

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

# DINO-VO: A Feature-based Visual Odometry Leveraging a Visual Foundation Model

Maulana Bisyr Azhari<sup>1</sup> and David Hyunchul Shim<sup>1</sup>

**Abstract**—Learning-based monocular visual odometry (VO) poses robustness, generalization, and efficiency challenges in robotics. Recent advances in visual foundation models, such as DINOv2, have improved robustness and generalization in various vision tasks, yet their integration in VO remains limited due to coarse feature granularity. In this paper, we present DINO-VO, a feature-based VO system leveraging DINOv2 visual foundation model for its sparse feature matching. To address the integration challenge, we propose a salient keypoints detector tailored to DINOv2’s coarse features. Furthermore, we complement DINOv2’s robust-semantic features with fine-grained geometric features, resulting in more localizable representations. Finally, a transformer-based matcher and differentiable pose estimation layer enable precise camera motion estimation by learning good matches. Against prior detector-descriptor networks like SuperPoint, DINO-VO demonstrates greater robustness in challenging environments. Furthermore, we show superior accuracy and generalization of the proposed feature descriptors against standalone DINOv2 coarse features. DINO-VO outperforms prior frame-to-frame VO methods on the TartanAir and KITTI datasets and is competitive on EuRoC dataset, while running efficiently at 72 FPS with less than 1GB of memory usage on a single GPU. Moreover, it performs competitively against Visual SLAM systems on outdoor driving scenarios, showcasing its generalization capabilities.

**Index Terms**—Deep Learning Methods; Localization; Visual Odometry; Visual Foundation Model.

## I. INTRODUCTION

**V**ISUAL ODOMETRY (VO) is a technique used to estimate the motion of a camera or robot by analyzing changes in visual data captured from consecutive images. VO is a critical component of many modern technologies, including autonomous vehicles, drones, augmented reality, and robotics. However, acquiring robust and accurate VO is challenging in environments with low visual texture, poor lighting, rapid motion, and dynamic objects. Traditional VO methods extract a set of keypoints using handcrafted features [1] or directly from pixel intensities [2], associate them in successive frames, recover the camera motion, and optimize the poses and structures. However, they lack robustness in low-texture

Manuscript received: March, 25, 2025; Revised June, 12, 2025; Accepted July, 13, 2025.

This paper was recommended for publication by Editor Sven Behnke upon evaluation of the Associate Editor and Reviewers’ comments. This work was financially supported by the Institute of Civil Military Technology Cooperation funded by the Defense Acquisition Program Administration and Ministry of Trade, Industry, and Energy of Korean government under grant No. UM222006RD2. (Corresponding author: David Hyunchul Shim.)

<sup>1</sup>The authors are with the School of Electrical Engineering, Korea Advanced Institute of Science and Technology, Daejeon 305701, South Korea (e-mail: mbazhari@kaist.ac.kr; hcshim@kaist.ac.kr).

Digital Object Identifier (DOI): see top of this page.

©2026 IEEE

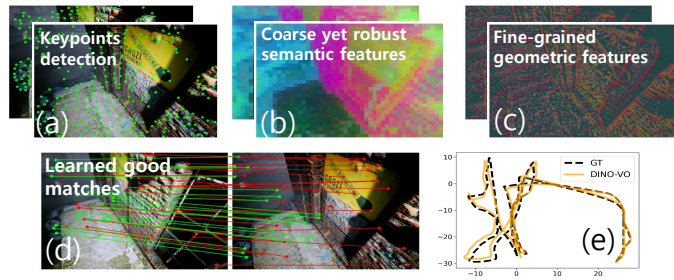


Fig. 1. High-level overview of DINO-VO. (a) The proposed keypoints detection query its corresponding (b) coarse yet robust semantic features from visual foundation model and (c) fine-grained localizable geometric features. (d) The keypoints are matched along with their predicted confidences (showing the 20 highest and 20 lowest-confidence matches). (e) Trajectory estimation on TartanAir dataset sequence MH000, where DINO-VO tracks the ego-motion in a frame-to-frame manner with low-drift.

environments, under variable lighting, and in dynamic object scenarios. Recent works have aimed to improve robustness by utilizing learning-based features [3], [4] or via end-to-end learning frameworks [5]–[10] with auxiliary objectives such as depth and/or optical flow. Nonetheless, these methods struggle with generalization [3]–[6] and are memory-intensive [8]–[10].

Recently, visual foundation models [11], [12] have demonstrated promising results in improving both robustness and generalization across various tasks. Specifically, DINOv2 [12], a strong pre-trained image encoder based on Vision Transformer (ViT) [13], has been demonstrated to perform well in other vision tasks [14]–[18]. DINOv2 has the potential of improving the robustness and generalization ability in VO. However, the direct application of DINOv2 on VO is not straightforward, as the ViT encoder results in coarse features due to patchification [15], [17].

In this paper, we introduce **DINO-VO**, a novel approach that leverages DINOv2 visual foundation model [12] to improve robustness and generalization for a sparse feature-based VO. The high-level overview of DINO-VO is depicted in Fig. 1. To address the issue of coarse features, we propose a lightweight keypoints detector specifically designed to align with DINOv2’s coarseness. We further employ a lightweight CNN encoder to enhance keypoint localization by providing fine-grained, low-level geometric features. Furthermore, we utilize transformer-based feature matching [19] to find keypoint correspondences. The utilization of learning-based matching layers, particularly transformer-based methods, has proven highly beneficial and demonstrated strong performance [15], [19], [20]. When coupled with a learned differentiable pose estimation layer [21], this approach allows us to focus directly on learning good matches for estimating camera motion.

To summarize, our contributions are as follows:

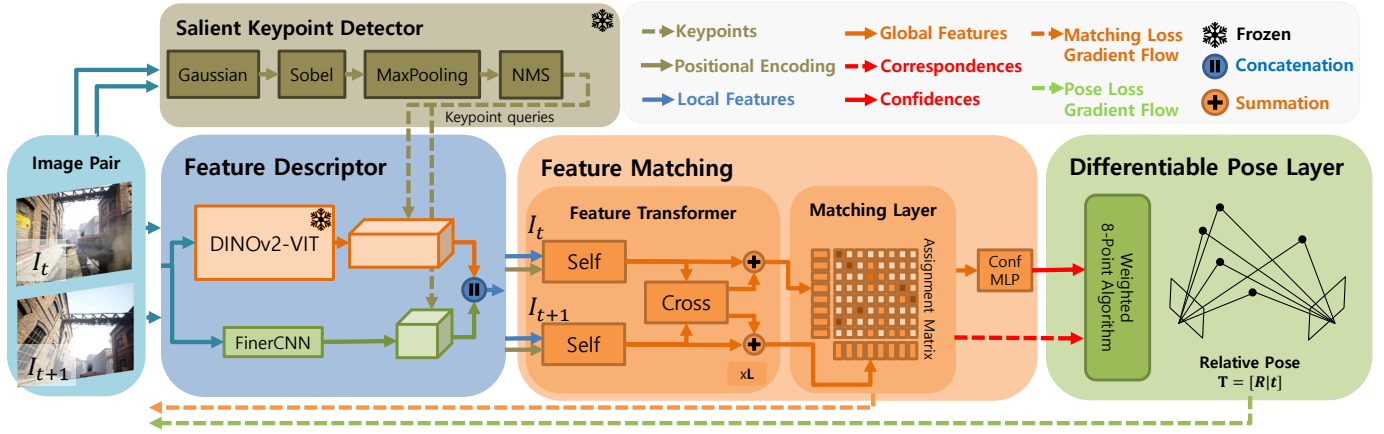


Fig. 2. Overall architecture of the proposed method. The visual odometry (VO) system receives a pair of images  $\{I_t, I_{t+1}\}$  and predict its relative transformation  $\mathbf{T} = [\mathbf{R}|\mathbf{t}]$ . The proposed method consists of 4 modules. First, the keypoints detector detects salient keypoints of each image that align with the coarseness of the DINOv2-ViT feature descriptor. Second, the keypoints query its feature descriptor generated by DINOv2-ViT Encoder –which provides robust and generalizable features– accompanied by a lightweight CNN Encoder to provide low-level geometrical and localizable features. Third, a transformer-based feature matching network is employed to get keypoint correspondences along with its confidence. Lastly, a differentiable pose layer predicts the image’s transformation given its keypoint correspondences and confidences.

- DINO-VO is a novel VO framework that directly leverages DINOv2 [12] visual foundation model features resulting in a robust and general system.
- We propose a salient keypoints detector and a fine-grained geometric descriptor to complement DINOv2 coarse features, which are lightweight yet significantly improve the VO performance.
- We validate DINO-VO on TartanAir, EuRoC, and KITTI datasets, achieving state-of-the-art performances compared to prior frame-to-frame VO methods. We further show the efficacy of our design choices on various sequences.

## II. RELATED WORK

### A. Learned Monocular Visual Odometry

The first notable learning-based VO is DeepVO [5], which presents an end-to-end VO using deep RCNNs. D3VO [22] integrates depth, pose, and uncertainty estimations in its sparse direct visual odometry framework. TartanVO [7] and iSLAM [23] estimate camera motion using optical flow prediction and an intrinsics camera layer, enabling out-of-domain evaluation and real-world use. DiffposeNet [24] predicts normal flow for pose estimation via a differentiable cheirality layer. XVO [25] utilizes multimodal supervision to achieve generalization across diverse driving scenarios. Other methods such as DROID-SLAM [8], DPVO [9] and DPV-SLAM [10] integrate RAFT [26] in an end-to-end system combining optical flow and uncertainty estimation using iterative GRU updates, followed by geometric optimization via bundle adjustment. However, these systems [8]–[10] mostly rely on dense optical flow prediction which requires intensive memory and/or computation. Despite the effort of using “sparse patch” in their optical flow matching processes, [9] and [10] will still require storing the dense feature map for the optical-flow matching process.

### B. Foundation Models for Computer Vision Tasks

Foundation models for computer vision have advanced rapidly in recent years [11], [12]. DINOv2 [12], in particular,




Fig. 3. Visualization of salient keypoints detector. **left**: input image. **middle**: gradient map and salient point selection on each grid, shown in blue. **right**: the red, orange, and black points correspond to removed points after NMS, thresholding, and top-k selections, respectively. The green points are the chosen keypoints and will be used to query feature descriptors.

introduced a scalable self-supervised approach that produces high-quality visual features without labeled data, achieving state-of-the-art results across various perception tasks [14]–[18]. RoMa [15] utilized DINOv2 for coarse matching in a two-stage dense image matching task, refined by a specialized ConvNet, achieving state-of-the-art performance across multiple datasets. Similarly, Dino-tracker [17] used DINOv2 for a dense pixel-tracking task. MambaVO [27] uses DINOv2 features in its demi-dense matching which are then refined using the Geometric Mamba Module. In sparse matching task, DINOv2 guided the matching layer in OmniGlue [16], significantly improving zero-shot generalization. However, DINOv2’s coarse feature map complicates its integration into tasks requiring pixel-level accuracy, such as visual odometry. Unlike previous approaches [15]–[17], [27], we aim to use DINOv2 directly in one-stage sparse image matching to meet the real-time requirement of visual odometry.

## III. METHOD

The objective of our work is to predict the relative transformation  $\mathbf{T} = [\mathbf{R}|\mathbf{t}]$ , composed of rotation matrix  $\mathbf{R} \in SO(3)$  and up-to-scale translation vector  $\mathbf{t} \in \mathbb{R}^3$ , between two consecutive images  $\{I_t, I_{t+1}\} \in \mathbb{R}^{H \times W \times 3}$  for a VO frontend system. As shown in Fig. 2, our proposed VO system integrates four key components: a salient keypoints detector, a feature descriptor networks that utilize DINOv2 along with a lightweight CNN encoder, a transformer-based feature matching layer, and a differentiable pose layer.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

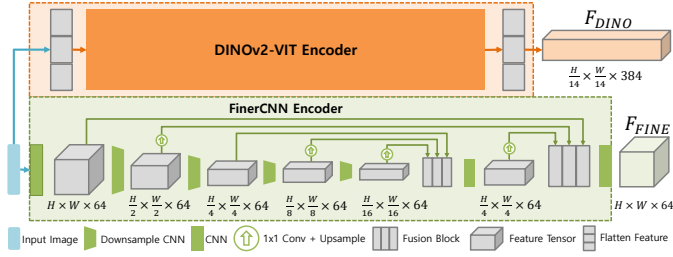


Fig. 4. Feature descriptor networks. It consists of two encoders: a DINOv2-ViT and a specialized lightweight CNN encoder, FinerCNN.

### A. Salient Keypoints Detector

As noted in [16], combining SuperPoint [28] with DINOv2’s coarse features poses challenges in sparse matching. RoMa [15] showed DINOv2 lacks precise localization, prompting two stage dense matching, but its high computational cost limits real-time VO. Our proposed keypoints detector queries DINOv2 features from sparse salient points, enabling precise, one-stage matching with low computational demands. The visualization of the proposed keypoints detector is depicted in Fig. 3.

The keypoints detector computes the gradient map using a GaussianFilter and a SobelFilter. The gradient map is then divided into  $H/r_p \times W/r_p$  grids, aligning keypoints with DINOv2 coarse features, where  $r_p = 14$  corresponds to the patch size of the DINOv2-ViT. We select the point with the highest gradient magnitude in each grid using a MaxPooling operator with a kernel size of  $r_p$  and a stride of  $r_p$ , resulting in the most salient point in each grid (shown in blue). Oftentimes, the salient points are located at the grid borders, leading to keypoints that are very close to one another, making them redundant (shown in red). To address this, we apply NonMaxSuppression (NMS) [28] with a radius of  $r_{NMS}$ . In addition to removing redundant points, a large  $r_{NMS}$  ensures that keypoints are well-distributed across the image, which improves the robustness and accuracy of pose estimation. Finally, we filter out keypoints with very small gradient magnitudes (shown in orange) and select the top- $k$  keypoints based on their gradient magnitudes (shown in green). The output, keypoint locations  $\mathcal{K} = \{(x_i, y_i)\}_{i=1}^K$  where  $x_i \in [1, H]$  and  $y_i \in [1, W]$ , is then used to query features described in the next subsection.

### B. Feature Descriptor

As shown in Fig. 4, the proposed feature descriptor consists of two networks: DINOv2 and FinerCNN. The DINOv2 encoder extracts generalizable and robust semantic features for our matching process and produces a feature map  $\mathbf{F}_{DINO} \in \mathbb{R}^{\frac{H}{14} \times \frac{W}{14} \times 384}$ . To complement DINOv2 coarse features, we employ FinerCNN to provide fine-grained low-level geometric features. FinerCNN is built from CNN blocks called basic layers, inspired by XFeat [29], which has shown great performance in image matching while having lightweight design. We adopt the feature pyramid approach to expand the receptive field until the resolution is reduced to  $H/16 \times W/16$ . The intermediate features are then fused at  $H/4 \times W/4$  and  $H \times W$  resolution levels as shown in Fig. 4. FinerCNN generates a

feature map  $\mathbf{F}_{FINE} \in \mathbb{R}^{H \times W \times 64}$  that has the same spatial dimension as the input image providing the optimal localizable features at pixel-level resolution.

For each keypoint location  $(x_i, y_i) \in \mathcal{K}$ , we directly retrieve its corresponding feature from the feature map by indexing it as  $\mathbf{f}_{FINE}^i := \mathbf{F}_{FINE}(x_i, y_i, :) \in \mathbb{R}^{64}$ . To query its corresponding feature in  $\mathbf{F}_{DINO}$ , we apply division and floor operations to the keypoint location  $\mathbf{f}_{DINO}^i := \mathbf{F}_{DINO}(\lfloor \frac{x_i}{14} \rfloor, \lfloor \frac{y_i}{14} \rfloor, :) \in \mathbb{R}^{384}$ . Finally, the two features are concatenated, and a Linear projector is used to reduce the dimensionality.

$$\mathbf{f}_i = \text{Linear}([\mathbf{f}_{DINO}^i | \mathbf{f}_{FINE}^i]) \in \mathbb{R}^{192}. \quad (1)$$

where  $[\cdot | \cdot]$  concatenates two vectors in the channel dimension.

### C. Feature Matching

Given two consecutive images  $I_t$  and  $I_{t+1}$ , we extract  $K$  features  $\mathbf{f}_i^{I_t} \in \mathbb{R}^{192}$  indexed by  $\mathcal{K}_t = \{(x_i, y_i)\}_{i=1}^K$  for image  $I_t$  and  $\mathbf{f}_j^{I_{t+1}} \in \mathbb{R}^{192}$  indexed by  $\mathcal{K}_{t+1} = \{(x_j, y_j)\}_{j=1}^K$  for image  $I_{t+1}$ , using feature descriptor explained in previous subsection. Features from the two images are matched by predicting a partial assignment  $\mathcal{M} = \{(i, j)\} \subset \mathcal{K}_t \times \mathcal{K}_{t+1}$ . We use learnable feature matching based on Transformer architecture, as it has been shown to be more robust compared to traditional feature matching methods [16], [19], [30]. Specifically, following [19], we use  $L$  identical transformer layers consisting of self-attention and cross-attention units to process the local features into features with global context. In each attention unit, the feature descriptor  $\mathbf{f}_i^T$  is updated using the following formulas:

$$\mathbf{f}_i^T \leftarrow \mathbf{f}_i^T + \text{MLP}([\mathbf{f}_i^T | \mathbf{m}_i^{T \leftarrow S}]), \quad \text{with} \quad (2)$$

$$\mathbf{m}_i^{T \leftarrow S} = \sum_{j \in S} \text{Softmax}(a_{ij}^{TS}) \mathbf{v}_j, \quad \text{and} \quad (3)$$

$$a_{ij}^{TS} = \begin{cases} \mathbf{q}_i^\top \mathbf{R}_P(\mathbf{x}_j - \mathbf{x}_i) \mathbf{k}_j, & \text{for self-attention} \\ \mathbf{k}_i^{T \top} \mathbf{k}_j^S \stackrel{!}{=} a_{ji}^{ST}, & \text{for cross-attention} \end{cases} \quad (4)$$

where  $\mathbf{m}_i^{T \leftarrow S}$  is the aggregated message from a source image  $S \in \{I_t, I_{t+1}\}$  to a target image  $T \in \{I_t, I_{t+1}\}$  using attention mechanism.  $a_{ij}^{TS}$  is the attention score between feature  $i$  of image  $T$  and feature  $j$  of image  $S$  which is computed differently in the self- and cross-attention unit using Eq. 4, with  $S = T$  in the self-attention and  $S = \{I_t, I_{t+1}\} \setminus T$  in the cross-attention.  $\mathbf{k}_i$ ,  $\mathbf{q}_i$ , and  $\mathbf{v}_i$  are the key, query and value vector of linearly projected  $\mathbf{f}_i$ , respectively.  $\mathbf{R}_P(\cdot) \in \mathbb{R}^{d \times d}$  is a rotary positional encoding [31] and  $\mathbf{p}_i := (x_i, y_i) \in [0, 1]^2$  is the normalized keypoint location of feature  $i$ .

The point correspondences can be obtained by combining the feature similarity  $\mathbf{S}_{ij}$  and feature-wise matchability scores  $\sigma_i, \sigma_j$  into a soft partial assignment matrix  $\mathbf{P} \in [0, 1]^{K \times K}$  as

$$\mathbf{P}_{ij} = \sigma_i \sigma_j \text{Softmax}(\mathbf{S}_{kj})_i \text{Softmax}(\mathbf{S}_{ik})_j, \quad \text{where} \quad (5)$$

$$\mathbf{S}_{ij} = \text{Linear}(\mathbf{f}_i^{I_t})^\top \text{Linear}(\mathbf{f}_j^{I_{t+1}}), \quad \text{and} \quad (6)$$

$$\sigma_i = \text{Sigmoid}(\text{Linear}(\mathbf{f}_i^{I_t})) \in [0, 1], \quad (7)$$

$$\sigma_j = \text{Sigmoid}(\text{Linear}(\mathbf{f}_j^{I_{t+1}})) \in [0, 1], \quad (8)$$

**IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.**

A pair of points  $(i, j)$  yields a correspondence when both points are predicted as matchable and when their similarity is higher than any other point in both images. Furthermore, for each pair of matching points  $(i, j)$  a confidence weight  $w_{ij}$  is predicted from the final layer descriptor  $\mathbf{f}_i^{I_t}$  and  $\mathbf{f}_j^{I_{t+1}}$  as:

$$w_{ij} = \text{ConfMLP} \left( \left[ \mathbf{f}_i^{I_t} | \mathbf{f}_j^{I_{t+1}} \right] \right). \quad (9)$$

#### D. Differentiable Pose Estimation

We utilize a confidence-weighted eight-point algorithm as the differentiable pose layer to estimate the relative pose between two images given its corresponding points. The weighted eight-point algorithm has shown to be useful and efficient for providing supervision to the matching layer to learn good matches for pose estimation [21]. For each pair of corresponding points  $\mathbf{x}_i^{\mathcal{K}_t} = [x_i, y_i, 1]^T$  and  $\mathbf{x}_j^{\mathcal{K}_{t+1}} = [x_j, y_j, 1]^T$  with  $x, y \in [0, 1]$ , the essential matrix  $\mathbf{E} \in \mathbb{R}^{3 \times 3}$  satisfies the following equation:

$$\mathbf{x}_j^T \mathbf{E} \mathbf{x}_i = 0, \quad (10)$$

By stacking the known coordinates of correspondences into a matrix  $\Phi$ , where each row  $[x_i x_j, x_i y_j, x_j, y_i x_j, y_i y_j, y, x_j, y_j, 1]$  describes one correspondence and flattening  $\mathbf{E}$  into  $\text{flat}(\mathbf{E}) = [e_{11}, e_{12}, e_{13}, e_{21}, e_{22}, e_{23}, e_{31}, e_{32}, e_{33}]^T$ , we can rewrite the confidence-weighted version of the eight-point algorithm into a matrix equation

$$\text{diag}(\mathbf{w}) \Phi \text{flat}(\mathbf{E}) = \mathbf{0}, \quad (11)$$

where  $\mathbf{w}$  is the weight vector obtained by stacking the corresponding  $w_{ij}$  of every match from (9).

We solve the Eq. (11) by using Singular value decomposition (SVD) of  $\text{diag}(\mathbf{w}) \Phi$ , which minimizes  $\|\text{diag}(\mathbf{w}) \text{flat}(\mathbf{E})\|_2$  and enforcing  $\mathbf{E}$  to have rank 2.  $\mathbf{E}$  can be decomposed into multiple relative transformations  $\mathbf{T}$ , the best  $\mathbf{T}$  is then chosen by using cheirality condition [32].

#### E. Supervision

Following [19], we supervise the assignment matrix  $\mathbf{P}$  using ground truth labels derived from two-view transformations and pixel-wise depth. We minimize the log-likelihood of the assignment predicted at each layer  $\ell$ . The matching objective  $\mathcal{L}_m$  is defined as:

$$\begin{aligned} \mathcal{L}_m = & -\frac{1}{L} \sum_{\ell} \left( \frac{1}{|\mathcal{M}|} \sum_{(i,j) \in \mathcal{M}} \log^{\ell} P_{ij} \right. \\ & + \frac{1}{2|\bar{\mathcal{K}}_t|} \sum_{i \in \bar{\mathcal{K}}_t} \log \left( 1 - \sigma_i^{\mathcal{K}_t} \right) \\ & \left. + \frac{1}{2|\bar{\mathcal{K}}_{t+1}|} \sum_{j \in \bar{\mathcal{K}}_{t+1}} \log \left( 1 - \sigma_j^{\mathcal{K}_{t+1}} \right) \right). \end{aligned} \quad (12)$$

where  $\bar{\mathcal{K}}_t \subseteq \mathcal{K}_t$  and  $\bar{\mathcal{K}}_{t+1} \subseteq \mathcal{K}_{t+1}$  are unmatchable keypoints.

We also minimize the difference between the predicted transformation  $\hat{\mathbf{T}} = [\hat{\mathbf{R}} | \hat{\mathbf{t}}] \in \text{Sim}(3)$  and the ground-truth

$\mathbf{T} = [\mathbf{R} | \mathbf{t}] \in \text{SE}(3)$  as the up-to-scale pose objective  $\mathcal{L}_p$  following [7] to handle scale ambiguity in monocular setting, which is defined as:

$$\begin{aligned} \mathcal{L}_p = & \lambda_t \left\| \frac{\hat{\mathbf{t}}}{\max(\|\hat{\mathbf{t}}\|, \epsilon)} - \frac{\mathbf{t}}{\max(\|\mathbf{t}\|, \epsilon)} \right\| \\ & + \lambda_r \|\text{Log}(\hat{\mathbf{R}}) - \text{Log}(\mathbf{R})\|. \end{aligned} \quad (13)$$

where  $\epsilon = 1e^{-6}$  to stabilize the numerical division,  $\lambda_t$  and  $\lambda_r$  are the weighting factor for the translation and rotation objective, respectively.  $\text{Log}(\cdot)$  is a mapping function from  $\text{SO}(3)$  to  $\text{so}(3)$ . The pose objective guides the matching layer to down-weight matches that are detrimental to pose accuracy. The overall objective  $\mathcal{L}_t$  is defined as a weighted summation of both objectives:

$$\mathcal{L}_t = (1 - \lambda_p) \mathcal{L}_m + \lambda_p \mathcal{L}_p. \quad (14)$$

where  $\lambda_p$  is the weighting factor of the pose objective.

#### F. DINO-VO

DINO-VO predicts the relative pose to the latest keyframe, which is selected when the mean pixel displacement of predicted correspondences exceeds 24px. This ensures sufficient parallax, thereby mitigating instabilities from pure rotation or minimal translation of the weighted 8-point algorithm [33]. Furthermore, alternate frames are processed for EuRoC and KITTI datasets, this step is unnecessary for the TartanAir MH dataset due to its inherently significant camera motion. All frame's poses are saved relative to its previous keyframe such that the full camera trajectory is obtained.

Since DINO-VO only predicts up-to-scale transformation, we follow [7] to scale the predicted translation  $\mathbf{t} \in \mathbb{R}^3$  with the ground-truth. DINO-VO uses the image resolution of  $476 \times 630$ ,  $476 \times 742$ , and  $364 \times 1008$  for TartanAir, EuRoC, and KITTI sequences, respectively.

## IV. EXPERIMENTS

### A. Implementation Details

1) *Detector and Networks Details:* We use GaussianFilter with kernel size 5 and  $std = 2.0$ , MaxPooling with kernel size of  $r_P = 14$  following the patch-size of DINOv2-ViT-s encoder [12],  $r_{\text{NMS}} = 8$ , gradient threshold of 0.01. We select the top 512 keypoints for image matching and pose estimation. We utilize LightGlue [19] as the matching layer with 12 layers, 3 attention heads with 64 head dimensions.

2) *Training Procedure:* We use only TartanAir dataset [34] for training. We train the model with matching objective  $\mathcal{L}_m$  for the first 4 epochs and total objective  $\mathcal{L}_t$  for the next 10 epochs. We use  $\lambda_r = 180$ ,  $\lambda_t = 400$ , and  $\lambda_p = 0.0$  up to  $\lambda_p = 0.9$  with  $1.5 \times 10^{-4}$  increment per step starting from epoch 5. We use a computer equipped with an RTX4090 GPU and an intel i9-12900k CPU with PyTorch. The entire training procedure takes about 3 days.

3) *Metrics:* Following prior works [7], [9], we report Absolute Trajectory Error (ATE) in meter for TartanAir MH, EuRoC, and KITTI datasets. Additionally, we report RMSE drift translational  $t_{rel}$  (%) and rotational  $r_{rel}$  (deg/100m) on KITTI dataset [35] sequences 6, 7, 9, and 10.

## IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

TABLE I  
ATE[M]↓ ON THE MH SEQUENCES OF TARTANAIR DATASET.

Sequence (MH)	Multi-Frame VO					Frame-to-Frame VO		
	SVO [36]	DSO [2]	DeepV2D [6]	DROID-VO [8]	DPVO [9]	TartanVO [7]	DiffPoseNet [24]	Ours (DINO-VO)
000	14.42	9.65	6.15	0.32	0.21	4.88	<u>2.56</u>	<b>0.87</b>
001	1.17	0.35	2.12	0.13	0.04	<u>0.26</u>	0.31	<b>0.19</b>
002	4.97	7.96	4.54	0.08	0.04	2.00	<u>1.57</u>	<b>0.28</b>
003	6.88	3.46	3.89	0.09	0.08	0.94	<u>0.72</u>	<b>0.25</b>
004	X	X	2.71	1.52	0.58	<u>1.07</u>	<b>0.82</b>	1.19
005	19.2	12.58	11.55	0.69	0.17	3.19	<u>1.83</u>	<b>0.92</b>
006	11.27	8.42	5.53	0.39	0.11	<u>1.00</u>	1.32	<b>0.23</b>
007	17.68	7.50	3.76	0.97	0.15	2.04	<u>1.24</u>	<b>0.73</b>
Avg	X	X	5.03	0.52	0.17	1.92	<u>1.29</u>	<b>0.58</b>

- The best and second best of frame-to-frame VO results are marked as **bold** and underlined.

TABLE II  
ATE[M]↓ ON THE EUROC MAV DATASET.

Sequence	Multi-Frame VO					Frame-to-Frame VO			
	SVO [36]	DSO [2]	DeepV2D [6]	DROID-VO [8]	DPVO [9]	TartanVO [7]	iSLAM-VO <sup>+</sup> [23]	MAC-VO <sup>+</sup> [37]	Ours (DINO-VO)
MH01	0.100	0.046	1.614	0.163	0.087	0.693	0.320	<u>0.240</u>	<b>0.150</b>
MH02	0.120	0.046	1.492	0.121	0.055	0.325	0.462	<u>0.256</u>	<b>0.118</b>
MH03	0.410	0.172	1.635	0.242	0.158	0.550	0.380	<u>0.260</u>	<b>0.254</b>
MH04	0.430	3.810	1.775	0.399	0.137	1.153	<u>0.962</u>	<b>0.496</b>	0.982
MH05	0.300	0.110	1.013	0.270	0.114	1.021	<u>0.500</u>	<b>0.560</b>	<b>0.454</b>
V101	0.070	0.089	0.717	0.102	0.050	0.447	0.366	<b>0.197</b>	<u>0.298</u>
V102	0.210	0.107	0.695	0.165	0.140	0.389	0.414	<b>0.096</b>	<u>0.260</u>
V102	X	0.903	1.483	0.158	0.086	0.622	<u>0.313</u>	<b>0.198</b>	0.559
V201	0.110	0.044	0.839	0.102	0.057	0.433	0.478	<b>0.145</b>	<u>0.334</u>
V202	0.110	0.132	1.052	0.115	0.049	0.749	0.424	<b>0.270</b>	<u>0.278</u>
V202	1.080	1.152	0.592	0.204	0.211	1.152	1.176	<b>0.425</b>	<u>0.763</u>
Avg	X	0.601	1.173	0.186	0.105	0.680	0.527	<b>0.286</b>	<u>0.404</u>

- The best and second best of frame-to-frame VO results are marked as **bold** and underlined.

- <sup>+</sup> The method uses stereo camera setup.

### B. Monocular Visual Odometry Results

In this section, we evaluate the accuracy and robustness of DINO-VO on TartanAir and EuRoC MAV datasets. For a fair comparison, we only numerically compare our method with frame-to-frame VO methods such as TartanVO [7], DiffPoseNet [24], MAC-VO [37], and a stereo VO frontend of iSLAM [23], named iSLAM-VO. We also provide the results of the SOTA VO methods which perform multi-frame optimization such as SVO [36], DSO [2], DeepV2D [6], DROID-VO [8], and DPVO [9].

In Table I, we present the evaluation of DINO-VO on TartanAir MH sequences, which span a variety of simulated environments with numerous challenges, including illumination changes, dynamic objects, low-texture scenes, and significant viewpoint variations—conditions that are challenging for many traditional methods [34]. Compared to other frame-to-frame VO approaches, DINO-VO demonstrates better accuracy in 7 out of 8 sequences. On average, DINO-VO reduces the ATE by 70% and 55% compared to TartanVO, and DiffPoseNet, respectively.

In Table II, we evaluate DINO-VO on EuRoC dataset,

TABLE III  
GENERALIZATION ASSESSMENT AGAINST OTHER LEARNED FRAME-TO-FRAME VO METHODS ON KITTI ODOMETRY SEQUENCE 06, 07, 09, AND 10 USING RMSE DRIFT↓  $t_{rel}$  AND  $r_{rel}$ .

Method	Met.	Sequence				Avg
		06	07	09	10	
DeepVO [5]	$t_{rel}$	5.42	3.91	-	8.11	-
	$r_{rel}$	5.82	4.60	-	8.83	-
DFVO [33]	$t_{rel}$	-	-	2.47	1.96	-
	$r_{rel}$	-	-	<b>0.30</b>	<b>0.31</b>	-
HDVO <sup>+</sup> [38]	$t_{rel}$	-	-	<u>1.97</u>	<u>1.89</u>	-
	$r_{rel}$	-	-	0.71	0.56	-
XVO [25]	$t_{rel}$	3.92	5.96	-	3.31	-
	$r_{rel}$	<u>1.46</u>	3.96	-	1.52	-
DualRefine [39]	$t_{rel}$	-	-	3.43	6.80	-
	$r_{rel}$	-	-	1.04	1.13	-
ALH-VO <sup>+</sup> [40]	$t_{rel}$	3.64	<b>2.53</b>	2.01	2.01	<u>2.55</u>
	$r_{rel}$	<u>1.88</u>	<b>1.46</b>	1.35	1.32	<u>1.50</u>
TartanVO [7]	$t_{rel}$	4.72	4.32	6.00	6.89	5.48
	$r_{rel}$	2.95	3.41	3.11	2.73	3.05
DiffPoseNet [24]	$t_{rel}$	2.93	4.06	4.02	3.95	3.74
	$r_{rel}$	1.76	2.35	<u>0.51</u>	1.23	1.46
iSLAM-VO <sup>+</sup> [23]	$t_{rel}$	4.30	3.23	4.22	4.63	4.09
	$r_{rel}$	1.16	<u>1.75</u>	1.41	1.40	<b>1.43</b>
MAC-VO <sup>+</sup> [37]	$t_{rel}$	2.91	3.56	6.89	8.03	5.34
	$r_{rel}$	1.74	2.58	2.69	3.82	2.70
<b>Ours (DINO-VO)</b>	$t_{rel}$	<b>1.65</b>	<u>2.57</u>	<b>1.64</b>	<b>0.96</b>	<b>1.70</b>
	$r_{rel}$	<b>0.82</b>	3.31	0.93	<u>0.71</u>	1.44

-  $t_{rel}$  (%) and  $r_{rel}$  (deg/100m) are the average translational and rotational RMSE drift on various segments with length 100-800 m.

- The best of each driving scenarios and general frame-to-frame VO results are underlined. The best overall results are **bold**.

- <sup>+</sup> The method uses stereo camera setup.

TABLE IV  
GENERALIZATION ASSESSMENT AGAINST OTHER GENERAL SLAM/VO METHODS ON KITTI ODOMETRY SEQUENCE 0-10 USING ATE[M]↓.

Sequence	SLAM				VO					
	ORB-SLAM3 [1]	LDSO [41]	DPV-SLAM++ [10]	Mamba VO++ [27]	SVO [36]	DSO [2]	DROID-VO [8]	DPVO [9]	Mamba VO [27]	Ours (DINO-VO)
00	6.77	9.32	8.30	<b>6.19</b>	X	126.7	98.4	113.2	112.4	39.8
01	X	11.68	11.3	<b>8.04</b>	X	165.0	84.2	12.7	<u>8.16</u>	24.8
02	30.5	31.98	39.6	<b>27.73</b>	X	138.7	108.8	123.4	93.78	<u>42.6</u>
03	<b>1.04</b>	2.85	2.50	1.94	X	4.77	2.58	2.09	1.80	<u>1.51</u>
04	0.93	1.22	0.78	<u>0.59</u>	58.4	1.08	0.93	0.68	0.66	<b>0.35</b>
05	5.54	5.1	5.74	<b>3.05</b>	X	49.5	59.2	58.9	56.51	<u>15.4</u>
06	16.6	13.55	<u>11.6</u>	11.79	X	113.6	64.4	54.8	57.19	<b>3.91</b>
07	9.70	2.96	<b>1.52</b>	1.7	X	27.9	24.2	19.2	17.9	<u>7.39</u>
08	60.7	129.0	110.9	105.4	X	120.2	64.5	115.9	116.0	<b>21.6</b>
09	<u>7.89</u>	21.64	76.7	63.24	X	74.3	71.8	75.1	73.56	<u>7.16</u>
10	<u>8.65</u>	17.36	13.7	10.51	X	16.3	16.9	13.6	14.37	<b>1.31</b>
Avg	X	22.42	25.7	<u>21.84</u>	X	76.3	54.2	53.6	50.21	<b>15.1</b>

- The best of each SLAM and VO results are underlined. The best overall results are **bold**.

which is well-known for its challenging sequence for learning-based methods due to lack of training data, challenging lighting conditions, and aggressive motion patterns. Compared to other frame-to-frame VO methods, DINO-VO achieves the lowest ATE on 4 out of 11 sequences in the EuRoC dataset, highlighting not only its generalization capabilities but also its accuracy and robustness, as it successfully completes all sequences without abrupt failures. DINO-VO reduces the ATE by 40% and 23% compared to TartanVO and iSLAM-VO, respectively. DINO-VO has a higher average ATE compared to MAC-VO, however, MAC-VO requires a stereo camera setup, 6× GPU memory, and 13× inference time (see Sec. IV-D).




Fig. 5. Inference time and memory usage comparison on EuRoC sequences using a single RTX4090 GPU and i9-12900k CPU. DINO-VO uses  $476 \times 742$  image input and half-precision inference (FP16) while other methods follow their default implementation.

TABLE V  
TIMING BREAKDOWN OF DINO-VO.

$t_{\text{Det}}$	$t_{\text{FinerCNN}}$	$t_{\text{DINOv2}}$	$t_{\text{Match}}$	$t_{\text{Pose}}$	$t_{\text{Total}}$
0.65 ms	1.10 ms	4.05 ms	3.85 ms	4.15 ms	13.8 ms

### C. Generalization to Outdoor Driving Scenarios

We evaluate the generalization ability of DINO-VO on KITTI dataset [35] which contains outdoor driving sequences, featuring dynamic objects such as cars and pedestrians which provides significant domain gap with TartanAir dataset. In addition to comparing with general VO methods, we also provide the comparison with frame-to-frame VO methods which are specifically trained and/or designed for outdoor driving scenarios [5], [25], [33], [38]–[40].

In Table III, we provide comparison with other frame-to-frame VO methods using RMSE drift [35]. Compared to the general methods, DINO-VO achieves the lowest  $t_{rel}$  on 3 out of 4 sequences and the lowest  $r_{rel}$  on 2 out of 4 sequences. On average, DINO-VO achieves the lowest  $t_{rel}$  with 54% error reduction compared to the second best results, DiffPoseNet. Our method achieves the second lowest average  $r_{rel}$  with less than 1% increase compared to iSLAM-VO. Despite never having seen these scenarios, DINO-VO provides significant improvement against general frame-to-frame VO methods [7], [23], [24], [37] and even methods which are designed/trained for driving scenarios [5], [25], [33], [38]–[40], showcasing its generalization ability.

In Table IV, we compare DINO-VO with other general multi-frame VO/SLAM methods in term of ATE on KITTI odometry sequence 0-10, specifically reported in [10]. DINO-VO has the lowest ATE in 5 out of 11 sequences and achieves the lowest average ATE despite not using multi-frames optimization and loop-closure. Note that on sequence 00, 02, 05, and 07, the visual SLAM systems benefit from loop-closure resulting in better accuracy. On Sequence 01, 03, 06, 08, 09, and 10 which contain significant appearances of dynamic objects, DINO-VO achieves the lowest ATE on 4 out of 6 sequences despite having never seen the objects and using no explicit matches filtering, showcasing its robustness to dynamic objects.

### D. Runtime and Memory

Fig. 5 shows the inference time and memory comparison of DINO-VO with other methods on EuRoC dataset using a single RTX4090 GPU and i9-12900k CPU. DINO-VO runs at 13.8 ms ( 72 FPS) and uses only 0.96GB GPU memory. All of the reported DINO-VO results are run using FP16 half-precision inference with PyTorch software. DINO-VO

TABLE VI  
ABLATION STUDY RESULT.

Det / Desc	Matching Performance On TartanAir			Traj Eval.		
	R	MMA @1/3/5/10 px		TartanAir ATE[m] $\downarrow$	EuRoC ATE[m] $\downarrow$	KITTI ATE[m] $\downarrow$
(1)	SP / SP	<b>92.7</b>	<u>36.8</u> / 70.4 / 78.3 / 83.3	1.14	0.616	<u>28.2</u>
	SP / D	89.1	36.4 / 69.1 / 76.4 / 80.7	<u>1.03</u>	0.630	56.1
	SP / F	<u>92.6</u>	36.7 / 71.0 / <u>78.6</u> / 83.3	1.11	0.634	30.3
	SP / D+F	92.2	<b>38.0</b> / <b>73.0</b> / <b>80.9</b> / <b>84.6</b>	<b>0.68</b>	<b>0.420</b>	<b>20.1</b>
(2)	SP / D+F	<u>92.2</u>	<b>38.0</b> / <b>73.0</b> / 80.9 / 84.6	<u>0.68</u>	<u>0.420</u>	20.1
	SP-A / D+F	<b>92.6</b>	36.9 / 71.3 / 80.1 / 83.3	0.71	0.460	28.1
	SIFT / D+F	83.4	33.8 / 67.0 / 75.1 / 77.3	1.14	0.584	25.9
	SIFT-A / D+F	88.7	31.8 / 65.6 / 78.6 / 83.7	0.92	0.565	<u>18.2</u>
	SAL / D+F	88.3	33.9 / 71.6 / <u>82.2</u> / 83.6	0.78	0.528	19.8
	SAL-A / D+F	91.0	33.1 / 68.5 / <b>82.8</b> / <b>87.0</b>	<b>0.58</b>	<b>0.404</b>	<b>15.1</b>

- The best and second best results of each setting (1) and (2) are **bold** and underlined.

- SP: SuperPoint (baseline); D: DINOv2; F: FinerCNN, D+F: Concatenated DINOv2 and FinerCNN (Proposed in Sec. III-B).

- SP-A: SuperPoint Detector w/ grid alignment; SIFT-A: SIFT detector w/ grid-alignment; SAL: Saliient Detector w/o grid alignment; SAL-A: SAL w/ grid alignment (Proposed in Sec. III-A).

has up to  $9 \times$  and  $13.5 \times$  lower memory footprint and faster inference time, respectively. Compared to prior works, DINO-VO uses sparse feature matching instead of optical-flow [7]–[9], [23], [37]. Furthermore, DINO-VO benefits from highly efficient parallelization of attention mechanism in its feature descriptor [13] and feature matching networks [19]. We further provide the timing breakdown of DINO-VO in Table V. The proposed keypoints detector and fine-grained CNN encoder have minimal overhead compared to the total inference time, using only 5% and 8% of the total runtime.

## V. ABLATION STUDIES

To evaluate the effectiveness of our keypoint detectors and feature descriptors, we conduct ablation studies with different settings. In setting (1), we built on top of SuperPoint (SP) joint detector-descriptor networks [28] as the baseline and ablate the feature descriptors. In setting (2), we ablate the keypoints detectors on top of setting (1) SP/D+F. \*-A denotes keypoints detector with grid-alignment proposed in Sec. III-A. We report the matching performance in term of recall at 5 pixel error threshold and mean matching accuracy (MMA) at various error thresholds [42] as well as the average ATE. The ablation study results are shown in Table VI. We further depict the qualitative comparisons in Fig. 6.

### A. Localizable Features Importance

As shown in Table VI, using DINOv2 descriptors alone (D) results in matching performance drop compared to using SuperPoint descriptors, similar results are observed in [16]. The combination of DINOv2 with FinerCNN (D+F) significantly improves the matching accuracies across all thresholds compared to SuperPoint descriptors. These results show the importance of having localizable features from FinerCNN. The FinerCNN alone has comparable matching performance to SuperPoint while having a  $2.4 \times$  faster inference speed. Beyond the matching performance, the proposed descriptors achieve 40%, 32% and 29% average ATE reduction on TartanAir, EuRoC, and KITTI, compared to using SuperPoint features.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

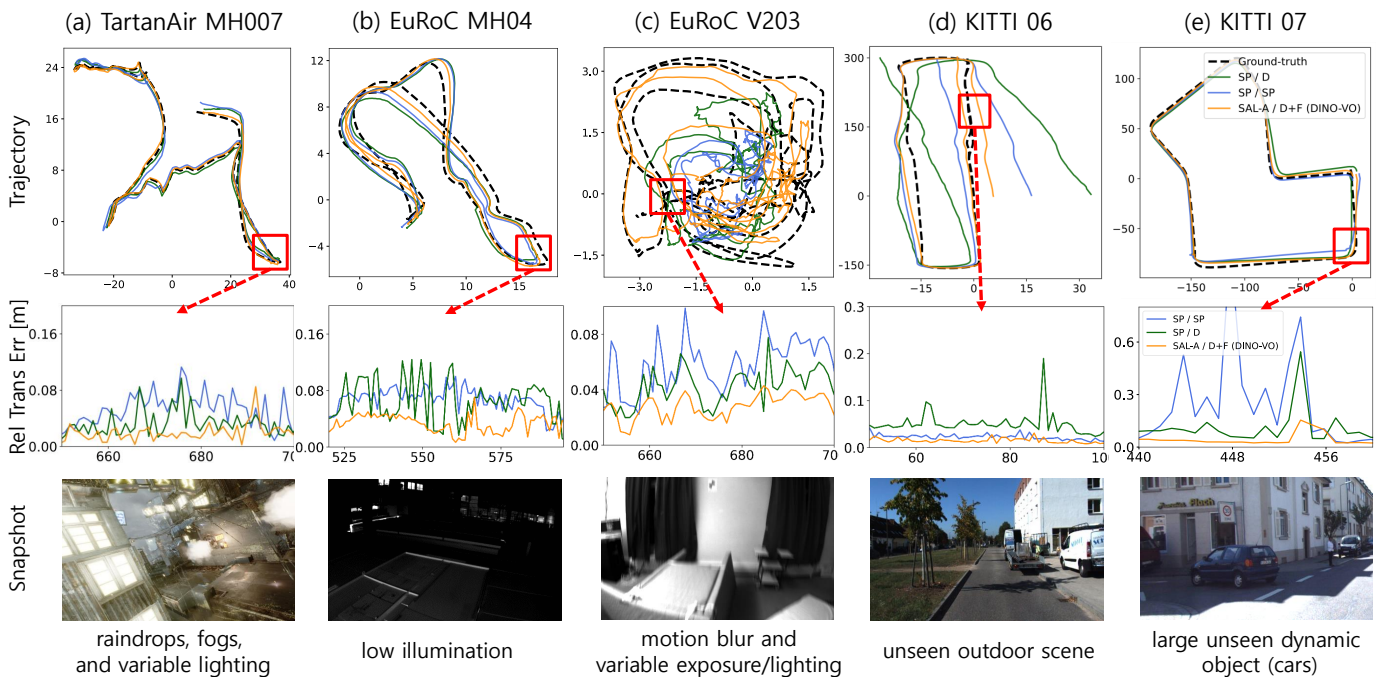


Fig. 6. Qualitative results. **top row**: Trajectory comparisons of the proposed method (DINO-VO) along with two ablated systems, namely ours with SP / SP and ours with SP / D. **middle row**: frame-to-frame relative translation error comparison on the red square area of the trajectory. **bottom row**: snapshot of the image input along with its challenging conditions which disturb the VO systems. DINO-VO has the lowest error on various challenging conditions compared to with the system with SuperPoint detector-descriptor and DINOv2 only features, showcasing the robustness of our system.

### B. Robustness and Generalization of DINOv2 Features

As provided in Table VI, the usage of DINOv2 without FinerCNN shows better accuracy on TartanAir MH sequences which contain significant disturbances [34] despite having less accurate matches. The semantic information provided by DINOv2 is observed to be beneficial to down-weight unfavorable matches for pose estimation. However, using DINOv2 features alone performs significantly worse when tested on out-of-domain sequences as shown in Table VI and Fig. 6. (d)-(e). We suspect that the matching networks learns to exploit specific bias of DINOv2 features in the TartanAir dataset hindering the generalization ability of DINOv2 features. Combining with FinerCNN resolves this issue as the matching network can optimally combine the geometric and semantic information encoded in FinerCNN and DINOv2 features. This combination further improve the robustness which are pronounce in sequences with various disturbances such as variable lighting, fogs, aggressive motion, motion blur and unseen dynamic object as shown Fig. 6. (a), (c), and (e).

### C. Keypoints Detector Effectiveness

The proposed keypoints detector (SAL-A) used in DINO-VO significantly improves the matching and VO accuracy compared to using SuperPoint and SIFT detectors. Without grid-alignment, salient detector (SAL) performs worse compared to SuperPoint detector, showing the importance of having aligned keypoints on DINOv2 coarse features. Introducing grid-alignment to SIFT detector (SIFT-A) significantly improves its matching and VO performance. Conversely, SuperPoint detector w/ alignment (SP-A) results in a slightly degraded performance. Furthermore, SAL-A offers more dis-

tributed keypoints across the image which is beneficial for pose estimation. Additionally, it has  $6\times$  and  $4\times$  faster detector inference time compared to SIFT and SP.

## VI. DISCUSSION AND CONCLUSION

In this paper we present DINO-VO, a novel feature-based visual odometry approach that leverages the robustness and generalization capabilities of the DINOv2 visual foundation model. We verify that DINO-VO is robust against various disturbances and generalizable when tested on out-of-domain datasets. DINO-VO achieves state-of-the-art accuracy against previous frame-to-frame VO methods on the TartanAir and KITTI datasets, while also demonstrating highly competitive performance on the EuRoC dataset. DINO-VO is designed to be lightweight and able to run at 72 FPS with less than 1GB memory usage on a single GPU. In addition, we have addressed the problem of directly utilizing DINOv2 coarse features in sparse image matching as highlighted in [16].

The main limitation of our approach is that it relies on frame-to-frame pose estimation such that it is more prone to drift compared to multi-frame VO methods. Nevertheless, our monocular method is closing the accuracy gap with multi-frame VO methods on TartanAir and EuRoC, and remains highly competitive even against top-performing stereo systems like MAC-VO when considering its significant advantages in efficiency and memory usage. Noticeably, our method outperforms multi-frames VO methods on KITTI dataset. Future work will investigate integrating DINO-VO on multi-frames VO schemes such as using bundle adjustment. These improvements could also benefits from metric depth estimation developed in [14] which also utilizes DINOv2 encoder.

## IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

## REFERENCES

- [1] C. Campos, R. Elvira, J. J. G. Rodríguez, J. M. Montiel, and J. D. Tardós, “Orb-slam3: An accurate open-source library for visual, visual-inertial, and multimap slam,” *IEEE Transactions on Robotics*, vol. 37, no. 6, pp. 1874–1890, 2021.
- [2] J. Engel, V. Koltun, and D. Cremers, “Direct sparse odometry,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 40, no. 3, pp. 611–625, 2017.
- [3] J. Tang, L. Ericson, J. Folkesson, and P. Jensfelt, “Gcnv2: Efficient correspondence prediction for real-time slam,” *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 3505–3512, 2019.
- [4] O. Ilter, I. Armeni, M. Pollefeys, and D. Barath, “Semantically guided feature matching for visual slam,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 12013–12019, IEEE, 2024.
- [5] S. Wang, R. Clark, H. Wen, and N. Trigoni, “Deepvo: Towards end-to-end visual odometry with deep recurrent convolutional neural networks,” in *2017 IEEE international conference on robotics and automation (ICRA)*, pp. 2043–2050, IEEE, 2017.
- [6] Z. Teed and J. Deng, “Deepv2d: Video to depth with differentiable structure from motion,” in *8th International Conference on Learning Representations, ICLR 2020*, 2020.
- [7] W. Wang, Y. Hu, and S. Scherer, “Tartanvo: A generalizable learning-based vo,” in *Conference on Robot Learning*, pp. 1761–1772, PMLR, 2021.
- [8] Z. Teed and J. Deng, “Droid-slam: Deep visual slam for monocular, stereo, and rgb-d cameras,” *Advances in neural information processing systems*, vol. 34, pp. 16558–16569, 2021.
- [9] Z. Teed, L. Lipson, and J. Deng, “Deep patch visual odometry,” *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [10] L. Lipson, Z. Teed, and J. Deng, “Deep Patch Visual SLAM,” in *European Conference on Computer Vision*, 2024.
- [11] M. Caron, H. Touvron, I. Misra, H. Jégou, J. Mairal, P. Bojanowski, and A. Joulin, “Emerging properties in self-supervised vision transformers,” in *Proceedings of the International Conference on Computer Vision (ICCV)*, 2021.
- [12] M. Oquab, T. Darcet, T. Moutakanni, H. V. Vo, M. Szafraniec, V. Khalidov, P. Fernandez, D. HAZIZA, F. Massa, A. El-Nouby, M. Assran, N. Ballas, W. Galuba, R. Howes, P.-Y. Huang, S.-W. Li, I. Misra, M. Rabbat, V. Sharma, G. Synnaeve, H. Xu, H. Jegou, J. Mairal, P. Labatut, A. Joulin, and P. Bojanowski, “DINOv2: Learning robust visual features without supervision,” *Transactions on Machine Learning Research*, 2024.
- [13] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, “An image is worth 16x16 words: Transformers for image recognition at scale,” in *International Conference on Learning Representations*, 2021.
- [14] L. Yang, B. Kang, Z. Huang, X. Xu, J. Feng, and H. Zhao, “Depth anything: Unleashing the power of large-scale unlabeled data,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10371–10381, 2024.
- [15] J. Edstedt, Q. Sun, G. Bökman, M. Wadenbäck, and M. Felsberg, “Roma: Robust dense feature matching,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 19790–19800, 2024.
- [16] H. Jiang, A. Karpur, B. Cao, Q. Huang, and A. Araujo, “Omniglu: Generalizable feature matching with foundation model guidance,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 19865–19875, 2024.
- [17] N. Tumanyan, A. Singer, S. Bagon, and T. Dekel, “Dino-tracker: Taming dino for self-supervised point tracking in a single video,” 2024.
- [18] N. Keetha, A. Mishra, J. Karhade, K. M. Jatavallabhula, S. Scherer, M. Krishna, and S. Garg, “Anyloc: Towards universal visual place recognition,” *IEEE Robotics and Automation Letters*, 2023.
- [19] P. Lindenberger, P.-E. Sarlin, and M. Pollefeys, “Lightglue: Local feature matching at light speed,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 17627–17638, 2023.
- [20] Y. Wang, X. He, S. Peng, D. Tan, and X. Zhou, “Efficient lofr: Semi-dense local feature matching with sparse-like speed,” *arXiv preprint arXiv:2403.04765*, 2024.
- [21] B. Roessle and M. Nießner, “End2end multi-view feature matching with differentiable pose optimization,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 477–487, 2023.
- [22] N. Yang, L. v. Stumberg, R. Wang, and D. Cremers, “D3vo: Deep depth, deep pose and deep uncertainty for monocular visual odometry,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 1281–1292, 2020.
- [23] T. Fu, S. Su, Y. Lu, and C. Wang, “islam: Imperative slam,” *IEEE Robotics and Automation Letters*, 2024.
- [24] C. M. Parameashwara, G. Hari, C. Fermüller, N. J. Sanket, and Y. Aloimonos, “Diffposenet: Direct differentiable camera pose estimation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6845–6854, 2022.
- [25] L. Lai, Z. Shangquan, J. Zhang, and E. Ohn-Bar, “Xvo: Generalized visual odometry via cross-modal self-training,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 10094–10105, 2023.
- [26] Z. Teed and J. Deng, “Raft: Recurrent all-pairs field transforms for optical flow,” in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, pp. 402–419, Springer, 2020.
- [27] S. Wang, W. Li, Y. Wang, Z. Fan, Z. Huang, X. Cai, J. Zhao, and D. Li, “Mambavo: Deep visual odometry based on sequential matching refinement and training smoothing,” in *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR)*, pp. 1252–1262, June 2025.
- [28] D. DeTone, T. Malisiewicz, and A. Rabinovich, “Superpoint: Self-supervised interest point detection and description,” in *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pp. 224–236, 2018.
- [29] G. Potje, F. Cadar, A. Araujo, R. Martins, and E. R. Nascimento, “Xfeat: Accelerated features for lightweight image matching,” in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.
- [30] P.-E. Sarlin, D. DeTone, T. Malisiewicz, and A. Rabinovich, “Superglue: Learning feature matching with graph neural networks,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 4938–4947, 2020.
- [31] J. Su, M. Ahmed, Y. Lu, S. Pan, W. Bo, and Y. Liu, “Roformer: Enhanced transformer with rotary position embedding,” *Neurocomputing*, vol. 568, p. 127063, 2024.
- [32] R. Hartley and A. Zisserman, *Multiple view geometry in computer vision*. Cambridge university press, 2003.
- [33] H. Zhan, C. S. Weerasekera, J.-W. Bian, and I. Reid, “Visual odometry revisited: What should be learnt?,” in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4203–4210, 2020.
- [34] W. Wang, D. Zhu, X. Wang, Y. Hu, Y. Qiu, C. Wang, Y. Hu, A. Kapoor, and S. Scherer, “Tartanair: A dataset to push the limits of visual slam,” in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 4909–4916, IEEE, 2020.
- [35] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, “Vision meets robotics: The kitti dataset,” *The International Journal of Robotics Research*, vol. 32, no. 11, pp. 1231–1237, 2013.
- [36] C. Forster, Z. Zhang, M. Gassner, M. Werlberger, and D. Scaramuzza, “Svo: Semidirect visual odometry for monocular and multicamera systems,” *IEEE Transactions on Robotics*, vol. 33, no. 2, pp. 249–265, 2016.
- [37] Y. Qiu, Y. Chen, Z. Zhang, W. Wang, and S. Scherer, “Mac-vo: Metrics-aware covariance for learning-based stereo visual odometry,” in *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 3803–3814, 2025.
- [38] Z. Liu, E. Malis, and P. Martinet, “A new dense hybrid stereo visual odometry approach,” in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 6998–7003, 2022.
- [39] A. Bangunharcana, A. Magd, and K.-S. Kim, “Dualrefine: Self-supervised depth and pose estimation through iterative epipolar sampling and refinement toward equilibrium,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 726–738, June 2023.
- [40] Z. Liu, E. Malis, and P. Martinet, “Adaptive learning for hybrid visual odometry,” *IEEE Robotics and Automation Letters*, vol. 9, no. 8, pp. 7341–7348, 2024.
- [41] X. Gao, R. Wang, N. Demmel, and D. Cremers, “Ldso: Direct sparse odometry with loop closure,” in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 2198–2204, IEEE, 2018.
- [42] Z. Zhan, D. Gao, Y.-J. Lin, Y. Xia, and C. Wang, “ismatching: Imperative correspondence learning,” in *European Conference on Computer Vision*, pp. 183–200, Springer, 2024.