# LiPo-LCD: Combining Lines and Points for Appearance-based Loop Closure Detection

Joan P. Company-Corcoles joanpep.company@uib.es Emilio Garcia-Fidalgo emilio.garcia@uib.es

Alberto Ortiz alberto.ortiz@uib.es Department of Mathematics and Computer Science, University of the Balearic Islands, and IDISBA (Institut d'Investigacio Sanitaria de les Illes Balears), Palma de Mallorca, Spain

#### Abstract

Visual SLAM approaches typically depend on loop closure detection to correct the inconsistencies that may arise during the map and camera trajectory calculations, typically making use of point features for detecting and closing the existing loops. In low-textured scenarios, however, it is difficult to find enough point features and, hence, the performance of these solutions drops drastically. An alternative for human-made scenarios, due to their structural regularity, is the use of geometrical cues such as straight segments, frequently present within these environments. Under this context, in this paper we introduce LiPo-LCD, a novel appearance-based loop closure detection method that integrates lines and points. Adopting the idea of incremental Bag-of-Binary-Words schemes, we build separate BoW models for each feature, and use them to retrieve previously seen images using a late fusion strategy. Additionally, a simple but effective mechanism, based on the concept of island, groups similar images close in time to reduce the image candidate search effort. A final step validates geometrically the loop candidates by incorporating the detected lines by means of a process comprising a line feature matching stage, followed by a robust spatial verification stage, now combining both lines and points. As it is reported in the paper, LiPo-LCD compares well with several state-of-the-art solutions for a number of datasets involving different environmental conditions.

# **1** Introduction

It is well known that many visual SLAM solutions rely on point features because of their wider applicability in general terms [23, 28]. Human-made environments, however, can lack

It may be distributed unchanged freely in print or electronic forms.



Figure 1: General overview of the proposed loop closure detection system.

texture and thus give rise to a low number of detected features. Nevertheless, precisely because of their nature, these environments usually exhibit structural regularities that can be described using richer features such as lines, which can be more robust and less sensitive to illumination changes. Several solutions can be found in the literature describing approaches combining both kinds of features, points and lines [1], [1]]. However, despite their success, most of them rely exclusively on feature points during the LCD stage, discarding information about lines that may be useful to improve the association performance for textureless environments. Other approaches opt for using holistic image representations [1], [2], [3], which can be faster to compute but less tolerant to visual changes, while, lately, solutions based on Convolutional Neural Networks (CNNs) [2], [3], [4], [5]] have shown to exhibit enhanced robustness and general performance, although they are still disengaged from real-time SLAM problems [1], [5]]. This is because they tend to require significant computational resources, e.g. on-board GPU, which makes them not suitable for mobile robotics in all cases.

The Bag of Words (BoW) model [29, 54], in combination with an inverted file, is arguably the most used indexing scheme for appearance-based loop closure detection [14, 25]. Depending on how the visual vocabulary is generated, BoW-based solutions can be classified into off-line and on-line approaches. Off-line solutions generate the visual dictionary during a training phase [11, 12, 22], what can be high time-consuming, while the general application of the resulting vocabulary becomes highly dependent on the diversity of the training set. As an alternative, there are approaches that propose to generate the dictionary on-line [11, 12, 21, 53, 59]. Moreover, binary descriptors [12, 13, 21] have emerged recently as an alternative to real-valued descriptors for BoW models [11, 12], since they offer advantages in terms of computational time and memory requirements. Additionally, similarity calculations can be performed using the Hamming distance or an of its variations, what can be efficiently implemented in modern processors.

Under this context, in this paper, we introduce *Lines and Points Loop Closure Detection* (LiPo-LCD), a novel appearance-based loop closure detection approach which combines points and lines. For a start, both features are described using binary descriptors. Next, an incremental BoW scheme is used for feature indexing. Lines and points are maintained separately into two incremental visual vocabularies and employed in parallel to obtain loop closure candidates efficiently. To combine the information provided by the two vocabularies, we propose a late fusion method based on a ranked voting system. Finally, to discard false positives, we improve the typical spatial verification step integrating lines into the procedure through: (1) a line matching strategy which includes structural information to achieve a higher number of matching candidates; and (2) a line representation tolerant to occlusions, which is combined with points into an epipolarity analysis step. A set of experiments validating LiPo-LCD and characterizing its performance against several state-of-the-art solutions is



Figure 2: (left) A human-made environment including a high number of lines and a low number of points. (right) An outdoor environment presenting the opposite situation.

reported at the end of the paper.

Our approach follows a dual scheme to combine points and lines, such as the solutions proposed by [12], 13]. Nonetheless, LiPo-LCD takes advantage of an incremental BoW strategy and incorporates lines into the spatial verification procedure that does not require map information, increasing its ability to be adapted to the operating environment, requiring only a monocular camera, and improving the performance in several datasets, as shown later.

# 2 Overview of the Loop-Closure Detection Approach

Figure 1 illustrates the approach proposed for loop closure detection. As can be observed, incremental visual vocabularies, along with the corresponding inverted files, are maintained independently for each visual feature. When a new image is sampled, a set of line and point binary descriptors is computed and used to (1) update the corresponding visual vocabulary and (2) obtain a list of the most similar images from each vocabulary. Next, the two lists are fused using a ranked voting procedure to obtain a final set of loop-closing candidates. To avoid adjacent images from competing with each other as loop candidates, we group images close in time using the concept of *dynamic island* [**L**]. Among the resulting islands, the one best corresponding with the query image is selected, while its representative image is geometrically assessed against the query to accept/reject the loop. The details about the aforementioned processes can be found next.

## 2.1 Image Description

As stated previously, LiPo-LCD describes images using lines and points. The rationale behind this approach is that the combination of multiple, complementary description techniques is a way leading to improving the performance and robustness of the loop closing method [1]. In our solution, the image  $I_t$  at time t is described by  $\phi(I_t) = \{P_t, L_t\}$ , being  $P_t$  a set of local keypoint descriptors and  $L_t$  a set of line descriptors, both deriving from  $I_t$ . These two descriptions complement each other to make image representation more robust: while some environments may be described more distinctively using lines than points, i.e. textureless scenes, others lacking structure will benefit from keypoints, and the net result is a joint descriptor of a wider scope. Figure 2 illustrates this issue for two environments.

#### 2.1.1 Point Description

Given the above-mentioned advantages about binary descriptors, in this work, we have opted for detecting and describing points using ORB [ $\square$ ]. Although the proposed strategy can be used with any other binary descriptor, we employ ORB because of its robustness to rotation, scale and illumination changes [ $\square$ ]. The *m* ORB descriptors found at image  $I_t$  define the point descriptor as  $P_t = \{p_0^t, p_1^t, \dots, p_{m-1}^t\}$ .

#### 2.1.2 Line Description

Lines are found using the Line Segment Detector (LSD) [ $\square$ ]. LSD is a linear-time line segment detector that provides high-precision results and subpixel accuracy without parameter tuning. On the one hand, detected lines are described using a binary form of the Line Band Descriptor (LBD) [ $\square$ ]. In the original implementation, a rectangular region centred on each line is considered. Such region is divided into a set of bands  $B_i$ , from which a descriptor  $BD_i$  is computed contrasting  $B_i$  with its neighbouring bands. On the other hand, the binary descriptor is finally obtained considering 32 possible pairs of band descriptors  $BD_i$  within the support region. Each pair is compared bit by bit, generating an 8-bit string per pair. A final 256-bit descriptor is generated concatenating the resulting strings for all pairs. The set of *n* LBD binary descriptors for image  $I_t$  defines the line descriptor  $L_t = \{l_0^t, l_1^t, \dots, l_{n-1}^t\}$ .

#### 2.2 Searching for Loop Closure Candidates

To index and retrieve loop closure candidates, we rely on the *OBIndex2* approach [**L3**], a hierarchical tree structure to manage an increasing number of binary descriptors in an efficient way. This structure can then be used as an incremental BoW scheme and combined with an inverted file for fast image retrieval. The reader is referred to [**L3**] for further detail.

Given that LiPo-LCD describes all visual features using binary descriptors, we maintain two instances of *OBIndex2*, one for points and one for lines. Each instance builds an incremental visual vocabulary along with an index of images for each feature. Given an image  $I_t$ , a parallel search is performed on each index to retrieve the most similar images of points and lines. As a result, two lists are obtained: (1) the *m* most similar images using points  $C_p^t = \{I_{p_0}^t, \dots, I_{p_{m-1}}^t\}$  and (2) the *n* most similar images using lines  $C_l^t = \{I_{l_0}^t, \dots, I_{l_{n-1}}^t\}$ . Each list is sorted by, respectively, their associated scores  $s_p^t(I_t, I_j^t)$  and  $s_l^t(I_t, I_j^t)$ , which measure the similarity between the query image  $I_t$  and the image  $I_j$ . Since the range of these scores varies depending on the distribution of visual words for each vocabulary, they are mapped onto the range [0,1] using *min-max normalization* as follows:

$$\tilde{s}_{k}^{t}\left(I_{t}, I_{j}^{t}\right) = \frac{s_{k}^{t}\left(I_{t}, I_{j}^{t}\right) - s_{k}^{t}\left(I_{t}, I_{min}^{t}\right)}{s_{k}^{t}\left(I_{t}, I_{max}^{t}\right) - s_{k}^{t}\left(I_{t}, I_{min}^{t}\right)},\tag{1}$$

where  $s_k^t(I_t, I_{min}^t)$  and  $s_k^t(I_t, I_{max}^t)$  respectively corresponds to the minimum and the maximum scores of an image candidate list, being  $k \in \{p, l\}$ . Images whose normalized score  $\tilde{s}_k^t$  is lower than a threshold are discarded to limit the maximum number of candidates. Additionally, the current image descriptors are used to update the visual vocabularies appropriately.

### 2.3 Merging Lists of Candidates

The two resulting lists  $C_p^t$  and  $C_l^t$  provide loop closure candidates from each individual perspective. Thus, the next step is to combine both lists to obtain an overall overview of possible candidates but considering lines and points altogether. In this regard, the literature comprises multiple techniques to combine multimodal information for image retrieval  $[\Box]$ . These can generally be categorized into two schemes, namely *early* and *late fusion*: while the former combines all features into a single representation before being processed, the latter works at the decision level, combining the outputs produced by different retrieval systems. In our proposal, given the heterogeneity of the features to combine, we rely on a late fusion approach that employs a ranked voting system based on the Borda count [5] to merge lists of candidates. This is a simple data fusion form based on democratic election strategies: first, a set of voters rank a list of fixed candidates on the basis of their preferences; scores are next given to each candidate in inverse proportion to their ranking; finally, once all votes have been emitted, the candidate with the highest number of votes wins. In LiPo-LCD, two independent voters, one for each visual vocabulary, emit an different-size ordered list of candidates  $C_{k}^{t}$ . The number of candidates c to vote for is set as the minimum length of the two lists. Next, top-*c* images on each list  $C_k^t$  are ranked with a score  $b_k$  as:

$$b_k(I_i^t) = (c-i)\tilde{s}_k^t \left(I_t, I_i^t\right), \qquad (2)$$

where *i* denotes the order of the image  $I_i$  in the list  $C_k^t$  and  $\tilde{s}_k^t(I_t, I_i^t)$  is the normalized score of the image in that list. For each image that appears in both lists, a combined Borda score  $\beta$  is computed as the geometric mean of the individual scores:

$$\beta(I_i^t) = \sqrt{b_p(I_i^t) b_l(I_i^t)}.$$
(3)

We employ the geometric mean instead of the arithmetic mean to reduce the influence of false positives in one of the lists. An integrated image list  $C_{pl}^t$  results next by sorting the scores  $\beta(I_i^t)$  of all the retrieved images. This list merges information from the two visual vocabularies, independently of the number of features detected in the current environment. Finally, to deal with the fact that some environments mostly exhibit one type of feature, images that only appear in one list are also incorporated into  $C_{pl}^t$ , although penalized.

## 2.4 Dynamic Islands Computation

In pursuit of selecting a final loop closure candidate, in this stage we verify the temporal consistency of the images retrieved in  $C_{pl}^t$ . To this end, we rely on the concept of *dynamic islands* used by iBoW-LCD [ $\square$ ]. This method permits to avoid images competing among them as loop candidates when they come from the same area of the environment. A dynamic island  $\Upsilon_n^m$  groups the images whose timestamps range from *m* to *n*. Initially, a set of islands  $\Gamma_t$  for the current image  $I_t$  is computed considering images in the list  $C_{pl}^t$  sequentially: every image  $I_i \in C_{pl}^t$  is either associated to an existing island  $\Upsilon_n^m$  if the image timestamp lies in the [m, n] interval or else is used to create a new island. After processing all images in  $C_{pl}^t$ , a global score g is computed for each island as:

$$g(\Upsilon_n^m) = \frac{\sum_{i=m}^n \beta(I_i^i)}{n-m+1}.$$
(4)

Unlike  $[\square]$ , where only points are considered, in LiPo-LCD, score g is the average of the Borda scores of the images belonging to the island, integrating both points and lines. Finally, the resulting set of islands  $\Gamma_t$  is sorted in descending order according to g. Next step is to select one of the resulting islands, denoted by  $\Upsilon^*(t)$ , to determine which area of the environment is the one most likely closing a loop with  $I_t$ . iBoW-LCD makes use of the concept of *priority islands*, defined as the islands in  $\Gamma_t$  that overlap in time with the island selected at time t - 1,  $\Upsilon^*(t - 1)$ . This is inspired by the fact that consecutive images should close loops with areas of the environments where previous images also closed a loop. iBoW-LCD selects, as a final island, the priority island with the highest score g, if any. However, this approach is just based on the appearance of the images and, therefore, due to perceptual aliasing, it might produce incorrect island associations in some human-made environments. For this reason, LiPo-LCD proposes a simple but effective modification of the original approach that only retains an island for the next time step if the final selected loop candidate satisfies the spatial verification procedure explained in Section 2.5. Once the best island  $\Upsilon^*(t)$  has been determined, the image  $I_c$  with the highest Borda score  $\beta$  of  $\Upsilon^*(t)$ is selected as its representative and evaluated in the next verification stage.

## 2.5 Spatial Verification

Although the BoW scheme is a good starting point to find loop closure candidates, to finish, we perform a final geometric check to take into account the spatial arrangement of the image features and avoid perceptual aliasing. This final step comprises an epipolarity analysis between the current image  $I_t$  and the loop candidate  $I_c$  on the basis of the number of inliers that support the roto-translation of the camera (after computing the fundamental matrix F using RANSAC). If the number of inliers is not high enough, the loop hypothesis is rejected.

The epipolarity analysis is typically carried out using a putative set of point matchings. However, as stated along this paper, point features might not be helpful because of the nature of the environment, and hence integrating lines into the geometric check can be useful, apart from the fact that straight segments can tolerate partial occlusions. To this end, LiPo-LCD makes use of (1) a novel line feature matching approach and (2) incorporates these line matchings, together with point matchings, into the geometric check. To match points, we make use of the available ORB descriptors, the Hamming distance and the Nearest Neighbour Distance Ratio (NNDR) [2].

#### 2.5.1 Line Feature Matching

Although NNDR is normally useful to discard false matchings between keypoints, it performs poor in respect to line descriptors matching, especially in human-made environments where line descriptors tend to be affected by perceptual aliasing [1]. To enhance line matching performance, the authors of [1] combine structural and appearance information in a relational graph. Despite their good results, their approach requires a high amount of memory and does not escalate well with the number of lines. In this work, we propose a much simpler but effective method to combine structural and appearance information for line feature matching. First, for each line descriptor  $l_i^t$  in the current image  $I_t$ , we retrieve an ordered list of the most similar line descriptors of the candidate image  $I_c$ . Next, to deal with camera rotations, we compute a global rotation  $\theta_g$  between the two frames as explained in [1].  $\theta_g$ 



Figure 3: Loop closure detections found in the L6I dataset using different visual features (Points, Lines, Points + Lines), and the associated ground truth. White dots represent a loop closure detected.

is next used to compute the relative orientation  $\alpha_i^j$  between each pair of lines as:

$$\boldsymbol{\alpha}_{i}^{j} = |\boldsymbol{\theta}_{i}^{t} - \boldsymbol{\theta}_{j}^{c} + \boldsymbol{\theta}_{g}|, \qquad (5)$$

being  $\theta_i^l$  the orientation of the line on the current image and  $\theta_j^c$  the orientation of their corresponding line in the list. For each list, all line matchings with high values of  $\alpha_i^j$  are discarded, and, as a result, a filtered list of line candidate matchings is obtained. To generate the final set of line matchings, we choose the two most similar surviving nearest neighbours from each list and apply the NNDR test.

#### 2.5.2 Epipolar Geometry Analysis Combining Points and Lines

Works described in [**f**, **[f]**] compute the fundamental matrix F from homographies estimated from line segment matchings across images, provided these segments lie in at least two different planes. LiPo-LCD makes use of a simpler but effective approach that avoids this constraint. On the one hand, differently to other representations that can be found in the literature [**[22**, **[11**], **[12**]], in this work, line segments are represented by their endpoints. On the other hand, endpoints are first matched between matching lines and next regarded as additional point correspondences for F computation. To associate segment endpoints (taking into account that a starting point of a line might correspond to the end point of the line in the other image), we select that pair that minimizes the rotation between lines using lines orientation and the global rotation  $\theta_g$ , as computed in Eq. 5. We consider a candidate line matching as an inlier if at least one endpoint pair supports the geometric model.

# **3** Experimental Results

In this section, we evaluate the performance of LiPo-LCD using several public datasets. LiPo-LCD is also compared against some state-of-the-art solutions. All experiments were performed on an Intel Core i7-9750H (2.60 GHz) processor with 16 GB RAM.

## 3.1 Methodology

Precision-recall metrics are used to evaluate the system. Given that false detections can be critical if LiPo-LCD is used in a real SLAM solution, we are especially interested in observing the maximum recall that can be achieved at 100% precision. OBIndex2 and iBoW-LCD



Figure 4: P-R curves for each dataset. P is 1.0 for all R values lower than 0.75.

	CC	EuR5	K00	K06	L6I	L60	MLG
NNDR	8.54	8.91	15.06	12.35	7.43	3.47	4.24
Proposed	18.15	19.21	25.37	22.51	10.14	8.45	11.34

Table 1: Average number of line inliers after the epipolar geometric analysis using NNDR and the proposed line feature matching method.

were configured as explained in [13]. The rest of the approaches shown in this section were executed using the default parameters proposed by their original authors. The following datasets were considered to validate LiPo-LCD: CityCentre [11] (CC), EuRoC Machine Hall 05 [1] (EuR5), KITTI 00 [16] (K00), KITTI 06 [16] (K06), Lip6Indoor [10] (L6I), Lip6Outdoor [11] (L6O) and Malaga 2009 Parking 6L [12] (MLG). These datasets encompass a wide range of environments including, for instance, urban and indoor scenarios, which are usually rich in lines, or outdoor scenarios, where points predominate over lines. For each dataset, we use the ground truth provided by the original authors except for the KITTI sequences, where we employ the one provided by [1], and the EuR5 and MLG datasets, where we use the files provided by [13].

## 3.2 General Performance

First, we validate the combination of points and lines proposed in this work. To this end, Fig. 3 shows the loop closures detected by LiPo-LCD using points, lines and both features, as well as the ground truth for the L6I dataset, whose images are poor in feature points. As can be observed, system performance increases when points and lines are used together as visual features. To measure the global performance of the system, Fig. 4 shows precision-recall curves for LiPo-LCD and for each dataset. As can be observed, high recall rates are

	FE	VU	SC	SV
Points	18.05	183.16	146.73	-
Lines	17.60	23.76	18.13	-
Parallel	19.05	196.58	159.01	15.09

Table 2: Average response time (ms) per image, calculated for each part of the pipeline. These times were computed over the K00 dataset. FE: Feature Extraction; VU: Vocabulary Update; SC: Search for Candidates; SV: Spatial Verification.

	CC	EuR5	K00	K06	L6I	L60	MLG
Bampis [5]	71.14	n.a.	96.53	n.a.	52.22	58.32	87.56
Gálvez-López [🗖]	31.61	n.a.	n.a.	n.a.	n.a.	n.a.	74.75
Mur-Artal [🔼]	43.03	n.a.	n.a.	n.a.	n.a.	n.a.	81.51
Cummins [	38.77	n.a.	49.2	55.34	n.a.	n.a.	68.52
Stumm 🗳	38.00	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.
Gomez-Ojeda [🗳]	n.a.	1.61	75.93	56.94	n.a.	n.a.	n.a.
Tsintotas [59]	n.a.	83.7	97.5	n.a.	n.a.	50.0	85.0
Tsintotas [53]	n.a.	69.2	93.2	n.a.	n.a.	n.a.	87.9
Angeli 🔲	n.a.	n.a.	n.a.	n.a.	36.86	23.59	n.a.
Zhang [12]	41.2	n.a.	n.a.	n.a.	n.a.	n.a.	82.6
Gehrig 🗳	n.a.	71.0	93.1	n.a.	n.a.	n.a.	n.a.
Khan [🛄]	38.92	n.a.	n.a.	n.a.	41.74	25.58	78.13
Garcia-Fidalgo [🗳]	88.25	n.a.	76.50	95.53	83.18	85.24	n.a.
LiPo-LCD	89.30	81.94	97.80	97.38	85.24	97.31	75.73

# Table 3: Maximum recall at 100% precision for several off-line approaches (top), on-line approaches (middle) and the proposed solution (bottom). Winners are indicated in bold face.

always achieved while maintaining the precision at 100%. Moreover, LiPo-LCD exhibits very stable behaviour in all cases.

Next, we evaluate our novel line feature matching strategy. For that purpose, we compute the average number of line inliers on each dataset using either a classical NNDR approach for lines and our approach. Results are shown in Table 1. As can be seen, the proposed line matching technique achieves a higher number of inliers in all datasets, even in sequences with severe appearance changes.

Finally, we evaluate the performance of LiPo-LCD in terms of computational times. The results obtained can be found in Table 2. We show results for K00 since it is the largest dataset considered in this work. We measure the average execution time in milliseconds for each stage of the pipeline, not taking into account times for merging lists of candidates and island selection, since they are negligible. The average response time of the whole system per image turns out to be 389.79 ms using a parallel implementation. As can be observed, feature extraction steps are very fast in all cases. The vocabulary update and the search for candidates steps are slower for points, due to the number of features to handle on each case. The spatial verification stage is always performed using points and lines together, and, hence, times for each feature separately are not available.

### 3.3 Comparison with Other Solutions

In this last section, LiPo-LCD is compared with other solutions. Table 3 shows the maximum recall achieved at 100% precision for all approaches. The results reported come from the original works, except for [ $\square$ ], which was executed by ourselves using the vocabularies and the default parameters provided by their authors. Results not available are indicated by *n.a.* As can be observed, LiPo-LCD achieves, in most cases, a higher recall than the other solutions. This is particularly interesting regarding the L6I dataset, where the combination of points and lines allows us to increase the performance in a low-textured scenario. It is also worth mentioning that LiPo-LCD outperforms [ $\square$ ], which is perhaps the most similar solution to ours.

# 4 Conclusions

In this work, we have described LiPo-LCD, an appearance-based loop closure detection method that combines points and lines. This combination allows us to detect loops in environments poor of feature points. Moreover, points and lines are described using binary descriptors for execution time reduction. To obtain loop closure candidates from both visual clues, we rely on a dual incremental BoW scheme. A late fusion method for merging both lists of candidates, based on the Borda count, is also proposed. The loop candidate hypothesis is finally validated by means of a geometrical check, which involves both points and lines. LiPo-LCD compares favourably with several state-of-the-art methods under different environmental conditions.

# Acknowledgements

This work is partially supported by EU-H2020 projects BUGWRIGHT2 (GA 871260) and ROBINS (GA 779776), and by projects PGC2018-095709-B-C21 (MCIU/AEI/FEDER, UE), and PROCOE/4/2017 (Govern Balear, 50% P.O. FEDER 2014-2020 Illes Balears). This publication reflects only the authors views and the European Union is not liable for any use that may be made of the information contained therein.

# References

- [1] Adrien Angeli, David Filliat, Stéphane Doncieux, and Jean-Arcady Meyer. A fast and incremental method for loop-closure detection using bags of visual words. *IEEE Transactions on Robotics*, 24(5):1027–1037, 2008. ISSN 15523098.
- [2] R. Arandjelović, P. Gronat, A. Torii, T. Pajdla, and J. Sivic. NetVLAD: CNN architecture for weakly supervised place recognition. In *International Conference on Computer Vision and Pattern Recognition*, 2016.
- [3] R. Arroyo, P. F. Alcantarilla, L. M. Bergasa, and E. Romera. Fusion and binarization of CNN features for robust topological localization across seasons. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 4656–4663, 2016.
- [4] Roberto Arroyo, Pablo. F Alcantarilla, Luis. M. Bergasa, J. Javier Yebes, and Sebastian Bronte. Fast and effective visual place recognition using binary codes and disparity information. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3089–3094, 2014.
- [5] Loukas Bampis, Angelos Amanatiadis, and Antonios Gasteratos. Fast loop-closure detection using visual-word-vectors from image sequences. *International Journal of Robotics Research*, 37(1):62–82, 2018.
- [6] H. Bay, V. Ferraris, and L. Van Gool. Wide-baseline stereo matching with line segments. In *International Conference on Computer Vision and Pattern Recognition*, volume 1, pages 329–336 vol. 1, 2005.

- [7] N. Bhowmik, R. González V., V. Gouet-Brunet, H. Pedrini, and G. Bloch. Efficient fusion of multidimensional descriptors for image retrieval. In *IEEE International Conference on Image Processing*, pages 5766–5770, 2014.
- [8] Jose-Luis Blanco, Francisco-Angel Moreno, and Javier Gonzalez. A collection of outdoor robotic datasets with centimeter-accuracy ground truth. *Autonomous Robots*, 27 (4):327, 2009.
- [9] Michael Burri, Janosch Nikolic, Pascal Gohl, Thomas Schneider, Joern Rehder, Sammy Omari, Markus W Achtelik, and Roland Siegwart. The EuRoC micro aerial vehicle datasets. *International Journal of Robotics Research*, 35(10):1157–1163, 2016.
- [10] Mark Cummins and Paul Newman. FAB-MAP: probabilistic localization and mapping in the space of appearance. *International Journal of Robotics Research*, 27(6):647–665, 2008.
- [11] Mark Cummins and Paul Newman. Appearance-only SLAM at large scale with FAB-MAP 2.0. International Journal of Robotics Research, 30(9):1100–1123, 2011.
- [12] D. Galvez-López and J. D. Tardos. Bags of binary words for fast place recognition in image sequences. *IEEE Transactions on Robotics*, 28(5):1188–1197, 2012.
- [13] E. Garcia-Fidalgo and A. Ortiz. iBoW-LCD: an appearance-based loop-closure detection approach using incremental bags of binary words. *IEEE Robotics and Automation Letters*, 3(4):3051–3057, 2018.
- [14] Emilio Garcia-Fidalgo and Alberto Ortiz. Vision-based topological mapping and localization methods: A survey. *Robotics and Autonomous Systems*, 64:1 – 20, 2015. ISSN 0921-8890.
- [15] Mathias Gehrig, Elena Stumm, Timo Hinzmann, and Roland Siegwart. Visual place recognition with probabilistic voting. In *IEEE International Conference on Robotics* and Automation, pages 3192–3199, 2017.
- [16] A. Geiger, P. Lenz, and R. Urtasun. Are we ready for autonomous driving? The KITTI vision benchmark suite. In *International Conference on Computer Vision and Pattern Recognition*, pages 3354–3361, 2012.
- [17] Ruben Gomez-Ojeda, David Zuñiga-Noël, Francisco-Angel Moreno, Davide Scaramuzza, and Javier Gonzalez-Jimenez. PL-SLAM: a stereo SLAM system through the combination of points and line segments. arXiv preprint arXiv:1705.09479, 2017.
- [18] R. Grompone von Gioi, J. Jakubowicz, J. Morel, and G. Randall. LSD: A fast line segment detector with a false detection control. *IEEE Transactions on Pattern Analysis* and Machine Intelligence, 32(4):722–732, 2010. ISSN 1939-3539.
- [19] Stephen Hausler and Michael Milford. Hierarchical multi-process fusion for visual place recognition. *arXiv preprint arXiv:2002.03895*, 2020.
- [20] Sheraz Khan and Dirk Wollherr. IBuILD: Incremental bag of binary words for appearance based loop closure detection. In *IEEE International Conference on Robotics and Automation*, pages 5441–5447, 2015.

- [21] M. Labbé and F. Michaud. Appearance-based loop closure detection for online largescale and long-term operation. *IEEE Transactions on Robotics*, 29(3):734–745, 2013.
- [22] K. Li, J. Yao, M. Xia, and L. Li. Joint point and line segment matching on wide-baseline stereo images. In *IEEE Winter Conference on Applications of Computer Vision*, pages 1–9, 2016.
- [23] H. Lim, J. Lim, and H. J. Kim. Real-time 6-DOF monocular visual SLAM in a largescale environment. In *IEEE International Conference on Robotics and Automation*, pages 1532–1539, 2014.
- [24] David G Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2):91–110, 2004.
- [25] S. Lowry, N. Sünderhauf, P. Newman, J. J. Leonard, D. Cox, P. Corke, and M. J. Milford. Visual place recognition: A survey. *IEEE Transactions on Robotics*, 32(1): 1–19, 2016.
- [26] M. J. Milford and G. F. Wyeth. SeqSLAM: visual route-based navigation for sunny summer days and stormy winter nights. In *IEEE International Conference on Robotics* and Automation, pages 1643–1649, 2012.
- [27] R. Mur-Artal and J. D. Tardós. Fast relocalisation and loop closing in keyframe-based SLAM. In *IEEE International Conference on Robotics and Automation*, pages 846– 853, 2014.
- [28] R. Mur-Artal and J. D. Tardós. ORB-SLAM2: An open-source SLAM system for monocular, stereo, and RGB-D cameras. *IEEE Transactions on Robotics*, 33(5):1255– 1262, 2017.
- [29] David Nister and Henrik Stewenius. Scalable recognition with a vocabulary tree. In *International Conference on Computer Vision and Pattern Recognition*, volume 2, pages 2161–2168, 2006. ISBN 0769525970.
- [30] Oscar A. Pellejero, Carlos Sagüés, and J. Jesús Guerrero. Automatic computation of the fundamental matrix from matched lines. In *Current Topics in Artificial Intelligence*, pages 197–206. Springer, 2004. ISBN 978-3-540-25945-9.
- [31] A. Pumarola, A. Vakhitov, A. Agudo, A. Sanfeliu, and F. Moreno-Noguer. PL-SLAM: real-time monocular visual SLAM with points and lines. In *IEEE International Conference on Robotics and Automation*, pages 4503–4508, 2017.
- [32] Ethan Rublee, Vincent Rabaud, Kurt Konolige, and Gary Bradski. ORB: An efficient alternative to SIFT or SURF. In *International Conference on Computer Vision*, pages 2564–2571, 2011.
- [33] Seyoon Jeong, Kyuheon Kim, Byungtae Chun, Jaeyeon Lee, and Y. J. Bae. An effective method for combining multiple features of image retrieval. In *IEEE Region 10 Conference. TENCON*, volume 2, pages 982–985 vol.2, 1999.
- [34] Sivic and Zisserman. Video google: a text retrieval approach to object matching in videos. In *International Conference on Computer Vision*, pages 1470–1477 vol.2, 2003.

- [35] Elena S. Stumm, Christopher Mei, and Simon Lacroix. Building location models for visual place recognition. *International Journal of Robotics Research*, 35(4):334–356, 2016.
- [36] N. Sünderhauf and P. Protzel. BRIEF-Gist closing the loop by simple means. In IEEE/RSJ International Conference on Intelligent Robots and Systems, pages 1234– 1241, 2011.
- [37] N. Sünderhauf, S. Shirazi, F. Dayoub, B. Upcroft, and M. Milford. On the performance of ConvNet features for place recognition. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 4297–4304, 2015.
- [38] K. A. Tsintotas, L. Bampis, and A. Gasteratos. Assigning visual words to places for loop closure detection. In *IEEE International Conference on Robotics and Automation*, pages 5979–5985, 2018.
- [39] K. A. Tsintotas, L. Bampis, and A. Gasteratos. Probabilistic appearance-based place recognition through bag of tracked words. *IEEE Robotics and Automation Letters*, 4 (2):1737–1744, 2019. ISSN 2377-3774.
- [40] Fukai Zhang, Ting Rui, Chengsong Yang, and Jianjun Shi. Lap-SLAM: A line-assisted point-based monocular VSLAM. *Electronics*, 8(2):243, 2019.
- [41] G. Zhang and I. H. Suh. Building a partial 3d line-based map using a monocular SLAM. In *IEEE International Conference on Robotics and Automation*, pages 1497– 1502, 2011.
- [42] Guangcong Zhang, Mason J Lilly, and Patricio A Vela. Learning binary features online from motion dynamics for incremental loop-closure detection and place recognition. In *IEEE International Conference on Robotics and Automation*, pages 765–772, 2016.
- [43] Lilian Zhang and Reinhard Koch. An efficient and robust line segment matching approach based on LBD descriptor and pairwise geometric consistency. *Journal of Visual Communication and Image Representation*, 24(7):794 805, 2013. ISSN 1047-3203.
- [44] H. Zhou, D. Zou, L. Pei, R. Ying, P. Liu, and W. Yu. StructSLAM: visual SLAM with building structure lines. *IEEE Transactions on Vehicular Technology*, 64(4):1364– 1375, 2015.
- [45] X. Zuo, X. Xie, Y. Liu, and G. Huang. Robust visual SLAM with point and line features. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1775–1782, 2017.