the algorithm precision with other teach-andrepeat systems, we have decided to use accuracy ε_{acc} and the repeatability ε_{rep} measures as introduced in Chen and Birchfield (2009). The ε_{acc} and εrep are computed as the RMS of Euclidean distance or a standard deviation of the robot's final



(a)



(d)



Fig. 4. Images of the experiments and of the path traversed by the robot around and inside the Pabellón 1 building, Ciudad Universitaria, Buenos Aires, Argentina. 4(a), 4(b)and 4(c) shows the autonomous navigation phase with the ExaBot, 4(e), 4(f) and 4(e) with the Pioneer 3P-AT, and 4(d) shows the environment where the experiments took place and the traversed path, taken from Google Maps.

positions from the path start by equations

$$\varepsilon_{acc} = \sqrt{\frac{1}{n} \sum_{i=1}^{n} \|\mathbf{x}_{i} - \mathbf{x}_{f}\|^{2}} \qquad \varepsilon_{rep} = \sqrt{\frac{1}{n} \sum_{i=1}^{n} \|\mathbf{x}_{i} - \mu\|^{2}},$$
(18)

where \mathbf{x}_i is the robot position after completing the *i*-th loop, \mathbf{x}_f the final position after training and $\mu = \sum_{i=1}^{n} \mathbf{c}_i / n$.

Table 2 compares repeatability and accuracy of the our results to the ones presented in Krajník et al. (2010) and Chen and Birchfield (2009). For this analysis the second experiment with the P3-AT was taken.

		Our method	Other methods	
		P3-AT	Krajnik	Chen
Accuracy	[m]	0.19	0.26	0.66
Repeatability	[m]	0.08	0.14	1.47
Path length	[m]	150	1040	40

Table 2. Comparison of the method's accuracy and repeatability to other works.

This summary indicates that the navigation accuracy of our



Fig. 5. Screenshots extracted from the ExaBot robot: 5(a), 5(b) and 5(i) during the segmentation-based navigation and 5(c), 5(d), 5(e), 5(f), 5(g) and 5(h) during the landmark-based navigation.

method is comparable with the closely related methods presented in articles Krajník et al. (2010); Chen and Birchfield (2009), showing better results for both accuracy and repeatability analysis.

7. Conclusion

In this paper we propose a hybrid visual navigation system for mobile robots in indoor/outdoor environments. The method enables to use both segmentation-based and feature-based navigation as elementary movement primitives for the entire system. A topological map of the environment is defined that can be thought as a graph, where the edges are navigable paths and the nodes are open areas. As we already saw segmentationbased navigation fits very well to path following (edges) and landmark-based navigation is suitable for open areas (nodes).

The presented method is robust and easy to implement and does not require sensor calibration or structured environment, and its computational complexity is independent of the environment size. We also present a convergence analysis showing that despite the robot localization within a segment is based on odometry, the robot position error does not diverge over time due to the heading corrections caused by the properties of the two movement primitives. We proved both theoretically and empirically that if the robot traverses a closed path repeatedly, its position error does not diverge. The aforementioned properties of the method allow even low-cost robots equipped only with a single, off-the-shelf camera to effectively act in large in mixed indoor/outdoor environments.

Acknowledgements

The research was supported by EU-ICT project 600623 STRANDS and CZ-ARG project 7AMB14AR015.

References

- Agrawal, M., Konolige, K., Blas, M., 2008. Censure: Center surround extremas for realtime feature detection and matching. Computer Vision– ECCV 2008, 102–115.
- Augustine, M., Ortmeier, F., Mair, E., Burschka, D., Stelzer, A., Suppa, M., 2012. Landmark-Tree map: A biologically inspired topological map for long-distance robot navigation, in: Int. Conference on Robotics and Biomimetics (ROBIO), pp. 128–135.
- Bay, H., Ess, A., Tuytelaars, T., Van Gool, L., 2008. Speeded-up robust features (SURF). Computer vision and image understanding 110, 346–359.
- Borenstein, J., Koren, Y., 1988. Obstacle avoidance with ultrasonic sensors. Robotics and Automation, IEEE Journal of 4, 213–218.
- Calonder, M., Lepetit, V., Strecha, C., Fua, P., 2010. Brief: Binary robust independent elementary features. Computer Vision–ECCV 2010, 778–792.
- Chaumette, F., Hutchinson, S., 2006. Visual servo control. i. basic approaches. Robotics & Automation Magazine, IEEE 13, 82–90.
- Chen, Z., Birchfield, S., 2009. Qualitative vision-based path following. Robotics, IEEE Transactions on 25, 749–754.
- Dahlkamp, H., Kaehler, A., Stavens, D., Thrun, S., Bradski, G., 2006. Selfsupervised monocular road detection in desert terrain, in: Proc. of Robotics: Science and Systems (RSS).
- De Cristóforis, P., Nitsche, M.A., Krajník, T., Mejail, M., 2013. Real-time monocular image-based path detection. Journal of Real-Time Image Processing.
- F. Chaumette and S. Hutchinson, 2007. Visual servo control. ii. advanced approaches [tutorial]. Robotics & Automation Magazine, IEEE 14, 109–118.
- Felzenszwalb, P., Huttenlocher, D., 2004. Efficient graph-based image segmentation. International Journal of Computer Vision 59, 167–181.
- Furgale, P., Barfoot, T.D., 2010. Visual teach and repeat for long-range rover autonomy. Journal of Field Robotics 27, 534–560.
- Kaliyaperumal, K., Lakshmanan, S., Kluge, K., 2001. An algorithm for detecting roads and obstacles in radar images. Vehicular Technology, IEEE Transactions on 50, 170–182.
- Krajník, T., De Cristóforis, P., Faigl, J., Szücsová, H., Nitsche, M., Preucil, L., Mejail, M., 2013. Image features for long-term mobile robot autonomy, in: ICRA 2013, Workshop on Long-Term Autonomy.
- Krajník, T., Faigl, J., Vonásek, V., Košnar, K., Kulich, M., Přeučil, L., 2010. Simple yet stable bearing-only navigation. Journal of Field Robotics 27, 511–533.
- Krajník, T., Nitsche, M., Faigl, J., Vaněk, P., Saska, M., Přeučil, L., Duckett, T., Mejail, M., 2014. A practical multirobot localization system. Journal of Intelligent Robotic Systems, 1–24.
- Krajnik, T., Nitsche, M., Pedre, S., Preucil, L., Mejail, M.E., 2012. A simple visual navigation system for an uav, in: Systems, Signals and Devices (SSD), 2012 9th International Multi-Conference on, IEEE. pp. 1–6. doi:10.1109/SSD.2012.6198031.
- Lowe, D., 2004. Distinctive image features from scale-invariant keypoints. International journal of computer vision 60, 91–110.
- Nitsche, M., Pire, T., Krajnk, T., Kulich, M., Mejail, M., 2014. Monte carlo localization for teach-and-repeat feature-based navigation, in: Mistry, M., Leonardis, A., Witkowski, M., Melhuish, C. (Eds.), Advances in Autonomous Robotics Systems, Springer International Publishing. pp. 13–24. doi:10.1007/978-3-319-10401-0_2.
- Ok, K., Ta, D., Dellaert, F., 2012. Vistas and wall-floor intersection featuresenabling autonomous flight in man-made environments. Workshop on Visual Control of Mobile Robots (ViCoMoR): IEEE/RSJ International Conference on Intelligent Robots and Systems.
- Ostafew, C.J., Schoellig, A.P., Barfoot, T.D., 2013. Visual teach and repeat, repeat, repeat: Iterative Learning Control to improve mobile robot path tracking in challenging outdoor environments. 2013 IEEE/RSJ International Conference on Intelligent Robots and Systems, 176–181doi:10.1109/IROS. 2013.6696350.
- Pedre, S., Nitsche, M., Pessacg, F., Caccavelli, J., De Cristóforis, P., 2014. Design of a multi-purpose low-cost mobile robot for research and education, in: Towards Autonomous Robotic Systems, Springer.
- Pfrunder, A., Schoellig, A.P., Barfoot, T.D., 2014. A proof-of-concept demonstration of visual teach and repeat on a quadrocopter using an altitude sensor and a monocular camera, in: Computer and Robot Vision (CRV), 2014 Canadian Conference on, IEEE. pp. 238–245.
- Surmann, H., Lingemann, K., Nüchter, A., Hertzberg, J., 2001. A 3d laser range finder for autonomous mobile robots, in: Proceedings of the 32nd ISR (International Symposium on Robotics), pp. 153–158.