

Good Deep Features to Track: Self-Supervised Feature Extraction and Tracking in Visual Odometry

Sai Puneeth Gottam¹, Haoming Zhang², *Member, IEEE* and Eivydas Keras²

Abstract—Visual-based localization has made significant progress, yet its performance often drops in large-scale, outdoor, and long-term settings due to factors like lighting changes, dynamic scenes, and low-texture areas. These challenges degrade feature extraction and tracking, which are critical for accurate motion estimation. While learning-based methods such as SuperPoint and SuperGlue show improved feature coverage and robustness, they still face generalization issues with out-of-distribution data. We address this by enhancing deep feature extraction and tracking through self-supervised learning with task-specific feedback. Our method promotes stable and informative features, improving generalization and reliability in challenging environments.

Index Terms—Visual Odometry, Self-Supervised Learning

I. INTRODUCTION

Visual-based robot localization has achieved remarkable success with a wide range of approaches tailored to specific applications. However, the performance of state-of-the-art approaches often degrades in large-scale and outdoor environments during a long-term operation, where a variety of challenges arise, such as dynamic lighting conditions and objects, low-parallax features and textureless surfaces.

These factors dramatically impact the quality and robustness of one of the most crucial steps in visual-based methods: feature extraction and tracking. Fig. 1 exemplifies this issue in a maritime environment, where illumination and scene appearance can vary frequently even at the same location.

Generally, a well-distributed coverage of tracked features across the entire image is expected to provide better geometric constraints. Assuming that the scale information is given, distant features help constrain rotation, while nearby features are essential for accurately resolving translational motion due to their stronger parallax [1].

Keeping this in mind, classical feature extractors such as SIFT and GFTT tend to only focus on texture-rich and well-lit regions, resulting in densely clustered features and poor coverage across the image (see Fig. 1b and Fig. 1c).


In contrast, deep learning-based methods such as SuperPoint [2] produce more uniformly distributed results, even in low-texture or challenging lighting conditions (see Fig. 1d). However, since pre-trained model often degrade during inference when exposed to out-of-distribution data, obtaining high-quality features that can be robustly tracked to ensure accurate motion estimation in diverse real-world environments remains an open research question.

¹S. Gottam is with the Chair of Cyber-Physical-Systems, Montanuniversität Leoben, Leoben, Austria.


²H. Zhang and E. Keras are with the Institute of Automatic Control, Faculty of Mechanical Engineering, RWTH Aachen University, Aachen, Germany.
Corresponding: haoming.zhang@rwth-aachen.de




(a) Original image.



(b) Feature extraction using SIFT [3].



(c) Feature extraction using GFTT [4].



(d) Feature extraction using pre-trained SuperPoint [2].

Figure 1: Extracted features (yellow) of visual odometry with different state-of-the-art approaches using the Pohang dataset (Pohang 04) [5].




Figure 2: Proposed self-supervised pipeline to train the SuperPoint [2] and SuperGlue [9] models within a well-established visual SLAM framework RTAB-Map [11].

Although many methods exist to improve feature quantity — such as sophisticated feature matching and filtering [1], optimizing camera attributes (e.g., exposure time) [6], and removing dynamic objects [7] using semantic information [8] — they still struggle to generalize reliably for long-term motion estimation in diverse outdoor environments.

Moreover, we observe that the distinction between “good” and “poor” visual features is not always clear-cut, as it depends not only on geometric properties but also on factors such as scale and temporal consistency. In visually degraded scenarios, features previously considered “poor” may still be trackable over short time windows and, in some cases, are sufficient to prevent visual odometry (VO) failure.

Given these observations, we aim to improve the performance and generalization of the deep learning-based feature extractor and tracker by leveraging self-supervised learning with well-designed loss terms that provide situational feedback on what constitutes “good deep features to track”. We chose SuperPoint [2] and SuperGlue [9] as our backbone networks compared to other discriminative models, as they are capable of capturing structurally and semantically meaningful features thanks to their adaptive (e.g., homographic adaptation) training strategy.

Unlike [10], which treats the VO as a black box and only tunes hyperparameters (e.g., grid size) from a system-level perspective, we argue that it is (more) important to focus on the core components of a VO.

We build on these foundations by directly modifying and fine-tuning the core components of feature extraction and matching in a self-supervised way. Our work integrates recent datasets like Pohang [5] and builds upon proven frameworks such as RTAB-Map [11].

Our hypothesis is supported by preliminary results, which show that the proposed method can effectively update the model parameters and force the models to ignore unstable or uninformative features.

II. METHOD

We build on the RTAB-Map framework by replacing the feature extractor and matcher with SuperPoint [2] and SuperGlue [9].

Fig.2 illustrates the proposed self-supervised learning pipeline in our work. We employed RTAB-Map [11] as the

baseline visual odometry, where the feature extractor and tracker are replaced by SuperPoint [2] and SuperGlue [9].

During training, estimated relative pose, stereo and temporal correspondences are used to evaluate reprojection errors, which serve as geometric consistency cues. Features that are well-tracked and geometrically reliable are identified as “good” and used to compute a total weighted loss:

$$L = \underbrace{w_i \cdot L_i(\mathbf{X}, \mathbf{Y}) + w'_i \cdot L_i(\mathbf{X}', \mathbf{Y}') + w_{pk} \cdot L_{pk}}_{\text{training detector head}} + \underbrace{w_d \cdot L_d(\mathcal{D}, \mathcal{D}', \mathcal{S})}_{\text{training descriptor head}} \quad (1)$$

where the loss terms $L_i(\mathbf{X}, \mathbf{Y})$ and $L_i(\mathbf{X}', \mathbf{Y}')$ penalize non-keypoint score map $\mathbf{X}$ with their labels $\mathbf{Y}$ from image $[\mathbf{X}, \mathbf{Y}] \in \mathcal{I}$, and the corresponding warped image $[\mathbf{X}', \mathbf{Y}] \in \mathcal{I}'$. We use a peaky loss L_{pk} to penalize inconsistent or diffuse feature responses in the score maps, which are often caused by blur or ambiguous textures. In addition, a hinge loss L_d is employed to enforce descriptor discriminability by maximizing the distance between non-matching features and minimizing it for matching pairs.

III. PRELIMINARY RESULTS AND DISCUSSION

A. Feature Extraction

Fig. 4 and Fig. 3 present qualitative results of feature extraction using the pre-trained SuperPoint model and the proposed self-supervised version.

Compared to the former, features extracted by the self-supervised model tend to be more equitably spread across expected regions in the image, such as rigid structures, while suppressing features from less reliable areas like terrain and water surfaces.

B. Trajectory Estimation

To validate the proposed method, we use both the pre-trained and self-supervised SuperPoint models as feature extractors within the RTAB-Map framework to estimate the vessel’s trajectory. In addition, we employ the trajectory estimation using ORB features [13] as a baseline method for comparison.



Figure 3: Extracted visual features (purple and green) of Pohang data [5] in different scenes and datetime using original pre-trained (top row) and self-supervised (bottom row) SuperPoint model.



(a) Features from pre-trained SuperPoint.




(b) Features from self-supervised SuperPoint.

Figure 4: Extracted visual features (red) of KITTI data (seq. 06) [12].


We evaluate all methods on the Pohang dataset [5], which is divided into sequences recorded in narrow and wide water canals, during both daytime and nighttime (see Fig. 4). The results are presented in Fig.5.

In narrow areas (Fig.4a and 4c) where most features can be acquired from infrastructure objects during daytime, the proposed model outperforms the others. A similar trend is observed during nighttime, although all VO methods experience degradation due to insufficient visual features.


The problem of insufficient visual features becomes more pronounced in near-offshore areas. Surprisingly, the estimated trajectories using SuperPoint features exhibit less drift during nighttime compared to daytime, which stands in contrast to the behavior observed with ORB features. This may be due to fewer visual distractions at night, such as reflections and moving elements, which are more common during the day. In such cases, classic feature extractors, which are unable to uniformly explore the entire image, may occasionally avoid distracting regions.




(a) Trajectories in narrow canal during daytime.



(b) Trajectories in narrow canal during nighttime.



(c) Trajectories in off-shore area during daytime.



(d) Trajectories in off-shore area during nighttime.


Figure 5: Estimated trajectories of both narrow canal and offshore areas in Pohang [5] using different feature extractors implemented in RTAB-Map [11].

IV. CONCLUSION AND FUTURE WORK


This work aims to enhance the generalization of deep learning-based feature extraction and tracking within a self-supervised training framework. In the early stage of this research, we successfully demonstrate a proof of concept, showing that the model can be effectively improved using simple loss terms—without relying on label-intensive training procedures.

However, we did not train the feature matcher (e.g., Super-Glue), which remains pre-trained on the original SuperPoint. Our current results are also **not yet sufficient to fully answer the question of what defines “good deep features” to track.**

Additionally, since vision-based localization faces more challenges in outdoor environments than modalities like GNSS or LiDAR, we aim to explore **how VO can be better designed and integrated into a multisensory state estimator for resilient, long-term robot navigation** (see Fig. 6).



(a) gnssFGO without visual odometry.



(b) gnssFGO with visual odometry.

Figure 6: Demonstration of robot localization using gnssFGO [14]. The green, white, and blue trajectories represent the baseline, GNSS standalone, and gnssFGO estimates.

ACKNOWLEDGMENT

The presented research was supported by the German Federal Ministry of Economic Affairs and Climate Action (BMWK) under Project 19F1150B (Museas). The work was conducted at the Institute of Automatic Control (IRT), RWTH Aachen University. The authors also thank Space Applications Services NV/SA in Belgium for their support for providing hardware support during this work.

REFERENCES

- [1] I. Cvišić, I. Marković, and I. Petrović, “SOFT2: Stereo visual odometry for road vehicles based on a point-to-epipolar-line metric,” *IEEE Transactions on Robotics*, vol. 39, no. 1, pp. 273–288, 2023.
- [2] D. DeTone, T. Malisiewicz, and A. Rabinovich, “SuperPoint: Self-supervised interest point detection and description,” in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2018, pp. 337–33712.
- [3] D. G. Lowe, “Distinctive image features from scale-invariant keypoints,” *Int. J. Comput. Vision*, vol. 60, no. 2, pp. 91–110, Nov. 2004.
- [4] J. Shi and Tomasi, “Good features to track,” in *1994 Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 1994, pp. 593–600.
- [5] D. Chung, J. Kim, C. Lee, and J. Kim, “Pohang canal dataset: A multimodal maritime dataset for autonomous navigation in restricted waters,” *The International Journal of Robotics Research*, vol. 42, no. 12, pp. 1104–1114, 2023.
- [6] J. Kim, Y. Cho, and A. Kim, “Proactive camera attribute control using Bayesian optimization for illumination-resilient visual navigation,” *IEEE Transactions on Robotics*, vol. 36, no. 4, pp. 1256–1271, 2020.
- [7] B. Kitt, F. Moosmann, and C. Stiller, “Moving on to dynamic environments: Visual odometry using feature classification,” in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2010, pp. 5551–5556.
- [8] S. Wen, X. Li, X. Liu, J. Li, S. Tao, Y. Long, and T. Qiu, “Dynamic SLAM: A visual SLAM in outdoor dynamic scenes,” *IEEE Transactions on Instrumentation and Measurement*, vol. 72, pp. 1–11, 2023.
- [9] P.-E. Sarlin, D. DeTone, T. Malisiewicz, and A. Rabinovich, “Super-Glue: Learning feature matching with graph neural networks,” in *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020, pp. 4937–4946.
- [10] N. Messikommer, G. Cioffi, M. Gehrig, and D. Scaramuzza, “Reinforcement learning meets visual odometry,” in *Computer Vision – ECCV 2024: 18th European Conference, Milan, Italy, September 29–October 4, 2024, Proceedings, Part LIX*. Berlin, Heidelberg: Springer-Verlag, 2024, p. 76–92.
- [11] M. Labbé and F. Michaud, “RTAB-Map as an open-source lidar and visual simultaneous localization and mapping library for large-scale and long-term online operation,” *Journal of Field Robotics*, vol. 36, no. 2, pp. 416–446, 2019.
- [12] A. Geiger, P. Lenz, and R. Urtasun, “Are we ready for autonomous driving? the KITTI vision benchmark suite,” in *2012 IEEE Conference on Computer Vision and Pattern Recognition*, 2012, pp. 3354–3361.
- [13] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski, “ORB: An efficient alternative to sift or surf,” in *2011 International Conference on Computer Vision*, 2011, pp. 2564–2571.
- [14] H. Zhang, C.-C. Chen, H. Vallery, and T. D. Barfoot, “GNSS/Multisensor fusion using continuous-time factor graph optimization for robust localization,” *IEEE Transactions on Robotics*, vol. 40, pp. 4003–4023, 2024.