

GV-Bench: Benchmarking Local Feature Matching for Geometric Verification of Long-term Loop Closure Detection

Jingwen Yu^{1,2}, Hanjing Ye², Jianhao Jiao³, Ping Tan¹ and Hong Zhang^{2*}

Abstract— Visual loop closure detection is an important module in visual simultaneous localization and mapping (SLAM), which associates current camera observation with previously visited places. Loop closures correct drifts in trajectory estimation to build a globally consistent map. However, a false loop closure can be fatal, so verification is required as an additional step to ensure robustness by rejecting the false positive loops. Geometric verification has been a well-acknowledged solution that leverages spatial clues provided by local feature matching to find true positives. Existing feature matching methods focus on homography and pose estimation in long-term visual localization, lacking references for geometric verification. To fill the gap, this paper proposes a unified benchmark targeting geometric verification of loop closure detection under long-term conditional variations. Furthermore, we evaluate six representative local feature matching methods (handcrafted and learning-based) under the benchmark, with in-depth analysis for limitations and future directions.

I. INTRODUCTION

The main objective of loop closure detection is to facilitate robust navigation for long-term and large-scale autonomy [1], where the robot’s operation may span different day periods, weather conditions, and seasons. As a result, it becomes increasingly difficult to associate two images under challenging conditional variations. Visual loop closure detection first extracts compact global descriptors from images [2]; then, the similarities are calculated, and results are retrieved via nearest neighbor searching. Extensive works [3], [4], [5], [6] focus on learning conditional-invariant global features under changing environments. A potential loop closure pair is formed by the current query image with a retrieved database image, as the retrieval stage in Fig. 2. In detecting such a potential pair, only the appearance similarity between the images is used without considering the geometric relation between images. However, the appearance of places might change dramatically in the long term due to various conditions. Meanwhile, the appearance might be aliasing due to repetitive texture and similar structures. Therefore, missing correct loop closures (i.e., low recall) and detecting wrong

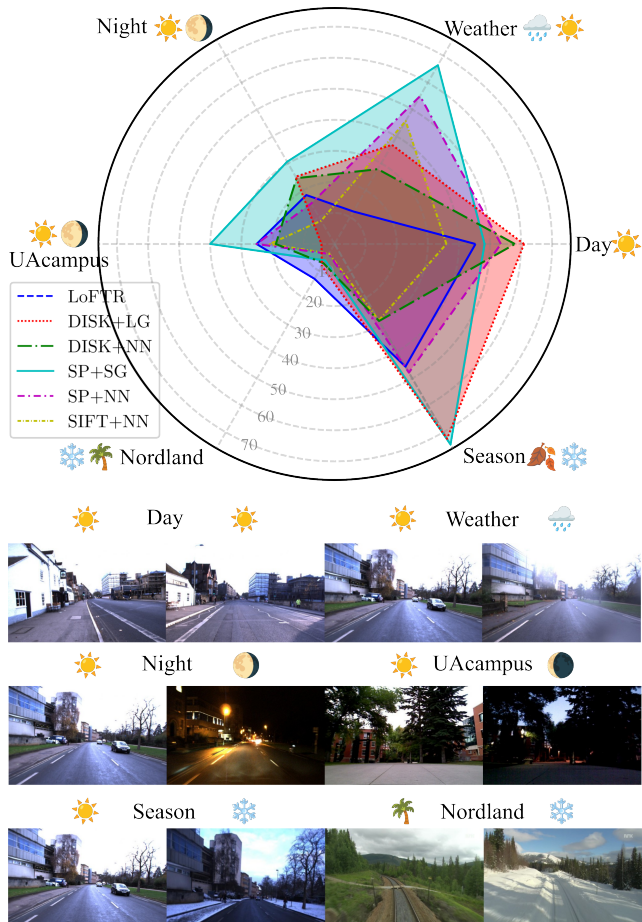


Fig. 1. **GV-Bench** enables evaluation of geometric verification over long-term loop closure detection. The benchmark contains six sequences representing different conditional variations. The benchmark is built upon three datasets [7], [8], [9]. We carefully select and compare six representatives of local feature matching, where SP+SG exhibits the best performance on average. However, the unfilled and unbalanced radar chart indicates promising improvement. The metric used in the chart is max recall @100 precision, which is explained in Sec. IV-A. The detailed construction of benchmark sequences is described in Sec. III-A.

loop closures (i.e., low precision) are the main challenges. To mitigate the second challenge, an additional step of geometric verification is employed to determine the correctness of each potential loop closure pair since false loop closures could lead to inaccurate localization.

Geometric verification refers to checking the spatial consistency of local feature correspondences between image pairs. As shown in Fig. 2, local features are first extracted and matched, followed by an outlier rejection method, such as RANSAC [10] by imposing epipolar constraint [11] between

*corresponding author: Hong Zhang hzhang@sustech.edu.cn

¹The Hong Kong University of Science and Technology, Hong Kong SAR, China. jingwen.yu@connect.ust.hk

²Shenzhen Key Laboratory of Robotics and Computer Vision, Southern University of Science and Technology, Shenzhen, China.

³Department of Computer Science, University College London, London, The United Kingdom.

The proposed benchmark is open-sourced at <https://github.com/jarvisyjw/GV-Bench>.

This work was supported by the Shenzhen Key Laboratory of Robotics and Computer Vision (ZDSYS20220330160557001) and the CETGC-ISA25EG01 project at The Hong Kong University of Science and Technology.

TABLE I. Six selected representatives of local feature matching. To simplify notation, we use the abbreviation below to represent each method.

Notation	Feature	Matcher
SIFT + NN	SIFT [13]	Nearest Neighbor
SP. + NN	SuperPoint [14]	Nearest Neighbor
SP. + SG.	SuperPoint [14]	SuperGlue [15]
DISK + NN	DISK [16]	Nearest Neighbor
DISK + LG.	DISK [16]	LightGlue [17]
LoFTR	—	LoFTR [18]

image pairs. The number of inliers or the inlier ratio is commonly employed as a binary classifier to validate a potential loop closure pair. A potential pair is classified as true if the number of inliers from local feature matching exceeds a pre-defined threshold. Therefore, conditional-invariant matches are the key to correctly verifying potential loop closure pairs. The long-term visual localization benchmark [12] provides a reference for choosing local features that are robust to appearance and viewpoint variations in pose estimation. However, there is no benchmark for evaluating their performance on geometric verification.

To fill the gap, we propose a unified benchmark for evaluating local feature matching methods under geometric verification for long-term loop closure detection. We aim to offer a reference for selecting a feature matching method for geometric verification in long-term autonomy. Extensive experiments are conducted to compare six representative local feature matching methods, summarized in Table I. We designed our benchmark with an easy-to-extend framework on top of three datasets [7], [8], [9] targeting long-term visual loop closure detection and place recognition. For a comparative study, we focus on three types of conditional variation: illumination, seasonal, and weather changes, as demonstrated in Fig. 1.

In summary, the main contributions of this paper are:

- **Fair and accessible geometric verification evaluation.** We open-source an out-of-box framework with a modular design as illustrated in Fig. 3, allowing for evaluating newly proposed methods on the common ground and extending to more diverse datasets.
- **A systematic analysis of geometric verification.** By employing the proposed benchmark, we point out possible future directions (e.g., training feature extractor and matcher with conditional variation data) through extensive experiments.

II. RELATED WORKS

A. Loop Closure Verification

Conventionally, geometric verification is a primary choice for verifying loop closure candidates [1]. Equivalently, it is also commonly adopted in visual localization systems, where it verifies the result from place recognition to help global localization [19]. In addition to the spatial constraints between images, a series of works attempts to leverage keypoints topology [20], scene graph [21], and semantics

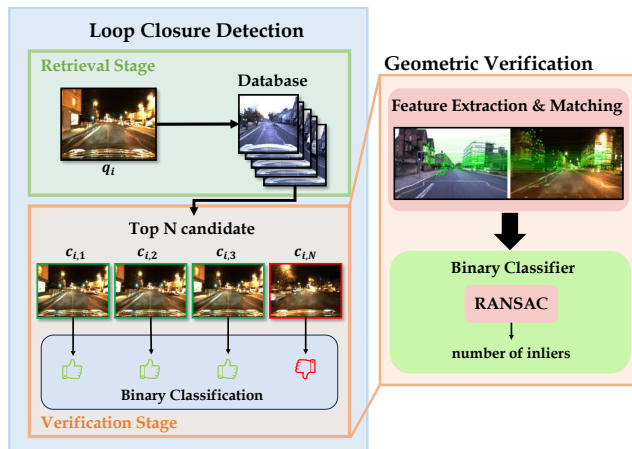


Fig. 2. **Loop Closure Detection** consists of two stages: retrieval and verification. Potential loop closure pairs $\{q_i, c_{i,j}\}$ detected by the retrieval stage are sent for verification. Each pair of images is examined under geometric constraints provided by local feature matching. RANSAC filters the matched correspondences to find the best inliers, which is used as the probability in binary classification.

[22]. In the 3D reconstruction, similar structures of buildings may cause mistakes in structure-from-motion. In a recent work [23], local feature matches combined with a convolutional neural network are employed as a binary classifier. However, it is limited by the domain gap and quality of image matching. Though many verification methods exist, a unified benchmark for evaluation on the common ground is missing. This paper proposes a benchmark to evaluate the geometric verification performance under long-term environmental changes, which are challenging for long-term robot autonomy. Besides, we specifically focus on evaluating different types of local feature matching methods under various conditions, aiming to reveal the existing challenges for future directions.

B. Local Feature Matching

Local feature matching is traditionally performed by i) detecting keypoints and computing keypoint descriptors, ii) matching descriptors using a nearest neighbor (NN) search, iii) filtering outliers by RANSAC [10]. SIFT [13], as a representative handcrafted feature, is widely used in object recognition, homography estimation, structure-from-motion, etc, for its scale and rotation invariance. With the advent of convolutional neural networks (CNNs), data-driven methods focus on learning better sparse detectors and local descriptors, such as SuperPoint [14] and DISK [16]. SuperPoint learns keypoint detection and description jointly by self-supervision. DISK proposes learning keypoint detection using a probabilistic approach via reinforcement learning.

Inspired by the success of learning-based local features, SuperGlue [15] achieves feature matching via a graph convolutional neural network (GCN), which performs particularly well with SuperPoint in homography and pose estimation. The following work, LightGlue [17], improves runtime performance and generalizes its usage to more feature keypoints and descriptors (including DISK). With the development

of transformers, LoFTR [18] predicts semi-dense matches and keypoints jointly, outperforming previous methods on pose estimation and claims to predict distinctive features under low-texture and repetitive patterns. In summary, we carefully select six representative local feature matching methods listed in Table I for evaluating geometric verification on the proposed benchmark.

III. EVALUATION METHODOLOGY

Here, we describe how a geometric verification method is typically constructed and how we evaluate different local feature matching algorithms. In Fig. 2, geometric verification deals with an image pair, computing the probability that the given pair is positive (i.e., representing the same spatial location). For each image pair, we first extract keypoints and compute correspondences from the feature matchers. These matches are further processed by computing the fundamental matrix [11] assuming a non-planar scene and applying RANSAC [10] to reject outliers. The detailed steps are provided in Section III-B. Moreover, a consistent list of image pairs should be used for fair comparison. In Section III-A, we describe how we generate a unified benchmark for evaluating geometric verification subject to conditional and viewpoint variations using three datasets [7], [8], [9]. In Fig. 3, we show the pipeline of benchmark construction, which is modular and easily expendable for more diverse conditional changes, local feature matching methods, and retrieval methods.

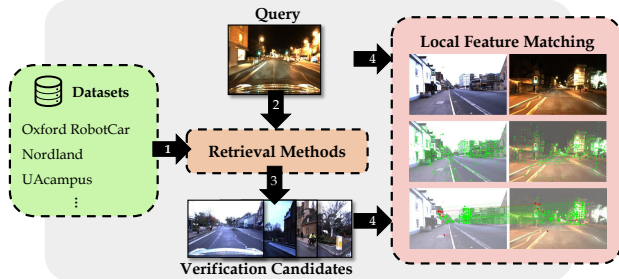


Fig. 3. **The pipeline of open-sourced benchmark** consists of i) pre-process dataset, ii) randomly select query set (if the dataset does not provide it), iii) retrieve verification candidates for each query, iv) match queries with candidates. The dashed modules (Datasets, Retrieval Methods, and Local Feature Matching) are expendable in the open-sourced framework, enabling easy customization for research purposes (i.e., enlarging sequences, using other retrieval methods, and evaluating new feature matching methods.)

A. Benchmark Introduction

Geometric verification examines a potential loop closure pair consisting of the current or query frame q_i with a database frame r_i . In loop closure detection, for the query frame q_i , we find the top N candidates from the database by sorting the similarity scores. In our benchmark, we choose N by the size of the image database. Our benchmark generates image pairs by mimicking the retrieval stage of the loop closure detection. In other words, for each query frame q_i , a list of potential loop closure pairs is generated as:

$$\{q_i, r_{i,j}\}_{j=1\dots N} \quad q_i \in \mathbf{Q}, r_{i,j} \in \mathbf{R}, N \in \mathbb{R} \quad (1)$$

where \mathbf{Q} is the query set and \mathbf{R} is the database set. We extract global descriptor \mathbf{V} by the popular method NetVLAD [3], which can be replaced by any retrieval method, thanks to the modular design illustrated in Fig 3.

Condition and viewpoint variations are the main challenges of long-term loop closure detection, where robot traversals may span over day and night, different weather conditions, and even seasons. However, it is hard to quantify or control any variation. To our best effort, we leverage the following datasets: Oxford RobotCar dataset [7], Nordland dataset [8], and UAcampus [9] to control the variables of variations. In general, we design the benchmark with three controlled variations: illumination (day and night), seasonal, and weather variations, as shown in Table II. In addition to conditional variations, we build a baseline sequence “Day”, for which only moderate environmental changes occur. Although this sequence is expected to be easily solvable, the results, as depicted in Fig. 5, are suboptimal, indicating that further improvement is possible.

1) *Oxford Robotcar Dataset [7]*: The data was collected by mounted cameras on a vehicle in Oxford, UK. The dataset includes over 100 repetitions of a consistent route spanning over a year. This collection covers various seasons, weather, and times of the day, offering various environmental conditions. We derive four sub-sequences: “Day”, “Night”, “Season”, and “Weather” from the original dataset as shown in Table II. We randomly choose 1720 images from a sequence¹ to form the query set \mathbf{Q} for all sub-sequences. Then, we retrieve the top 20 candidates from the other four sequences². The viewpoint changes are moderate since the images were captured on urban roads, as illustrated in Fig. 4(d) and 4(c). While conditional changes strongly over day and night. The ground truth labels for generated pairs are determined following the rule in [24], which selects positive pairs within a radius of 25 meters in translation and 40 degrees in orientation. Reference poses for images are provided by the RTK measurements. All the images are post-processed to avoid ambiguity³.

2) *Nordland Dataset [8]*: The images were collected by a mounted camera on the Nordland Railway in Norway across seasons. Following [25], we use images from winter as \mathbf{Q} and summer as \mathbf{R} . We extract the top 20 candidates for 1415 queries to generate 28300 potential pairs. The conditional variation is strong since the vegetation is covered in heavy snow, without artificial buildings as landmarks. However, the viewpoint change is approximated to none for a fixed camera posture on the train. If the reference image is within two frames relative to the query, it is selected as a positive pair.

3) *UAcampus Dataset [9]*: A forward-facing camera collected the images on a mobile robot teleoperated on a university campus covering the same route of 650 meters. The sequences are collected at different times (06:20 and

¹2014-12-09-13-21-02

²2014-11-18-13-20-12, 2014-12-16-18-44-24, 2015-02-03-08-45-10, 2014-11-25-09-18-32

³We undistort and remove the original images’ car hood (bottom 160 rows of pixels).

TABLE II. **Summary of sequences in proposed benchmark.** The benchmark consists of six sequences covering mainly three types of conditional changes: illumination (Night and UAcampus), seasonal (Season and Nordland), and weather changes in long-term loop closure detection. The “Day” sequence serves as the baseline challenge with moderate environmental changes over a short period.

Sequences	Traverse		Variations in	
	Query	Ref.	Condition	Viewpoint
Day [7]	Day	Day	moderate	moderate
Night [7]	Day	Night	very strong	moderate
Season [7]	Autumn	Winter	strong	moderate
Weather [7]	Sun	Rain	strong	moderate
Nordland [8]	Winter	Summer	very strong	moderate
UAcampus[9]	Day	Night	strong	strong

22:15) on the same day. We use the morning (06:20) as \mathbf{Q} and the night (22:15) as \mathbf{R} , generating the top 10 candidates for 647 query images. The conditional changes are very strong due to the illumination. The campus buildings are of similar appearance. In the meantime, the viewpoint variations are large. The ground truth is annotated manually.

B. Geometric Verification

Given an image pair, we extract keypoints and compute matches between them as demonstrated in Fig. 3. The matches are then transformed to point-wise correspondences denoted as \mathbf{x}_1 and \mathbf{x}_2 (in homogeneous coordinates), which are associated keypoint locations of the query and reference image, respectively.


$$\mathbf{x}_2^T \mathbf{F} \mathbf{x}_1 = 0 \quad (2)$$

We compute the fundamental matrix, denoted as \mathbf{F} , in conjunction with the RANdom SAMple Consensus (RANSAC) algorithm to solve for the best inliers \mathcal{I} . The fundamental matrix is a 3 by 3 matrix encapsulating the epipolar constraint between two views. It relates corresponding points in the two images through Equ. 2. The number of best inliers \mathcal{I} is used for binary classification presented in Fig. 2. For example, if the number of \mathcal{I} is larger than a given threshold, it is regarded as a positive pair.


Six representatives of local feature matching methods are carefully selected referring to long-term visual localization benchmarks [12], [26], and ‘Workshop on Long-Term Visual Localization under Changing Conditions⁴,’ where local features are benchmarked for pose estimation under changing conditions. The combination of SuperPoint [14] and SuperGlue [15] has been the lead for a long time till the detector-free local feature mather LoFTR [18] came out. DISK [16] is chosen for its competitive performance demonstrated in Image Matching Challenge [27]. In addition to the learning-based methods, SIFT [13] is chosen to represent the classical handcrafted features. As listed in Table I, for SuperPoint and DISK, we compare the nearest neighbor (i.e., representing the distinctiveness of feature descriptor) and learning-based matching network for a comprehensive study.

³<https://github.com/jarvisyiw/GV-Bench>

⁴<https://www.visuallocalization.net>



(a) Negative pair in “Nordland” with inliers: 9



(b) Positive pair in “Nordland” with inliers: 37



(c) Negative pair in “Day” with inliers: 45



(d) Positive pair in “Day” with inliers: 41

Fig. 4. **Example of local feature matching** under long-term challenging conditional changes. In the above figures, we visualize LoFTR matches with vanilla RANSAC. The inliers (green lines) and outliers (red lines) are highlighted. The number of inliers of Fig. 4(c) and 4(d) are counter-intuitive because RANSAC fails when false matches are dominant (more detailed analysis are provided in Sec. IV-D).

IV. EXPERIMENTAL RESULTS

A. Evaluation Metrics

Geometric verification can be regarded as a binary classification problem that determines whether a given potential loop closure pair is positive. Precision and recall are usually used for evaluation [1], where precision measures the proportion of correctly detected loop closures (true positive) out of all detected ones (true positives + false positives), and recall measures the proportion of correctly detected loop closures (true positives) out of all actual loop closures (true positives + false negatives). High precision indicates a low rate of false alarms, while high recall indicates the ability to detect most of the true loop closures. These two metrics are used jointly for the trade-off between false alarms and missed detections as a precision-recall curve in Fig. 5. Geometric verification should reject the false positive loop closure robustly. Therefore, we use two metrics for evaluation: maximum recall @100 precision (**MR**) [20] and

TABLE III. **Experimental results of all six sequences.** **Best**, **Second Best**, and **Third Best** results are highlighted. The last column presents the average performance of all six sequences. **MR** represents maximum recall @100 precision, **AP** stands for average precision. The results are round to three decimal places.

Matching Method	Day		Night		Season		Weather		Nordland		UAcampus		Average	
	AP	MR	AP	MR	AP	MR	AP	MR	AP	MR	AP	MR	AP	MR
SIFT + NN	98.1	36.0	62.1	8.7	98.9	28.4	99.7	45.7	13.6	3.5	53.5	20.9	71.0	23.9
SP. + NN	99.3	53.5	86.8	15.1	99.7	47.8	99.8	55.0	52.8	3.7	74.0	25.1	85.4	33.4
SP. + SG.	99.6	48.3	96.0	30.6	99.9	74.6	99.9	66.5	71.1	6.1	85.3	40.1	92.0	44.4
DISK + NN.	97.4	58.0	51.7	24.5	99.0	28.6	99.7	27.8	17.9	7.2	62.7	18.9	71.4	27.5
DISK + LG.	99.7	61.0	81.4	25.0	99.9	72.4	99.9	36.9	63.8	8.2	78.6	4.2	87.2	34.6
LoFTR	99.5	45.3	97.9	18.2	99.9	45.6	99.9	12.1	81.1	13.0	81.0	25.1	93.2	26.6




Fig. 5. **Precision-recall curve** of the “Day” sequence. The marker annotates the maximum recall @100 precision (MR).

average precision (AP). The MR represents the highest recall while keeping the precision to 100%, representing the ability to find true loop closures without false positives.

B. Comparison over Conditional Variations

The evaluation results of the proposed benchmark are presented in Table III. On average, the combination of SuperPoint and SuperGlue achieves the highest MR of 44.4%, while LoFTR reaches the highest AP of 93.2%, indicating the highest AP does not lead to a higher MR. The re-rank methods, like [28] and [29] in the computer vision community, focus on evaluating the re-rank performance over AP (AUC), which does not provide a reference for geometric verification by far. Although no method consistently ranks top on the benchmark under different challenging conditions, SP.+SG. achieves the highest MR on four sequences out of six, suggesting a decent choice for geometric verification. LoFTR and DISK+LG. take up the remaining two first places, which are also competitive. DISK+LG. gets second-best MR on average; however, its poor performance on “UAcampus” indicates it might be unreliable. Learning-based sparse feature extractors SuperPoint and DISK outperform SIFT by nearest neighbor matching, demonstrating the distinctiveness of learning-based descriptors under changing conditions. This is consistent with previous evaluations on

homography [14] and pose estimation [16].

All methods suffer from perceptual aliasing. As depicted in Fig. 4(a) and 4(b), in “Nordland” sequences, the surrounding is unstructured vegetation. Even for loop closure pairs, the number of inlier matches is low, making it hard to distinguish from the negative pairs. Fig. 6 demonstrated that the frequency density for LoFTR matches is more “distributed,” i.e., the average number of inliers of loop closures is higher than non-loop closures. The winter images are similar to a low-texture environment where the heavy white snow covers the rich texture of the environment. Our results are consistent with [30] that LoFTR works especially well under texture-less situations.




Fig. 6. **The number of inliers’ distribution** of the sequence “Nordland” of SP. + SG. (Top) and LoFTR (Bottom). The density of frequency indicates the distribution of the image pairs with the corresponding normalized inliers count. The ideal distribution is that there is no overlap between non-loop-closure pairs and loop-closure pairs. However, as figures above shows, it is rare in practice. From the plot of inliers’ distribution, we can observe which kind of samples are in dominant.

C. Comparison over Runtime

For robotics applications, time efficiency is critical on resource-constraint platforms. We measure the runtime of six

methods listed in Table I on NVIDIA GeForce RTX 3090 GPU and Intel i7-13700K CPU over 10K runs. The results are shown in Fig. 7 as inference time over performance, i.e., max recall @100 precision. We can conclude that the runtime of six local feature matching methods is at a millisecond level on a modern GPU. The choice can be made based on the trade-off between time efficiency and performance.




Fig. 7. **The efficiency vs. efficacy** of the six feature matching methods. Y-axis represents for performance in max recall @100 precision, x-axis is the inference speed in FPS.

D. Discussion

Based on the experiments, we conclude three key take-aways for the potential improvement of geometric verification: improving local feature extraction and matching, building a multi-condition image database, and enhancing outlier rejection.

1) **Leveraging conditional variation data:** Although feature extraction and matching methods like SP.+SG. and LoFTR perform well in long-term visual localization [26], they do not work very well in our proposed benchmark. One of the possible solutions is to learn more distinctive descriptors under various conditions. SuperPoint is trained on synthetic data and MS-COCO [31], while LoFTR is trained from MegaDepth [32]. These datasets are not designed for long-term visual tasks and do not cover challenging conditional variations presented in our proposed benchmark. A natural idea is to expand the diversity of the training set, i.e., training the feature detector and matcher with conditional variation data. A recent work GIM [33], shows the possibility of enhancing the robustness against long-term variations via learning generalizable image matching from large amounts of data, such as internet videos.

Moreover, in the training of feature matchers [15], [17], [18], the goal is to match two paired images for homography or pose estimation. Therefore, a visual overlap is assumed, which is not hypothetical in geometric verification. Unpaired image samples should be considered in the training of feature matchers for robust geometric verification.

2) **Building a multi-condition image database:** Reducing the appearance gap between query and reference set may potentially mitigate the issue of perceptual aliasing illustrated in Fig. 4(a) and 4(b). According to a robotic visual localization system across long-term operations [34], multiple experiences (i.e., image sequences over different conditions) are

stored in the database. It achieves robust localization under significant appearance variation by searching and matching the latest sequences. As an inspiration, storing multiple sequences based on conditional changes (e.g., illumination, season, and weather changes) might alleviate the perceptual aliasing due to appearance gap over a long period.

3) **Exploiting more powerful outlier rejection:** The vanilla RANSAC fails to work for scenarios where false matches dominate. As illustrated in Fig. 4(c), in an urban environment with a similar structure, the incorrect matches are uniformly distributed, leading to a “false fundamental matrix”. This happens to all the sequences regardless of conditional variation. RANSAC works under epipolar constraints only, which does not hold for two unpaired images. However, more powerful outlier rejectors [35], [36] outperform in fundamental matrix estimation, promising possible improvement in geometric verification.

V. CONCLUSIONS AND FUTURE WORK

This paper builds a unified benchmark for evaluating geometric verification of long-term visual loop closure detection. Specifically, we focus on conducting a comparative study of six carefully selected feature matching methods, covering handcrafted and learning-based features. Our benchmark is open-sourced with a modular design to support future research. This study allows us to draw the following conclusions. i) The combination of sparse keypoint detector SuperPoint [14] and feature matcher SuperGlue [15] in general outperforms others. ii) When subjected to challenging conditions, the learning-based keypoint descriptor is more distinctive than the handcrafted one. This result is consistent with previous evaluations on homography and pose estimation [14], [16]. iii) We conclude that the verification of loop closure detection is far from being solved. Further discussion suggests possible future directions, including training keypoint detectors with conditional variation data, expanding the dataset with images collected under different conditions, and exploring more powerful outlier rejection methods.

For future work, i) we will expand the benchmark to support other loop closure verification methods mentioned in Sec. II-A. Although we designed this benchmark to focus on geometric verification, it naturally supports evaluating other verification methods. By this expansion, we can comprehensively understand and explore beyond geometric verification. ii) we will focus on the directions proposed in Sec. IV-D, aiming at developing a novel and effective method of geometric verification.

ACKNOWLEDGMENT

The authors appreciate the discussion with Mr. Guangcheng Chen (Southern University of Science and Technology), Mr. Zhonghang Liu (Singapore Management University), Dr. Hengli Wang (Huawei Technologies), and Dr. Yuxuan Liu (TIER IV).

REFERENCES

- [1] K. A. Tsintotas, L. Bampis, and A. Gasteratos, "The revisiting problem in simultaneous localization and mapping: A survey on visual loop closure detection," *IEEE Trans. Intell. Transp. Syst.*, vol. 23, no. 11, pp. 19 929–19 953, 2022.
- [2] Y. Liu and H. Zhang, "Visual loop closure detection with a compact image descriptor," in *Proc. IEEE/RJS Int. Conf. Intell. Robots Syst.*, 2012, pp. 1051–1056.
- [3] R. Arandjelovic, P. Gronat, A. Torii, T. Pajdla, and J. Sivic, "Netvlad: Cnn architecture for weakly supervised place recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 5297–5307.
- [4] G. Berton, C. Masone, and B. Caputo, "Rethinking visual geolocalization for large-scale applications," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 4878–4888.
- [5] N. Keetha, A. Mishra, J. Karhade, K. M. Jatavallabhula, S. Scherer, M. Krishna, and S. Garg, "Anyloc: Towards universal visual place recognition," *IEEE Robot. Autom. Lett.*, 2023.
- [6] H. Ye, W. Chen, J. Yu, L. He, Y. Guan, and H. Zhang, "Condition-invariant and compact visual place description by convolutional autoencoder," *Robotica*, vol. 41, no. 6, pp. 1718–1732, 2023.
- [7] W. Maddern, G. Pascoe, C. Linegar, and P. Newman, "1 year, 1000 km: The oxford robotcar dataset," *Int. J. Robot. Res. (IJRR)*, vol. 36, no. 1, pp. 3–15, 2017.
- [8] S. Skrede, "Nordland dataset," <https://bit.ly/2QVBOym>, 2013.
- [9] Y. Liu, R. Feng, and H. Zhang, "Keypoint matching by outlier pruning with consensus constraint," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2015, pp. 5481–5486.
- [10] M. A. Fischler and R. C. Bolles, "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography," *Communications of the ACM*, vol. 24, no. 6, pp. 381–395, 1981.
- [11] R. Hartley and A. Zisserman, *Multiple view geometry in computer vision*. Cambridge university press, 2003.
- [12] C. Toft, W. Maddern, A. Torii, L. Hammarstrand, E. Stenborg, D. Safari, M. Okutomi, M. Pollefeys, J. Sivic, T. Pajdla, et al., "Long-term visual localization revisited," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 4, pp. 2074–2088, 2020.
- [13] D. G. Lowe, "Object recognition from local scale-invariant features," in *Proc. IEEE Int. Conf. Comput. Vis.*, vol. 2, 1999, pp. 1150–1157.
- [14] D. DeTone, T. Malisiewicz, and A. Rabinovich, "Superpoint: Self-supervised interest point detection and description," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshop*, 2018, pp. 224–236.
- [15] P.-E. Sarlin, D. DeTone, T. Malisiewicz, and A. Rabinovich, "Super-glue: Learning feature matching with graph neural networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 4938–4947.
- [16] M. Tyszkiewicz, P. Fua, and E. Trulls, "Disk: Learning local features with policy gradient," *Adv. Neural Inf. Process. Syst.*, vol. 33, pp. 14 254–14 265, 2020.
- [17] P. Lindenberger, P.-E. Sarlin, and M. Pollefeys, "Lightglue: Local feature matching at light speed," in *Proc. IEEE Int. Conf. Comput. Vis.*, October 2023, pp. 17 627–17 638.
- [18] J. Sun, Z. Shen, Y. Wang, H. Bao, and X. Zhou, "Loftr: Detector-free local feature matching with transformers," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 8922–8931.
- [19] P.-E. Sarlin, C. Cadena, R. Siegwart, and M. Dymczyk, "From coarse to fine: Robust hierarchical localization at large scale," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019.
- [20] H. Yue, J. Miao, Y. Yu, W. Chen, and C. Wen, "Robust loop closure detection based on bag of superpoints and graph verification," in *Proc. IEEE/RJS Int. Conf. Intell. Robots Syst.*, 2019, pp. 3787–3793.
- [21] N. Hughes, Y. Chang, and L. Carlone, "Hydra: A real-time spatial perception system for 3d scene graph construction and optimization," *arXiv preprint arXiv:2201.13360*, 2022.
- [22] S. Orhan, J. J. Guerrero, and Y. Baştanlar, "Semantic pose verification for outdoor visual localization with self-supervised contrastive learning," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshop*, 2022, pp. 3989–3998.
- [23] R. Cai, J. Tung, Q. Wang, H. Averbuch-Elor, B. Hariharan, and N. Snavely, "Doppelgangers: Learning to disambiguate images of similar structures," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2023, pp. 34–44.
- [24] F. Warburg, S. Hauberg, M. Lopez-Antequera, P. Gargallo, Y. Kuang, and J. Civera, "Mapillary street-level sequences: A dataset for lifelong place recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 2626–2635.
- [25] L. G. Camara and L. Přeučil, "Visual place recognition by spatial matching of high-level cnn features," *Robot. and Auton. Syst.*, vol. 133, p. 103625, 2020.
- [26] T. Sattler, W. Maddern, C. Toft, A. Torii, L. Hammarstrand, E. Stenborg, D. Safari, M. Okutomi, M. Pollefeys, J. Sivic, F. Kahl, and T. Pajdla, "Benchmarking 6dof outdoor visual localization in changing conditions," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018.
- [27] Y. Jin, D. Mishkin, A. Mishchuk, J. Matas, P. Fua, K. M. Yi, and E. Trulls, "Image Matching across Wide Baselines: From Paper to Practice," *Int. J. Comput. Vis.*, 2020.
- [28] F. Tan, J. Yuan, and V. Ordonez, "Instance-level image retrieval using reranking transformers," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2021, pp. 12 105–12 115.
- [29] G. Barbarani, M. Mostafa, H. Bayramov, G. Trivigno, G. Berton, C. Masone, and B. Caputo, "Are local features all you need for cross-domain visual place recognition?" in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 6154–6164.
- [30] X. He, J. Sun, Y. Wang, S. Peng, Q. Huang, H. Bao, and X. Zhou, "Detector-free structure from motion," *arXiv preprint arXiv:2306.15669*, 2023.
- [31] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft coco: Common objects in context," in *Proc. Eur. Conf. Comput. Vis.*, 2014, pp. 740–755.
- [32] Z. Li and N. Snavely, "Megadepth: Learning single-view depth prediction from internet photos," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018.
- [33] X. Shen, Z. Cai, W. Yin, M. Müller, Z. Li, K. Wang, X. Chen, and C. Wang, "Gim: Learning generalizable image matcher from internet videos," in *The Twelfth Int. Conf. on Learn. Representations*, 2024.
- [34] M. Paton, K. MacTavish, M. Warren, and T. D. Barfoot, "Bridging the appearance gap: Multi-experience localization for long-term visual teach and repeat," in *Proc. IEEE/RJS Int. Conf. Intell. Robots Syst.*, 2016, pp. 1918–1925.
- [35] T. Wei, Y. Patel, A. Shekhovtsov, J. Matas, and D. Barath, "Generalized differentiable ransac," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2023, pp. 17 649–17 660.
- [36] D. Barath, J. Noskova, M. Ivaschekin, and J. Matas, "Magsac++, a fast, reliable and accurate robust estimator," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 1304–1312.