GyF

This ACCV 2024 paper, provided here by the Computer Vision Foundation, is the author-created version. The content of this paper is identical to the content of the officially published ACCV 2024 LNCS version of the paper as available on SpringerLink: https://link.springer.com/conference/accv

Leveraging Semantic Cues from Foundation Vision Models for Enhanced Local Feature Correspondence

Felipe Cadar^{1,2}, Guilherme Potje¹, Renato Martins^{2,3}, Cédric Demonceaux^{2,3}, and Erickson R. Nascimento^{1,4}

 ¹ Computer Science Department, Universidade Federal de Minas Gerais, Brazil {cadar, guipotje, erickson}@dcc.ufmg.br
 ² ICB UMR CNRS 6303, Université de Bourgogne, France

{renato.martins,cedric.demonceaux}@u-bourgogne.fr ³ Inria, LORIA, CNRS, Université de Lorraine, France ⁴ Microsoft

Abstract. Visual correspondence is a crucial step in key computer vision tasks, including camera localization, image registration, and structure from motion. The most effective techniques for matching keypoints currently involve using learned sparse or dense matchers, which need pairs of images. These neural networks have a good general understanding of features from both images, but they often struggle to match points from different semantic areas. This paper presents a new method that uses semantic cues from foundation vision model features (like DI-NOv2) to enhance local feature matching by incorporating semantic reasoning into existing descriptors. Therefore, the learned descriptors do not require image pairs at inference time, allowing feature caching and fast matching using similarity search, unlike learned matchers. We present adapted versions of six existing descriptors, with an average increase in performance of 29% in camera localization, with comparable accuracy to existing matchers as LightGlue and LoFTR in two existing benchmarks. Both code and trained models are available at https://www.verlab.dcc.ufmg.br/descriptors/reasoning accv24/

Keywords: Image Correspondence \cdot Local Features \cdot Keypoint Detection and Description \cdot Semantic Cues \cdot Foundation Vision Models

1 Introduction

Visual correspondence is fundamental for important higher-level vision tasks like camera pose estimation, simultaneous localization and mapping (SLAM), and structure from motion (SfM). Recently, the pipeline for finding visual correspondences between pairs of images has been changing in favor of methods that provide different types of context aggregation, like learned sparse matchers [17,10] or dense correspondence networks [19]. These methods depend on gathering information from both perspectives to condition features for better



Fig. 1. Leveraging semantic information for improving visual correspondence. The figure illustrates the matching process using Mutual Nearest Neighbor (MNN) for the base descriptor XFeat [12] and for our approach, which employs semantic conditioning (shown in the top right). Correct matches are shown in green and wrong matches in red. We can also assess the interpretability and consistency of the descriptors by finding the closest 128 matches to a given query point in the image (red point in the bottom left) using either semantic or texture features. Hotter colors indicate higher similarities. Please notice the similarity ranking improvement with the conditioned features around the sink region.

correspondence prediction. Although they have been shown to provide better results in downstream tasks, they must be run for every pair of images, making it expensive to use in large tasks like SfM pipelines, where a single image will be matched multiple times to other images with a similar viewpoint. While the traditional single-view pipeline can pre-extract features for individual images and use an efficient similarity search such as mutual nearest neighbor (MNN), it does not perform as well as the context aggregation methods. This paper proposes an approach to semantically condition keypoint descriptors to find better and more consistent correspondences while maintaining the advantages of single-view extraction and caching. Foundations Models, such as DINOv2 [11] and SAM [7], can extract features that contain an understanding of semantic concepts in a scene to complement local texture patterns. These features can be adapted for various tasks, e.g., image classification, instance retrieval, video understanding, depth estimation, semantic segmentation, and semantic matching, by freezing the backbone and training new layers for a specific task [11]. In order to capture the meaning of scenes and objects, models such as DINOv2 have developed a strong invariance to local texture changes. However, the high level of invariance in these features can make them less sensitive when it comes to identifying pixel-level matches between images. Instead, they can offer a basis for agreement between regions, which can be used to filter connections between visually similar but semantically different regions, as shown in Fig. 1. In this paper, instead of relying on two-view context aggregation, we propose an effective technique to

leverage high-level features understanding coming from an LVM to semantically condition textured-based correspondences.

A key technical contribution of this paper is a novel learning-based method for integrating semantic context into local features, enabling efficient similarity search during matching and substantially improving matching accuracy. The experimental results demonstrate that our approach significantly enhances the performance of various detect-and-describe techniques in camera pose estimation and visual localization tasks within indoor environments.

2 Related Works

Semantics and object-level in description. Leveraging semantic features for their high invariance is not unprecedented. In nonrigid image matching, methods like GeoPatch [14], DEAL [15], and DALF [13] merge geometric and texture features to achieve invariance to deformations. Sim2Real [25] is a detector and describe approach that learns to find correspondences using two types of losses: the inter-objects, and the intra-objects, segmenting its features into two parts. The first part is responsible for finding the right object, and the second part matches a precise location within that object. Although it captures the concept of separating invariance and distinctiveness, its features are trained specifically for single object-matching. Our approach introduces a much more general formulation, significantly improving camera pose estimation in indoor scenes. SFD2 [21] also introduced a more general approach by explicitly distilling semantic segmentation features. To improve long-term outdoor visual localization, it uses the semantic cues from an off-the-shelf semantic segmentation model to classify keypoints in four levels of stability, including Volatile, Dynamic, Shortterm, and Long-term. The descriptor learning uses the same semantic backbone as the detector, and similar to Sim2Real, it optimizes both intra and inter-class distances. Although SFD2 proposes a more general matching descriptor, it still fixes the classes of keypoints that are reliable and depends on fixed semantic segmentation labels for learning descriptors. Our proposed network is different since we do not explicitly define any semantic behavior based on classes, making it more general and easier to train while also leveraging foundation vision models, which provides a fine-grained and more general semantic understanding of the scene. DeDoDe descriptor [4] is the closest competitor to our methodology and optimizes a dual-softmax matching loss while incorporating DINOv2 [11] features in the extraction pipeline. Although it uses the same semantic cues as our methodology, DeDoDe descriptors do not perform any aggregation between features. In our proposed model, we design an attention mechanism to refine the description based on the semantic cues of the scenes, in addition to learning a semantic-only feature vector for improving the matching.

Context aggregation. A recent trend in image correspondence for finding more robust matchings between views is to use information from both views to predict an assignment. The learned sparse matchers SuperGlue [17] and more

recently LightGlue [10] receive keypoints and descriptors from two images from a base extractor method, like SuperPoint [3] and predict an assignment of the keypoints. Using self and cross-attention, these networks iteratively refine the descriptors into "matching descriptors". In this process, every descriptor is aware of the other descriptors from its own image and from the corresponding pair view. The motivation is that the attention aggregation process can gather enough context information to estimate a more consistent and robust matching while discarding keypoints that are not visible in both images. Recently, Omniglue sparse matcher [6] extended LightGlue by adding semantic features from DINOv2 into the pipeline, still maintaining the two-view aggregation.

Methods like LoFTR [19] and EfficientLoFTR [20] adopt linear approximations of the attention aggregation and a coarse-to-fine strategy to predict an assignment, respectively. In the same direction, MESA [22] also uses a detectionfree matching pipeline, but including higher-level features from the Segment Anything Model [7]. Overall, methods that aggregate context from both views, either dense or sparse, cannot fully cache features between images.

This aggregation strategy in large-scale reconstruction can be expensive since the number of image pairs to be matched can grow quadratically. Our proposed method, on its turn, only uses a single view for the feature extraction and aggregation, allowing frameworks to cache features for later use in an efficient similarity search operation.

3 Methodology

In this section, we present the main concept of our methodology, detailing how to add semantic awareness to a local descriptor and the supervision designed to train it. A scheme of the overall training and inference stages is shown in Fig. 2. The proposed strategy first extracts two sets of descriptors: a set of **texture** features using an off-the-shelf local feature method, and the **semantic** features coming from a LVM model for context (as the selected DINOv2 in this work). To that end, we adopt a base method for extracting traditional, textured-focused features and a base method for extracting semantic-focused features. Following the base extraction, we refine the features using a self-attention *Reasoning* module. For finding matching image pairs, we use the two sets of **texture** and **semantic** features, extracted independently for each image, to calculate a similarity matrix for finding mutual matches using the *Semantic Conditioning*.

Base Feature Extraction. We extract two sets of base features: the texturebased (D_t^*) and the semantic-based (D_s^*) descriptors. For extracting texturebased descriptors and keypoints, we use an off-the-shelf detection and description method, like XFeat [12] or Superpoint [3]. Following, we extract a dense feature map from DINOv2 and sample the semantic-based descriptors at the positions of the detected keypoints via bicubic interpolation. Finally, we project both sets of descriptors to a shared dimension size, as shown in the top left side of Fig. 2 (*Base Extraction*).



Fig. 2. Semantic Conditioning Pipeline. Our method first extracts both low-level scene texture via a base local feature descriptor (XFeat) and semantically meaning-ful, high-level features via a foundation vision model (DINOv2) with the associated salient texture keypoints. Then, the Reasoning module is applied in both representations, where cross attention layers are used iteratively to enhance the representations of both texture and semantic features. Finally, the descriptor similarity is computed by combining both the texture and semantic similarity using element-wise product (\odot) .

Descriptor Reasoning. At this point, using the two sets of features (*i.e.*, texture-based and semantic-based) we refine them for the *Semantic Conditioning* stage. We initialize the refined texture descriptor (D_t) and the refined semantic descriptor (D_s) using their raw versions D_s^* and D_t^* . We iteratively refine D_t using attention aggregation while alternating the attention keys between the raw descriptors D_s^* and D_t^* . For each attention layer $\operatorname{aggr}_i(K, Q, V)$, we compute:

$$D_{t_{i+1}} = \operatorname{aggr}_i(D_{\mathbf{s}}^*, D_{t_i}, D_{t_i}) \quad \text{if } i \text{ is even},$$

$$D_{t_{i+1}} = \operatorname{aggr}_i(D_{\mathbf{t}}^*, D_{t_i}, D_{t_i}) \quad \text{if } i \text{ is odd},$$

$$(1)$$

where *i* is the number of the iteration, *K* the key, *Q* the query, *V* the value, and D_{t_i} and $D_{t_{i+1}}$ are the refined texture descriptors at the iterations *i* and *i*+1. The attention aggregation mechanism is the same as the Lightglue's cross-attention layer [10]. Similarly, we iteratively refine D_s , using only the original semantic features as keys:

$$D_{s_{i+1}} = \operatorname{aggr}_i(D_{\mathbf{s}}^*, D_{s_i}, D_{s_i}).$$

$$(2)$$

The strategy of using alternating texture and semantic keys aims to incorporate both types of information, providing the texture descriptor with scene context from the semantic descriptors. This approach enables the texture component to identify precise matches and to understand the type of structure it is describing, resulting in more semantically relevant aggregated information. Similarly, we do not want the refined semantic features to have a local understanding of the scene. The objective of this reasoning pipeline is to guide the matches for semantically

coherent areas. Finally, both sets of descriptors are normalized to the unit hypersphere with the L_2 normalization.

Semantic Conditioning. After obtaining D_t , D_s , and the base extractor's keypoints K, we cache them all for later use since there are no more steps before the visual matching. The extraction from a pair of images I^1 and I^2 will result in D_t^1 , D_s^1 , K^1 (descriptors and keypoints for the first image), and D_t^2 , D_s^2 , K^2 (descriptors and keypoints for the second image). With each pair of descriptor types, we can calculate a correlation matrix C_t and C_s , for texture and semantics, respectively. Individually, D_t and D_s are either too local or too global, resulting in discriminative descriptors that can wrongly match repetitive textures or highly invariant descriptors that can only find coarse matches. We define the Semantic Conditioning as the operation of combining both correlation matrices for adjusting the match probability between coherent regions while keeping the local discriminatory power. The final correlation matrix considering the product of features is then

$$C_f = C_t \odot C_s, \tag{3}$$

where \odot is the element-wise product operation. C_f is then used for finding matches with Mutual Nearest Neighbor (MNN) search. We conditioned the matching using the product operation so that even low semantic similarities can completely erase wrong matches, regardless of high texture similarity.

3.1 Supervision

For training, the parameters of the base extractors remain frozen and we optimize only the weights of the initial projections and the descriptor reasoning as depicted in Fig. 2. We freeze the weights because each base extractor can have a specific training strategy that works better. By using freezed, off-the-shelf extractors, we can accommodate more methods. The DINOv2 was also frozen following [11], which used it as a backbone for several tasks.

We trained the networks using the ScanNet [2] dataset, which contains exclusively indoor scenes. We focused on indoor environments because they are more semantically rich than outdoor scenes. We sample image pairs with a minimum estimated overlap of 40%, along with the depth maps and camera poses. For every pair, we use the base extractor to estimate keypoints from both images, project the keypoints from the first one to the second image, and find matches that are under three pixels distance from a correspondent keypoint. These ground-truth assignments are obtained for calculating the dual-softmax loss, following a similar strategy to XFeat [12]. Since our descriptors are refined iteratively, we can follow the supervision process of LightGlue [10] and calculate the dual-softmax loss for each iteration of the descriptors to seep up convergence. Assuming the ground-truth assignment M_{qt} of size $M \times 2$ containing the matching indices, we

calculate the loss \mathcal{L}_l for the layer l as:

$$\mathcal{L}_{l} = -\sum_{(i,j)\in M_{gt}} \log(\operatorname{softmax}_{r}(C_{f_{l}})_{(i,j)}) -\sum_{(i,j)\in M_{gt}} \log(\operatorname{softmax}_{r}(C_{f_{l}}^{\top})_{(j,i)}),$$
(4)

where softmax_r is a row-wise softmax function and the correlation matrix C_{f_l} for the layer l is computed as:

$$C_{t_l} = D_{t_l}^1 D_{t_l}^{2\top}, \quad C_{s_l} = D_{s_l}^1 D_{s_l}^{2\top}, \quad C_{f_l} = C_{t_l} \odot C_{s_l}.$$
 (5)

Following the deep supervision of LightGlue, the final loss is the mean of the intermediary layers' losses. For N layers, the final loss $\mathcal{L}_{\text{total}}$ is then:

$$\mathcal{L}_{\text{total}} = \frac{1}{N} \sum_{l \in [1,N]} \mathcal{L}_l.$$
(6)

4 Experiments

Experimental setup and training. We trained the method in approximately 1,000,000 pairs from the ScanNet dataset [2], with a fixed image size of 512×512 , extracting 2,048 keypoints and descriptors on each image. The training takes 10 hours on $4 \times V100$ 32GB GPUs with a batch size of 16 on each GPU. We used 5 layers of attention aggregation. The experiment with 9 layers used a batch size of 8 but kept the same total training paris. We used Adam with learning rate $1e^{-4}$ as the optimizer and the descriptor size is fixed to 256. We have adopted DINOv2 as the semantics provider as it has been shown to provide rich semantic features that generalize well in different benchmarks [4,6,11].

For the DINOv2 version, we tested all four types and settled on the smallest one (S), as discussed in Sec. 4.3. For the evaluation, we have adopted the Scannet1500 relative pose estimation benchmark, which inclides 1,500 image pairs from the test split of ScanNet, and the 7Scenes [5] visual localization benchmark, which contains 7 indoor scenes. Both benchmarks pose challenges for the visual correspondence task, as they contain motion blur, repetitive structures, and low-textured surfaces.

Inference and computational analysis. In inference time, we resize input images to 896 pixels on the longest edge for extracting DINOv2 features, but we keep the original shape for the base extractor. For all models, we normalize the output descriptors. All experiments were run without resizing the images. For a pair of images in the original ScanNet resolution of 968×1296 , and with a max keypoint detection count of 2,048, our network takes, on average, 23ms for texture features extraction, 83ms for semantic features extraction, 4ms for reasoning, and 11ms for matching, using XFeat as the base local extractor, for a total of 121ms, in which 110ms are processed in single images and can be cached. All experiments were executed in a NVIDIA GeForce RTX 3080 with 10GB.

7

Table 1. Pose estimation results compared to global visual matchers and existing single view competitors. The semantic conditioning yields results that are even competitive with learned matchers and the detector-free matcher LoFTR (half upper part of the table), as well as to existing single view descriptors with semantics (our direct competitors shown in the bottom part).

		AUC \uparrow	
Method	@5°	@10°	@20°
LoFTR	20.20	37.60	52.60
LightGlue	22.30	40.80	57.00
DeDoDe-G	9.50	19.40	30.20
SFD2	13.00	25.60	38.30
Ours	20.30	36.81	52.37

Table 2. Semantic conditioning capability of our method for different base descriptors. "Ours" accompanying a method's name means that we trained our network with the features of this base method. Pose estimation results on Scannet1500.

Method	@5°	AUC ↑ @10°	@20°	Average Gain
DeDoDe-B	5.70	11.30	17.80	+50.67%
Ours+DeDoDe-B	8.91	17.28	26.25	
DeDoDe-G	9.50	19.40	30.20	+9.07%
Ours+DeDoDe-G	11.13	21.29	32.04	
ALIKE	7.60	15.40	23.90	+35.27%
ALIKE Ours	10.56	20.96	31.92	
ALIKED	12.00	23.40	35.80	+24.37%
Ours+ALIKED	15.31	29.39	43.85	
$egin{array}{c} \operatorname{RELF} \\ \operatorname{Ours+RELF} \end{array}$	11.60 15.20	24.20 30.46	33.30 46.31	+24.12%
XFeat	15.80	30.80	45.80	+18.70%
Ours+XFeat	19.72	36.77	53.19	
SuperPoint	13.70	26.30	39.90	+37.02%
Ours+SuperPoint	20.30	36.81	52.37	

Baselines. In the selected baselines, we have traditional detect and describe methods as ALIKE [24], ALIKED [23], SuperPoint [3], RELF [8] and XFeat [12]. We selected DeDoDe [4], as the only other descriptor that uses semantic information from DINO, and SFD2 [21] that explicitly uses semantic information for supervision in training. Although our proposed method is not a matcher, we also included results for the learned matcher LightGlue [10] and the detector-free matcher LoFTR [19] for comparing methods that can aggregate information from both views.

9

4.1 Relative Pose Estimation

Setup. Following [19], we evaluated pose estimation in 1500 pairs of images from the validation split of the ScanNet [2] dataset. The indoor images contain strong viewpoint changes, poorly textured surfaces, and motion blur artifacts, making it challenging for the pose estimation task. As in previous works [12], we use LO-RANSAC [1] to estimate the essential matrix and search for the optimal threshold for each method.

Metrics. We follow the protocol of previous works [12,10,19], which report the area under the curve (AUC) of the maximum angular error in rotation and translation at the thresholds 5°, 10°, and 20°.

Results. The quantitative registration results are shown in Tab. 1 and Tab. 2. In Table 1, we included methods of dual-view context aggregation, like Light-Glue [10] and LoFTR [19], and other descriptors that also leverage semantic information like DeDoDe-G [4] and SFD2 [21]. It is worth noting that even doing only single view extraction, Superpoint combined with semantic conditioning can yield competitive results compared to Lightglue (without any awareness of the pairing view). Our designed strategy to leverage semantics to improve the matching capability of existing descriptors is described in Tab. 2. We can notice a significant improvement in all baselines when combined with our proposed semantic conditioning. Although many of them were only trained on outdoor images from MegaDepth dataset [9], like DeDoDe [4], SFD2 [21], ALIKE [24], and ALIKED [23], we could still improve their indoor pose estimation results by at least 24% without retraining neither the feature extraction nor the DINOv2 backbone. This result suggests that the extracted visual cues are not better than the original versions of these descriptors, only better conditioned by the semantics.

4.2 Visual Localization

Setup. We evaluated the visual localization task on the popular benchmark 7Scenes [5]. The 7Scenes dataset comprises seven different indoor scenes with annotated 6-DoF camera poses and dense 3D reconstructions. It is divided into training and test sequences, each with a distinct camera path. The images are acquired using a Kinect 1 and are often blurry, with several regions featuring repetitive indoor structures and textureless areas. We use the HLoc [16] framework for building a Structure-from-Motion (SfM) [18] database from the training sequences. Then, the images in the test set are localized by computing image correspondences between the current frame and the SfM models using the descriptors provided by each competitor.

Metrics. As a standard practice [16], for each localized image from the test set, we use the camera translation error in meters, and rotation errors in degrees, according to the ground-truth camera poses provided by the benchmark [5].

Table 3. Visual localization results in 7Scenes. For each scene we report the camera pose error achieved by the competing methods. We also include LightGlue as a learned matcher reference, providing an upper bound in matching quality.

	Translation errors (m) \downarrow , Rotation errors (°) \downarrow						
Methods \downarrow	Chess	Fire	Heads	Office	Pumpkin	Kitchen	Stairs
LightGlue	0.024, 0.803	0.019, 0.792	0.011, 0.715	0.027, 0.826	0.039, 1.044	0.033, 1.127	0.051, 1.324
DeDoDe-G	0.025, 0.902	0.018 , 0.793	0.011, 0.810	$0.\overline{0}2\overline{8}, 0.\overline{8}92$	0.046, 1.357	0.035, 1.250	$0.\overline{0}\overline{69}, 1.7\overline{0}8$
XFeat	0.027, 0.919	0.018 , 0.731	0.013, 0.865	0.028, 0.842	0.042, 1.072	0.035, 1.189	0.043, 1.108
Ours+XFeat	0.026, 0.913	0.018, 0.725	0.013, 0.881	0.027 , 0.840	0.043, 1.111	0.035, 1.215	0.048, 1.108
SuperPoint	0.025, 0.822	0.018 , 0.738	0.011, 0.759	0.027, 0.820	0.040, 1.052	0.034, 1.130	0.055, 1.427
Ours+SuperPoint	0.026, 0.876	0.019, 0.755	0.012, 0.802	0.027 , 0.826	0.041, 1.052	0.034 , 1.171	0.046, 1.167

Results. The visual localization benchmark results are shown in Tab. 3. An interesting observation is that our method is able to reduce XFeat errors in several cases. For SuperPoint, our approach was not able to provide meaningful improvements. We hypothesize that XFeat being a smaller backbone, provides more concise and less redundant features, making it less prone to overfitting and leverages the semantic information the most. We achieve the highest average percentage of correctly localized cameras when considering the percentage of localized cameras within different thresholds. From stricter thresholds of 1°, 1 cm up to 500 cm, 10°, LightGlue, the gold standard but expensive matcher, correctly localizes 66.97% of the cameras. This is followed by Ours (with Super-Point as base texture detector): 66.95%, Superpoint: 66.88%, XFeat: 66.36%, and DeDoDe-G: 64.44%. This demonstrates that semantic information can increase correspondences in ambiguous regions, as shown in Fig. 3.

4.3 Sensitivity and Ablation Studies

In this section, we investigate the design decisions that led to our final semantic condoning method. Every variant in these experiments was evaluated using relative pose estimation on the Scannet1500 benchmark, using XFeat [12] as the base extractor. The choice of XFeat in these experiments was guided since it is a recent lightweight descriptor with competitive performance along different visual correspondence benchmarks.

DINOv2 backbone size. As the DINOv2 model family has many versions with different capacities, they also require increasingly higher amount of computation. Following DeDoDe [4], we fixed the dimension of the input images to DINOv2 to mitigate the fast-growing computational complexity of the vision transformer, which is quadratic in the number of tokens. To study which type of model and image resolution, we trained four versions of our network, one for each DINO model, and evaluated fixing the input DINOv2 image size in inference to 518 and 896 on the long edge. Table 4 shows that, while fixing image size to 518, larger DINOv2 models can yield better pose estimation results; the results difference is smaller for the larger images. We chose to fix the input size to 896 and use

Table 4. Sensitivity analysis to semantic model capacity. We evaluate fixing DINO interval resizing dimension and varying model size.

DINOv2	AUC	$\uparrow @ siz$	e 518	Inference	AUC	$\uparrow @ siz$	e 896	Inference
Size	$@5^{\circ}$	@10°	@20°	time (ms)	@5°	@10°	@20°	time (ms)
S	18.80	36.33	53.11	54.80	19.72	36.77	53.19	121.11
В	19.87	37.57	53.88	88.86	20.70	37.93	54.30	225.96
L	19.54	37.13	53.95	174.54	19.88	37.70	55.00	561.63
G	20.41	38.20	55.22	446.01	20.89	38.58	55.66	1695.00

Table 5. Pose estimation evaluation on Scannet1500. SC means Semantic Conditioning, #AL is the number of aggregation layers, and K is the type of Key used in the texture reasoning. The version used in all experiments is marked in **bold**.

				AUC \uparrow				
Exp	SC	Κ	#AL	@5°	@10°	@20°		
1		Α	5	18.27	34.29	49.80		
2	\checkmark	А	3	18.71	35.57	51.68		
3	\checkmark	Α	5	19.72	36.77	53.19		
4	\checkmark	А	7	19.19	36.12	52.21		
5	\checkmark	А	9	19.35	36.34	52.66		
6	\checkmark	\overline{S}	$^{-}5^{-}$	19.18	36.19	$5\overline{2}.\overline{60}$		
7	\checkmark	Т	5	19.48	36.64	52.58		

the smaller version of DINOv2 as it already yields satisfactory results. This also leads to a lower inference time.

Semantic conditioning. When refining the texture descriptors (D_t) with the semantic keys (D_s^*) , they can also aggregate semantic information. On Experiment 1 of Tab. 5, we tested the capacity of a single descriptor playing both roles of semantic and texture guidance. We removed the semantic refinement and only used the texture correlation (C_t) for matching descriptors. We can see a large drop in performance in all thresholds compared to Experiment 3, which is our default setting.

Number of attention aggregation layers. Another factor that we investigated was the number of layers for attention aggregation. This evaluation follows an observation from LightGlue [10], where the network can collect information from both images to decide if the descriptors are already good for matching or need more refining. In single-view extraction we do not want to depend on other image descriptors, so it is difficult to know when to stop refining dynamically. For that reason, we decided to fix the number of aggregation layers. In Tab. 5, Experiments 2 through 5 evaluate networks trained with 3, 5, 7, and 9 layers. Although they all are significantly close to each other, the biggest gap was between 3 and 5 layers; for that reason, we kept the five layers as default.



Fig. 3. Visual matching results for three refined base descriptors in the two benchmarks. Green matches are inliers and red outlier matches. The left side of the figure shows the results for the original descriptors SuperPoint [3], XFeat [12] and ALIKED [23] when matching two different image pairs. The right side of the figure shows the matching using their semantically conditioned versions using our proposed methodology. We can observe from visual inspection that the matches are more consistent between the views. Please also notice the increased inlier ratio when considering the semantic conditioned versions.

Alternating aggregation key. We also evaluated the impact of alternating keys in the texture descriptor refinement. In Tab. 5, experiments 6 and 7 are versions of the default network that only use the semantic key and the texture key, respectively. We can notice that using only the semantic key is worse than using only the texture key. Since the network still has semantic conditioning in the matching stage, it should leverage the texture information for fine matching. Although the texture only is better than semantic only, alternating between both yields better results because the network can adjust the texture features while aggregating context from the semantics of the scene present in the image.

4.4 Qualitative Matching Analysis

The reason for using semantic cues to find matches is to reduce ambiguity between descriptors from semantically different areas by using both sets of infor-

13

mation. This makes it less likely to mistakenly match a similar texture to the wrong instance. As can be noticed in Fig. 3 with examples from the two considered benchmarks, and without any filtering, our method leverages the semantic information for a much more consistent matching between the views. By using semantical conditioning, we not only filter semantically wrong matches, but we find a greater number of correct correspondences between coherent areas.

Interpretability of the features. We can also check the interpretability and consistency of the descriptors by computing the closest 128 matches to a given query point in the image using either solely the semantic, the texture or the semantically conditioned features as shown in Figs. 1 and 4. Given a query keypoint (highlight in red in the first column) we find the top 128 corresponding keypoints ranked by feature correlation. The semantic correlation selects only keypoints in a similar area, but without local precision. This can be seen in Fig. 1 where the areas belonging to the sink are selected as the closest ones. The same behavior is displayed in the top example in Fig. 4, where the query point located in the cable also activates points located along the entire cable. On the other hand, the texture description correlation selects keypoints locally similar but has not regard for the context of the location (as shown in the third column in Figs. 1 and 4). The conditioned features can find correspondences that are relevant both in semantics and texture. In some cases, we can have high texture and semantic correlations for non corresponding keypoints, for instance when the correct correspondence is not visible, as shown in the last row of Fig. 4 where the selected keypoint is not visible in the pairing view. Yet, the closest matches in this case are the ones located in the two other handles. Other representative visualization examples of the correlations for four other scenes are provided in the supplementary material.

5 Conclusion

This work introduces a learning-based technique for visual feature description that is capable of utilizing semantic cues present in the image. We design a network performing information aggregation that leverages semantic features to refine and condition off-the-shelf descriptors to improve indoor visual matching. Our method performance in camera pose estimation is superior to existing state-of-the-art models exploring semantic cues, and it is also competitive even with that of recent learned matchers (such as LightGlue), while only using a single image for feature extraction and Nearest Neighbor search for matching. With extensive experiments, we show that our method can improve the pose estimation results of six different base descriptors by 25% on average. The improved descriptors can be used in large-scale SfM reconstructions by utilizing a single view for image extraction, as MNN matching is much faster than running learned matchers for thousands of image pairs.

7Scenes Benchmark

Query LocationSemanticsTextureConditioned FeaturesImage: SemanticsImage: SemanticsImage:

Scannet1500 Benchmark



Fig. 4. Interpretability and consistency of the conditioned features. We show the closest 128 matches to a given query keypoint (red point in the first column) for the different descriptors with either solely semantics, the refined texture descriptor or with the proposed semantic conditioned features (fourth column). Hotter colors indicate higher similarity. Please notice the similarity ranking improvement with the conditioned features for finding matches such as in the mouse cable (first row). Our approach consistency is highlighted in the estimated closest keypoints when selecting the drawer handle of the kitchen (fourth row) which is occluded in the paired view.

Acknowledgments. This work was partially supported by grants from CAPES, CNPq, FAPEMIG, Google, ANER MOVIS from Conseil Régional BFC and ANR (ANR-23-CE23-0003-01), to whom we are grateful. This project was also provided with AI computing and storage resources by GENCI at IDRIS thanks to the grant 2024-AD011015289 on the supercomputer Jean Zay's V100 partitions.

References

- 1. Chum, O., Matas, J., Kittler, J.: Locally optimized ransac. In: Pattern Recognition. Springer Berlin Heidelberg (2003)
- Dai, A., Chang, A.X., Savva, M., Halber, M., Funkhouser, T., Niessner, M.: Scannet: Richly-annotated 3d reconstructions of indoor scenes. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (July 2017)
- DeTone, D., Malisiewicz, T., Rabinovich, A.: Superpoint: Self-supervised interest point detection and description. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW). pp. 224–236 (2018)
- Edstedt, J., Bökman, G., Wadenbäck, M., Felsberg, M.: DeDoDe: Detect, Don't Describe — Describe, Don't Detect for Local Feature Matching. In: 2024 International Conference on 3D Vision (3DV). IEEE (2024)
- Glocker, B., Izadi, S., Shotton, J., Criminisi, A.: Real-time rgb-d camera relocalization. In: International Symposium on Mixed and Augmented Reality (ISMAR). IEEE (October 2013)
- Jiang, H., Karpur, A., Cao, B., Huang, Q., Araujo, A.: Omniglue: Generalizable feature matching with foundation model guidance. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 19865– 19875 (June 2024)
- Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., Dollár, P., Girshick, R.: Segment anything. arXiv:2304.02643 (2023)
- 8. Lee, J., Kim, B., Kim, S., Cho, M.: Learning rotation-equivariant features for visual correspondence. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2023)
- Li, Z., Snavely, N.: Megadepth: Learning single-view depth prediction from internet photos. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
- Lindenberger, P., Sarlin, P.E., Pollefeys, M.: LightGlue: Local Feature Matching at Light Speed. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2023)
- Oquab, M., Darcet, T., Moutakanni, T., Vo, H.V., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., Howes, R., Huang, P.Y., Xu, H., Sharma, V., Li, S.W., Galuba, W., Rabbat, M., Assran, M., Ballas, N., Synnaeve, G., Misra, I., Jegou, H., Mairal, J., Labatut, P., Joulin, A., Bojanowski, P.: Dinov2: Learning robust visual features without supervision (2023)
- Potje, G., , Cadar, F., Araujo, A., Martins, R., Nascimento, E.R.: Xfeat: Accelerated features for lightweight image matching. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2024)
- Potje, G., Cadar, F., Araujo, A., Martins, R., Nascimento, E.R.: Enhancing deformable local features by jointly learning to detect and describe keypoints. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2023)
- Potje, G., Martins, R., Cadar, F., Nascimento, E.R.: Learning geodesic-aware local features from rgb-d images. Computer Vision and Image Understanding (CVIU) 219, 103409 (Jun 2022)

- 16 F. Cadar et al.
- Potje, G., Martins, R., Chamone, F., Nascimento, E.: Extracting deformationaware local features by learning to deform. Advances in Neural Information Processing Systems (NeurIPS) 34, 10759–10771 (2021)
- Sarlin, P.E., Cadena, C., Siegwart, R., Dymczyk, M.: From coarse to fine: Robust hierarchical localization at large scale. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 12716–12725 (2019)
- Sarlin, P.E., DeTone, D., Malisiewicz, T., Rabinovich, A.: Superglue: Learning feature matching with graph neural networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4938–4947 (2020)
- Schonberger, J.L., Frahm, J.M.: Structure-from-motion revisited. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4104–4113 (2016)
- Sun, J., Shen, Z., Wang, Y., Bao, H., Zhou, X.: Loftr: Detector-free local feature matching with transformers. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 8922–8931 (2021)
- Wang, Y., He, X., Peng, S., Tan, D., Zhou, X.: Efficient loftr: Semi-dense local feature matching with sparse-like speed. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 21666–21675 (June 2024)
- Xue, F., Budvytis, I., Cipolla, R.: Sfd2: Semantic-guided feature detection and description. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2023)
- Zhang, Y., Zhao, X.: Mesa: Matching everything by segmenting anything. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 20217–20226 (June 2024)
- Zhao, X., Wu, X., Chen, W., Chen, P.C.Y., Xu, Q., Li, Z.: Aliked: A lighter keypoint and descriptor extraction network via deformable transformation. IEEE Transactions on Instrumentation and Measurement 72, 1–16 (2023)
- Zhao, X., Wu, X., Miao, J., Chen, W., Chen, P.C., Li, Z.: Alike: Accurate and lightweight keypoint detection and descriptor extraction. IEEE Transactions on Multimedia (2022)
- Zhong, C., Yang, C., Qi, J., Sun, F., Liu, H., Mu, X., Huang, W.: Sim2real objectcentric keypoint detection and description. In: AAAI Conference on Artificial Intelligence (AAAI) (2022)