

A uniform distribution of landmarks for efficient map compression \star

Youssef Bouaziz, Eric Royer, Guillaume Bresson, Michel Dhome

▶ To cite this version:

Youssef Bouaziz, Eric Royer, Guillaume Bresson, Michel Dhome. A uniform distribution of landmarks for efficient map compression \star . ICVS, Lecture Notes in Computer Science, 14253, 2023, 978-3-031-44136-3. hal-04338699

HAL Id: hal-04338699 https://uca.hal.science/hal-04338699

Submitted on 12 Dec 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers. L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A uniform distribution of landmarks for efficient map compression*

Youssef Bouaziz^{1,2}, Eric Royer¹, Guillaume Bresson², and Michel Dhome¹

¹ Clermont Auvergne University, Institut Pascal, Clermont-Ferrand, France name.surname@uca.fr
² Institut VEDECOM, Versailles, France name.surname@vedecom.fr

Abstract. In this paper, we address the challenge of visual-based localization in dynamic outdoor environments characterized by continuous appearance changes. These changes greatly affect the visual information of the scene, resulting in significant performance degradation in visual localization. The issue arises from the difficulty of mapping data between the current image and the landmarks on the map due to environmental variations. One approach to tackle this problem is continuously adding new landmarks to the map to accommodate diverse environmental conditions. However, this leads to map growth, which in turn incurs high costs and resource demands for localization. To address this, we propose a map management approach based on an extension of the state-of-theart technique called Summary Maps. Our approach employs a scoring policy that assigns scores to landmarks based on their appearance in multiple localization sessions. Consequently, landmarks observed in multiple sessions are assigned higher scores. We demonstrate the necessity of maintaining landmark diversity throughout map compression to ensure reliable long-term localization. To evaluate our approach, we conducted experiments on a dataset comprising over 100 sequences encompassing various environmental conditions. The obtained results were compared with those of the state-of-the-art approach, showcasing the effectiveness and superiority of our proposed method.

Keywords: Visual-Based Navigation · Computer Vision for Transportation · Long-Term SLAM.

1 Introduction

Maps play a crucial role in self-driving applications, particularly in the context of high-precision localization using on-board cameras and reconstructed 3D points. However, the computational and memory limitations of mobile computing platforms pose significant challenges for real-time processing of these maps.

^{*} This work has been sponsored by the French government research program "Investissements d'Avenir" through the IMobS3 Laboratory of Excellence (ANR-10-LABX-16-01), by the European Union through the Regional Competitiveness and Employment program 2014-2020 (ERDF - AURA region) and by the AURA region.

Therefore, effective map management techniques are essential to optimize processing power and memory usage while maintaining localization accuracy.

While substantial progress has been made in the visual SLAM community for static environments or those with minimal changes, the challenge of localization in dynamic environments with varying conditions has only recently been addressed. Achieving reliable lifelong navigation in such dynamic environments poses a major hurdle for visual SLAM. In this paper, our focus is on realtime visual-based localization in outdoor environments for autonomous shuttles. These shuttles traverse the same path repeatedly but encounter diverse environmental conditions, leading to potential degradation in localization performance, even at familiar locations.

In such scenarios, environmental changes create significant difficulties in associating data between the current image and landmarks in the map. It is imperative for autonomous shuttles to adapt to such changes in order to ensure reliable long-term localization.

One approach to improve localization performance under changing conditions is to build a map that encompasses all environmental variations by continuously adding landmarks. However, this leads to continual map growth, directly proportional to the number of shuttle traversals. Consequently, localization after multiple traversals becomes infeasible due to the excessive memory requirements to store the map and the computational demands to match points between the current image and an extensive landmark database. In essence, achieving long-term real-time localization becomes unattainable after a certain number of localization sessions.

The objective of this paper is to propose a real-time solution that addresses the map growth issue for long-term localization. To this end, we introduce a map management strategy to reduce the map size offline. The proposed strategy focuses on removing redundant data from the map, significantly reducing the computational cost of the SLAM algorithm in long-term scenarios.

Several recent works have also explored the map management problem, including studies by [1], Mühlfellner *et al.* [14], Dymczyk et al. [7], Burki et al. [5], Krajník et al. [10], Halodová et al. [8], among others. Some of these approaches reduce map size based on a scoring policy that suggests removing landmarks with the lowest observation rate [14], [7]. Others eliminate landmarks associated with localization failures [5] or those with a high incorrect matching rate [8].

In this paper, we present an improvement to the Summary Maps approach proposed by Mühlfellner *et al.* [14]. Their method scores landmarks based on the number of different localization sessions in which they appear and removes landmarks with the lowest scores in an offline process. Our improved version of Summary Maps introduces a new constraint on landmark removal, ensuring a uniform number of landmarks per traversal after map summarization (i.e., compression).

We evaluate our enhanced approach using the publicly available IPLT (Institut Pascal Long-Term) dataset [3]. This dataset comprises 128 sequences recorded over a 16-month period, featuring a vehicle repeatedly following the same path around a parking lot with slight lateral and angular deviations. The IPLT dataset encompasses various environmental conditions, including changes in luminance, weather, seasons, and the presence of parked vehicles. Each sequence spans approximately 200 meters in length (see Fig. 1). We utilize this dataset to assess the performance of our approach and compare it against an existing state-of-the-art method.



Fig. 1. Example of 6 sequences recorded in a parking lot.

Our experiments demonstrate that our approach significantly improves localization performance, enabling successful localization even in challenging conditions.

2 Related work

In most real-world robotic scenarios, the ability to operate in dynamic and everchanging environments for extended periods is crucial. Simultaneous Localization and Mapping (SLAM) is a fundamental capability required in such scenarios. However, existing SLAM frameworks are often evaluated in static environments or scenes with minimal dynamic objects, such as moving people.

Recently, efforts have been made to extend the performance of localization in dynamic environments. Traditional feature-based comparison techniques in vision-based SLAM are considered unsuitable for long-term operations due to their vulnerability to changing conditions. As an alternative, Murillo and Kosecka proposed an image-based approach using the Gist representation of panoramic images to improve localization in dynamic environments [15]. Unlike local feature descriptors like SIFT and SURF, Gist is a global descriptor calculated using the entire image, representing an abstract scene representation. However, this approach involves an extensive search in the database to find the corresponding image, which becomes computationally expensive in large-scale environments. Milford and Wyeth introduced SeqSLAM, which enhances the performance of global image descriptors by matching sequences of images instead of individual images [13]. Although SeqSLAM demonstrates impressive results on various seasonal datasets, it remains sensitive to viewpoint changes.

Mühlfellner *et al.* proposed Summary Maps (SM) as a map management technique, where landmarks are ranked based on the number of different localization sessions in which they appear, and low-score landmarks are removed in an offline process [14]. However, SM has limitations in that landmarks in rarely visited areas receive low scores and are expected to be removed during map summarization. Dymczyk et al. addressed this bias towards frequently visited regions by designing a scoring policy that considers the expected number of trajectories to observe a landmark [7].

Krajník et al. developed a system that predicts the current state of the environment based on learned temporal patterns, constructing a new independent map in each run and integrating them into a spatio-temporal occupancy grid [10]. Bürki et al. improved appearance-based landmark selection by formulating a new ranking function [5]. They introduced two types of sessions, "rich sessions" and "observation sessions," assuming that the environmental condition of a session performing worse than a predefined threshold is not covered in the map. They added new landmarks to cover the new encountered environmental conditions and employed offline map summarization to produce a reliable map with a fixed size. Halodová et al. extended previous work by presenting an adaptive map update scheme that removes or adds features based on their past influence on localization quality [8]. They introduced an adaptive scoring policy that increments or decrements feature scores based on correct, incorrect, or unmatched matches. This strategy requires an accurate landmark retrieval technique to avoid penalizing incorrectly matched features caused by inaccurate retrievals. Bouaziz et al. proposed a keyframe retrieval technique that employs a ranking function considering factors such as Euclidean distance and underlying environmental conditions to search for keyframes with a higher number of inliers [2].

In this paper, we present an improvement to the Summary Maps (SM) approach proposed by Mühlfellner et al. [14]. Our enhanced version, named Uniform Summary Maps (USM), imposes a new constraint on landmark removal. The objective is to ensure a uniform number of landmarks in each traversal (sequence) after map summarization (compression).

3 Methodology

In this section, we introduce our map management approach, which extends the work proposed by Mühlfellner *et al.* [14] called Summary Maps (SM). We have re-implemented SM on our mapping framework.

Mühlfellner *et al.* [14] defined a landmark scoring policy that assigns scores to landmarks based on the number of different localization sessions in which they appear. Landmarks observed in multiple sessions receive high scores, indicating their value for localization. Conversely, landmarks with low scores are considered irrelevant. The Summary Maps approach utilizes these scores to summarize the map in an offline process conducted after each localization session, removing the least significant landmarks. Figure 2 illustrates the scoring function strategy proposed by Mühlfellner et al., as described in [14]. It presents an example of localization and showcases the landmarks observed in two different sessions, denoted as i and j.



Fig. 2. The landmark scoring policy proposed by Mühlfellner *et al.*. The figure illustrates a localization example where some landmarks $(l_{i,j})$ were observed in two different sessions (session *i* and session *j*) and accordingly, their scores were increased by the scoring policy.

An inherent limitation of this approach is that landmarks in rarely visited areas are assigned low scores since they are infrequently observed. Consequently, they are more likely to be removed during the map summarization process. Dymczyk *et al.* [7] referred to this issue as the bias towards regions that were more frequently visited and aimed to address this problem. They enhanced the scoring policy proposed in [14] by considering the expected number of trajectories required to observe a landmark.

However, we have identified another limitation of the Summary Maps approach, which is somewhat similar to the aforementioned bias. This limitation can be defined as a bias towards more experienced environmental conditions. It occurs when a set of sequences with similar environmental conditions is combined with a sequence having a distinct environmental condition in the same map. For example, a collection of daytime sequences incorporated into a map with one nighttime sequence. In such cases, landmarks observed under the odd environmental condition receive low scores from the scoring policy due to their infrequent observations compared to the landmarks from the other sequences. Consequently, these landmarks are filtered out during the map summarization step.

Table 1 presents an example of the bias towards more experienced environmental conditions that arises when summarizing a map containing 10 sequences using the Summary Maps approach. This map was constructed using the 10 sequences depicted in Figure 4 and includes only one nighttime sequence (the odd sequence). The table demonstrates the compression of the map using different compression ratios (1, 1.5, 2, 3, 5, and 10) with the Summary Maps approach. Compressing the map with a compression ratio r entails sorting all landmarks based on their scores and removing the lowest-scoring landmarks, representing 100 * (1 - 1/r) percent of the total (e.g., removing the 50% lowest scored landmarks for r = 2).

Table 1. Bias towards more experienced environmental conditions in Summary Maps.

Traversal	Compression ratio r							
	1 (no compression)	1.5	2	3	5	10		
1 🔅	140,524	113,922	100,371	75,087	50,304	20,817		
2 i	127,687	93,539	72,205	48,527	28,803	14,071		
3 🬧	149,065	83,821	52,025	34,912	18,913	9,324		
4 🬧	140,900	72,955	37,769	26,157	15,868	8,946		
5 🚢	122,122	97,989	86,896	55,751	29,495	14,022		
6 🜻	124,643	89,807	76,360	44,106	18,436	9,495		
7 🗳	72,044	41,333	28,777	16,310	5,092	2,709		
8 🔆	116,091	67,204	44,438	32,717	24,424	12,825		
9 🍊	127,972	96,797	78,262	51,610	34,322	16,692		
10 🍊	143,640	85,758	55,241	36,385	27,280	17,567		
total	1,264,688	843,125	632,344	421,562	252,937	126,469		

The Table shows the number of landmarks observed in each traversal after compressing the map using the Summary Maps approach at different compression ratios. The second column (r = 1) represents the initial map with no compression (100 * (1 - 1/1) = 0%). The last row represents the total number of landmarks on each map.

According to the table, a significant number of landmarks from the 7^{th} traversal (corresponding to the nighttime sequence) were removed after compressing the map. This demonstrates that the Summary Maps approach excludes landmarks observed in rarely experienced environmental conditions. This can pose a serious problem when localizing using nighttime sequences on the compressed map, as most of the nighttime landmarks have been filtered.

Therefore, in this paper, we propose an improvement for the Summary Maps technique called Uniform Summary Maps (USM). The goal of USM is to ensure a uniform distribution of landmarks across all traversals after map summarization. To achieve this, we introduce a constraint that imposes the distribution of a uniform number of landmarks across traversals. The algorithmic procedure for implementing USM is explained in Algorithm 1.

To further illustrate the execution of the algorithm and the effectiveness of our improved version of Summary Maps, we provide an example and compare the results with the original Summary Maps approach. Figure 3 presents an execution example of our algorithm on a map with N = 4 traversals. Each traversal l is associated with a specific number of landmarks denoted by n_{land}^l . The objective is to remove a total of $n_{\text{tot}} = 260$ landmarks from the map while ensuring a uniform distribution across the traversals.

A uniform distribution of landmarks for efficient map compression

Algorithm 1 Uniform Summary Maps

1: Parameters:					
: The total number of landmarks to remove: $n_{\rm tot}$					
3: The number of traversals in the map: N					
4: Steps:					
5: repeat					
6: Compute the number of landmarks n_{land}^{l} observed on each traversal l , with					
$l \in [1, N]$					
7: Sort the n_{land}^{l} landmarks of each traversal l according to the scoring policy					
8: Compute the highest number of landmarks: $n_{\max} \leftarrow \max(\{n_{land}^1, \dots, n_{land}^N\})$					
9: Find the set of traversals $S = \{L_{\max}^1, \dots, L_{\max}^s\}$ having n_{\max} landmarks					
10: Find the traversal L_{smax} having the second-highest number of landmarks:					
$n_{ ext{smax}} \leftarrow \max(\{n_{land}^1, \dots, n_{land}^N\} \setminus \{n_{ ext{max}}\})$					
11: Compute $n_{\text{diff}} \leftarrow n_{\text{max}} - n_{\text{smax}}$					
12: Compute the number of landmarks to remove:					
$n_{\text{rem}} \leftarrow \min(n_{\text{diff}} * \mathcal{S} , n_{\text{tot}}) / *$ the function min is used to make sure that we					
do not remove more than $n_{ m tot}$ landmarks */					
13: for each traversal $L_{\max} \in S$ do					
14: Remove the lowest scored $\frac{n_{\text{rem}}}{ \mathcal{S} }$ landmarks from traversal L_{max}					
15: end for					
16: Update the total number of landmarks to remove: $n_{\text{tot}} \leftarrow n_{\text{tot}} - n_{\text{rem}}$					
17: until $n_{\text{tot}} \leq 0$					

In the first iteration, the algorithm computes the number of landmarks $n_{\rm rem}$ to remove from the traversal with the highest number of landmarks, which is the 4th traversal.

In the second iteration, it is observed that there are two traversals, S = 2, 4, with the highest number of landmarks. Consequently, the algorithm calculates the number of landmarks to be removed $n_{\rm rem}$, and removes $n_{\rm rem}/|S|$ landmarks from each traversal in S.

In the third iteration, the set S contains three traversals. The algorithm calculates the number of landmarks to be removed in this iteration as $n_{\text{diff}} \times |S| = 150$. However, this value exceeds the total number of landmarks to be removed, which is $n_{\text{tot}} = 120$. To ensure that we do not remove more landmarks than n_{tot} , we use the statement $\min(n_{\text{diff}} \times |S|, n_{\text{tot}})$.

Table 2 presents the results of compressing the map using our improved version of Summary Maps, while Table 1 displays the results obtained with the original Summary Maps approach.

The comparison in Table 2 reveals that after compressing the map with different compression ratios, the number of landmarks remains uniform across the different traversals. This demonstrates the effectiveness of our proposed technique in achieving a balanced distribution of landmarks while reducing the overall number of landmarks in the map.



Fig. 3. An example illustrating the execution mechanism of our proposed algorithm. The yellow colored rows designate n_{max} and the green colored ones designate n_{smax} .

4 Experiments and results

To evaluate the performance of our approach, we conducted experiments using the IPLT dataset [3]. The IPLT dataset is a publicly available dataset that provides a diverse range of sequences for benchmarking visual localization algorithms. The dataset consists of over 100 sequences, each approximately 200 meters in length, captured under various environmental conditions including day, night, dusk, rain, and overcast. In all sequences, the vehicle follows a consistent path within a parking lot, as shown in Figure 1.

Although we acknowledge the importance of evaluating our approach on other widely used datasets such as the Oxford RobotCar dataset [11] or NCLT dataset [6], these datasets do not offer a substantial number of sequences traversing the same path, which is crucial for testing the effectiveness of our work. Therefore, we cannot rely on these datasets for our evaluation.

The IPLT dataset used in our study was generated using recorded images from two grayscale cameras mounted on an experimental vehicle. Each camera has a field of view of 100° . In addition to the image data, there is also wheel odometry information for localization purposes. The dataset includes 103 sequences captured at different times of the day and under various weather conditions, such as rain, sun, and overcast. To introduce additional challenges, some sequences feature lateral and angular deviations. Out of these 103 sequences, we utilized 10 sequences to construct the global map, while the remaining 93 sequences were reserved for evaluation and analysis.

8

Traversal	Compression ratio r							
	1 (no compression)	1.5	2	3	5	10		
1 🔅	140,524	85,676	63,235	42,157	25,294	12,647		
$2 \not\approx$	127,687	$85,\!676$	63,235	42,157	25,294	12,647		
3 🌧	149,065	$85,\!676$	63,235	42,157	25,294	12,647		
4 👚	140,900	$85,\!676$	63,235	42,157	25,294	12,647		
5 🍊	122,122	$85,\!676$	63,235	42,157	25,294	12,647		
6 🛒	124,643	$85,\!676$	63,235	42,157	25,294	12,647		
7 📞	72,044	72,044	63,235	42,157	25,294	12,647		
8 🔅	116,091	$85,\!676$	63,235	42,157	25,294	12,647		
9 🍊	127,972	$85,\!676$	63,235	42,157	25,294	12,647		
10 🍊	143,640	$85,\!676$	63,235	42,157	25,294	$12,\!647$		
total	1,264,688	843,128	632,350	421,570	252,940	126,470		

Table 2. Compressing the map with the Uniform Summary Maps.

We evaluated our approach on the IPLT dataset using the created global map which incorporates 10 traversals. In Figure 4, we present an overview of images from the 10 mapping sequences.

We conducted the evaluation using two different compression ratios: r = 3 and r = 2, where r represents the amount of landmark removal from the map. For instance, for r = 3, we removed 2/3 of the landmarks from the map. To assess the performance, we measured the average number of inliers observed in each test sequence, as well as the average number of localization failures per kilometer, which served as evaluation criteria.

We determined that reliable localization could be achieved when at least 30 points were matched between the current image and the database. Below this threshold, we considered it a localization failure. We have set a conservative threshold based on a study conducted by Royer et al. [16] to ensure the safety of the autonomous shuttle. This threshold, which we have adopted, is considered conservative and helps maintain a high level of security during the localization process.

To evaluate the impact of map compression on localization performance, we compared the localization performance on the initial global map, denoted as M_0 , which consisted of the 10 traversals, with the performance on two compressed maps: M_{SM} generated using the Summary Maps approach [14], and M_{USM} generated using our improved approach. In Figure 5, we present the average number of inliers per image and the average number of localization failures per kilometer observed during re-localization on these maps. In our experiments, we use Harris corner detector [9] for extracting key-points which are matched with ZNCC — Zero-mean Normalized Cross-Correlation — computed on 11×11 pixel windows around each key-point. However, our method can still be applied in the same way using other descriptors.



Fig. 4. An overview of images from the mapping sequences taken with the front camera. For each sequence, we provide the acquisition date and represent the environmental condition using a small icon.

The test sequences were manually classified into five different classes: "sun," "overcast," "rain," "dusk," and "night," with the "global" class containing all 93 testing sequences.

The Summary Maps approach (M_{SM}) exhibits a notable weakness when it comes to reducing the map size. This weakness is particularly more relevant in night sequences. The limitation arises from the fact that the 10 mapping sequences depicted in Figure 4 comprise only a single night sequence. As a result, landmarks observed during the night traversal receive a lower score compared to others, making them more likely to be removed during the map summarization process employed by the Summary Maps approach. Consequently, the nocturnal localization performance is adversely affected. This effect is evident in both subfigures (a) and (b) of Figure 5, where the number of localization failures in the night class is significant (166.6 localization failures/km with r = 3 and 49 with r = 2).

Conversely, our proposed technique, referred to as USM, effectively addresses this issue by ensuring a balanced representation of landmarks in the map after compression. As depicted in the figure, USM successfully increases the overall number of inliers and reduces the number of localization failures per kilometer, particularly during nighttime scenarios. This improvement in localization performance demonstrates the efficacy of our approach.

4.1 Discussion

Comparing the performance of our approach with existing state-of-the-art techniques poses a significant challenge. The difficulty arises from a fundamental conceptual difference between our approach and other techniques, wherein each

11



Fig. 5. Localization performance comparison on M_0 , M_{SM} and M_{USM} . Each color refers to a map as indicated in the legend, and the boxes represent the mean +/the standard deviation of inliers or localization failures on all the sequences of the corresponding class. Sub-figures (a) and (b) represent the localization performance when choosing r = 3 and r = 2 respectively.

0 10.06 0.18

(b) r = 2

0.27 49 12

0.04 7.63 1.45

100 0

000

0.03

000

state-of-the-art method is designed and applied within a specific mapping framework that utilizes distinct feature representations [4], [12]. This discrepancy in feature representation makes direct comparisons challenging, especially when comparing approaches applied to different types of SLAM frameworks, such as filter-based SLAM versus keyframe-based SLAM.

The variations in feature representation introduce complexities in evaluating and benchmarking different approaches. Each mapping framework has its own set of assumptions, algorithmic choices, and performance metrics. As a result, the performance evaluation of one approach within a specific framework may not be directly comparable to the evaluation of another approach in a different framework.

While comparing our approach with specific state-of-the-art techniques may be challenging, it is still valuable to highlight the unique aspects and contributions of our approach within the context of our chosen mapping framework. By focusing on the performance and evaluation metrics relevant to our framework, we can assess the effectiveness and efficiency of our approach compared to baseline methods or previous versions of our own approach. Additionally, conducting extensive experiments and providing thorough quantitative and qualitative analyses can help demonstrate the strengths and limitations of our approach in a more comprehensive manner.

In future research, it would be beneficial to establish standardized evaluation frameworks that can facilitate fair comparisons across different mapping approaches and frameworks. Such frameworks could define common datasets, performance metrics, and evaluation protocols, allowing for more direct comparisons between approaches. By adopting standardized evaluation practices, researchers can collectively advance the field and facilitate better understanding and comparison of different SLAM techniques.

5 Conclusions

In this study, we have proposed an enhancement to a state-of-the-art map management approach that leverages landmark observation information. Our improved technique introduces refinements to the scoring policy, allowing for more effective compression of maps while maintaining long-term localization performance. In our experiments, we conducted a comprehensive evaluation of our approach in various environmental conditions.

By comparing our improved technique with the original state-of-the-art approach, we have observed substantial performance gains. Our approach consistently outperformed the baseline method in terms of localization accuracy and robustness, even when operating with compressed maps. This highlights the effectiveness of our proposed enhancements and their positive impact on the overall localization performance.

Looking ahead, our future research endeavors will focus on further refining the scoring policy of our approach. We aim to incorporate additional information, such as past localization successes and failures, into the compression process. By considering a broader range of factors, we expect to enhance the efficiency and effectiveness of our map compression technique, ultimately improving the overall localization capabilities in challenging real-world scenarios.

References

- Biber, P., Duckett, T., et al.: Dynamic maps for long-term operation of mobile service robots. In: Robotics: science and systems. pp. 17–24 (2005)
- Bouaziz, Y., Royer, E., Bresson, G., Dhome, M.: Keyframes retrieval for robust long-term visual localization in changing conditions. In: 2021 IEEE 19th World Symposium on Applied Machine Intelligence and Informatics (SAMI). pp. 000093– 000100. IEEE (2021)

- Bouaziz, Y., Royer, E., Bresson, G., Dhome, M.: Over two years of challenging environmental conditions for localization: The iplt dataset. In: Proceedings of the 18th International Conference on Informatics in Control, Automation and Robotics - Volume 1: ICINCO, pp. 383–387. INSTICC, SciTePress (2021). https://doi.org/10.5220/0010518303830387
- Bürki, M., Cadena, C., Gilitschenski, I., Siegwart, R., Nieto, J.: Appearance-based landmark selection for visual localization. Journal of Field Robotics 36(6), 1041– 1073 (2019)
- Bürki, M., Dymczyk, M., Gilitschenski, I., Cadena, C., Siegwart, R., Nieto, J.: Map management for efficient long-term visual localization in outdoor environments. In: 2018 IEEE Intelligent Vehicles Symposium (IV). pp. 682–688. IEEE (2018)
- Carlevaris-Bianco, N., Ushani, A.K., Eustice, R.M.: University of Michigan North Campus long-term vision and lidar dataset. International Journal of Robotics Research 35(9), 1023–1035 (2015)
- Dymczyk, M., Lynen, S., Cieslewski, T., Bosse, M., Siegwart, R., Furgale, P.: The gist of maps-summarizing experience for lifelong localization. In: 2015 IEEE International Conference on Robotics and Automation (ICRA). pp. 2767–2773. IEEE (2015)
- Halodová, L., Dvoráková, E., Majer, F., Vintr, T., Mozos, O.M., Dayoub, F., Krajník, T.: Predictive and adaptive maps for long-term visual navigation in changing environments. In: 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 7033–7039. IEEE (2019)
- Harris, C.G., Stephens, M., et al.: A combined corner and edge detector. In: Alvey vision conference. vol. 15, pp. 10–5244. Citeseer (1988)
- Krajník, T., Fentanes, J.P., Hanheide, M., Duckett, T.: Persistent localization and life-long mapping in changing environments using the frequency map enhancement. In: 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 4558–4563. IEEE (2016)
- Maddern, W., Pascoe, G., Linegar, C., Newman, P.: 1 year, 1000km: The oxford robotcar dataset. The International Journal of Robotics Research (IJRR) **36**(1), 3–15 (2017). https://doi.org/10.1177/0278364916679498, http://dx.doi.org/10.1177/0278364916679498
- Magnago, V., Palopoli, L., Passerone, R., Fontanelli, D., Macii, D.: Effective landmark placement for robot indoor localization with position uncertainty constraints. IEEE Transactions on Instrumentation and Measurement 68(11), 4443– 4455 (2019)
- Milford, M.J., Wyeth, G.F.: Seqslam: Visual route-based navigation for sunny summer days and stormy winter nights. In: 2012 IEEE International Conference on Robotics and Automation. pp. 1643–1649. IEEE (2012)
- Mühlfellner, P., Bürki, M., Bosse, M., Derendarz, W., Philippsen, R., Furgale, P.: Summary maps for lifelong visual localization. Journal of Field Robotics 33(5), 561–590 (2016)
- Murillo, A.C., Kosecka, J.: Experiments in place recognition using gist panoramas. In: 2009 IEEE 12th International Conference on Computer Vision Workshops, ICCV Workshops. pp. 2196–2203. IEEE (2009)
- 16. Royer, E., Marmoiton, F., Alizon, S., Ramadasan, D., Slade, M., Nizard, A., Dhome, M., Thuilot, B., Bonjean, F.: Lessons learned after more than 1000 km in an autonomous shuttle guided by vision. In: 2016 IEEE 19th International Conference on Intelligent Transportation Systems (ITSC). pp. 2248–2253. IEEE (2016)