



HAL
open science

Are Semi-Dense Detector-Free Methods Good at Matching Local Features?

Matthieu Vilain, Rémi Giraud, Hugo Germain, Guillaume Bourmaud

► **To cite this version:**

Matthieu Vilain, Rémi Giraud, Hugo Germain, Guillaume Bourmaud. Are Semi-Dense Detector-Free Methods Good at Matching Local Features?. VISAPP 2024: 19th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications., Feb 2024, Rome, Italy. <hal-04456159>

HAL Id: hal-04456159

<https://hal.science/hal-04456159v1>

Submitted on 13 Feb 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

Are Semi-Dense Detector-Free Methods Good at Matching Local Features?

Matthieu Vilain¹, Rémi Giraud¹, Hugo Germain, and Guillaume Bourmaud¹

¹*Univ. Bordeaux, CNRS, Bordeaux INP, IMS, UMR 5218, F-33400 Talence, France*
{matthieu.vilain, remi.giraud, guillaume.bourmaud}@u-bordeaux.fr

Keywords: Image matching, Transformer, Pose estimation

Abstract: Semi-dense detector-free approaches (SDF), such as LoFTR, are currently among the most popular image matching methods. While SDF methods are trained to establish correspondences between two images, their performances are almost exclusively evaluated using relative pose estimation metrics. Thus, the link between their ability to establish correspondences and the quality of the resulting estimated pose has thus far received little attention. This paper is a first attempt to study this link. We start with proposing a novel structured attention-based image matching architecture (SAM). It allows us to show a counter-intuitive result on two datasets (MegaDepth and HPatches): on the one hand SAM either outperforms or is on par with SDF methods in terms of pose/homography estimation metrics, but on the other hand SDF approaches are significantly better than SAM in terms of matching accuracy. We then propose to limit the computation of the matching accuracy to textured regions, and show that in this case SAM often surpasses SDF methods. Our findings highlight a strong correlation between the ability to establish accurate correspondences in textured regions and the accuracy of the resulting estimated pose/homography. Our code will be made available.

1 Introduction

Image matching is the task of establishing correspondences between two partially overlapping images. It is considered to be a fundamental problem of 3D computer vision, as establishing correspondences is a precondition for several downstream tasks such as structure from motion (Heinly et al., 2015; Schönberger and Frahm, 2016), visual localisation (Svärm et al., 2017; Sattler et al., 2017; Taira et al., 2018) or simultaneous localization and mapping (Strasdat et al., 2011; Sarlin et al., 2022). Despite the abundant work dedicated to image matching over the last thirty years, this topic remains an unsolved problem for challenging scenarios, such as when images are captured from strongly differing viewpoints (Arnold et al., 2022) (Jin et al., 2020), present occlusions (Sarlin et al., 2022) or feature day-to-night changes (Zhang et al., 2021).

An inherent part of the problem is the difficulty to evaluate image matching methods. In practice, it is often tackled with a proxy task such as relative or absolute camera pose estimation, or 3D reconstruction. In this context, image matching performances were recently significantly improved with the advent of attention layers (Vaswani et al., 2017). Cross-

attention layers are mainly responsible for this breakthrough (Sarlin et al., 2020) as they enable the local features of detected keypoints in both images to communicate and adjust with respect to each other. Prior siamese architectures (Yi et al., 2016; Ono et al., 2018; Dusmanu et al., 2019), (Revaud et al., 2019), (Germain et al., 2020; Germain et al., 2021) had so far prevented this type of communication. Shortly after this breakthrough, a second significant milestone was reached (Sun et al., 2021) by combining the usage of attention layers (Sarlin et al., 2020) with the idea of having a detector-free method (Rocco et al., 2018; Rocco et al., 2020a; Rocco et al., 2020b; Li et al., 2020), (Zhou et al., 2021; Truong et al., 2021). In LoFTR (Sun et al., 2021), low-resolution dense cross-attention layers are employed that allow semi-dense low-resolution features of the two images to communicate and adjust to each other. Such a method is said to be detector-free, as it matches semi-dense local features instead of sparse sets of local features coming from detected keypoint locations (Lowe, 1999).

Semi-dense Detector-Free (SDF) methods (Chen et al., 2022; Giang et al., 2023; Wang et al., 2022; Sun et al., 2021), (Mao et al., 2022; Tang et al., 2022), such as LoFTR, are among the best performing image matching approaches in terms of pose estimation met-




Figure 1: Given query locations within textured regions of the source image (left), we show their predicted correspondents in the target image (right) for: **(Top row)** SAM (Proposed) - Structured Attention-based image Matching, **(Bottom row)** LoFTR (Sun et al., 2021)+QuadTree (Tang et al., 2022) - a semi-dense detector-free approach (line colors indicate the distance in pixels with respect to the ground truth correspondent). We report: (MA@2) - the matching accuracy at 2 pixels computed on all the semi-dense locations of the source image with available ground truth correspondent (which includes both textured and uniform regions), (MA_{text}@2) - the matching accuracy at 2 pixels computed on all the textured semi-dense locations of the source image with available ground truth correspondent (*i.e.*, uniform regions are ignored), (errR and errT) - the relative pose error. SAM has a better pose estimation but a lower matching accuracy (MA@2), which seems counter-intuitive. However, if we consider only textured regions (MA_{text}@2), then SAM outperforms LoFTR+QuadTree.

rics. However, to the best of our knowledge, the link between their ability to establish correspondences and the quality of the resulting estimated pose has thus far received little attention. This paper is a first attempt to study this link.

Contributions - We start with proposing a novel Structured Attention-based image Matching architecture (SAM). We evaluate SAM and 6 SDF methods on 3 datasets (MegaDepth - relative pose estimation and matching, HPatches - homography estimation and matching, ETH3D - matching).

We highlight a counter-intuitive result on two datasets (MegaDepth and HPatches): on the one hand SAM either outperforms or is on par with SDF methods in terms of pose/homography estimation metrics, but on the other hand SDF approaches are significantly better than SAM in terms of Matching Accuracy (MA). Here the MA is computed on all the semi-dense locations (of the source image) with available ground truth correspondent, which includes both textured and uniform regions.

We propose to limit the computation of the matching accuracy to textured regions, and show that in this case SAM often surpasses SDF methods. Our find-

ings highlight a strong correlation between the ability to establish accurate correspondences in textured regions and the accuracy of the resulting estimated pose/homography (see Figure 1).

Organization of the paper - Sec. 2 discusses the related work. The proposed SAM architecture is introduced in Sec. 3 and the experiments in Sec. 4.

2 Related work

Since this paper focuses on the link between the ability of SDF methods to establish correspondences and the quality of the resulting estimated pose, we only present SDF methods in this literature review and refer the reader to (Edstedt et al., 2023; Zhu and Liu, 2023; Ni et al., 2023) for a broader literature review.

To the best of our knowledge, LoFTR (Sun et al., 2021) was the first method to perform attention-based detector-free matching. A siamese CNN is first applied on the source/target image pair to extract fine dense features of resolution 1/2 and coarse dense features of resolution 1/8. The source and target coarse features are fed into a dense attention-based module, interleaving self-attention layers with cross-attention layers as proposed in (Sarlin et al., 2020). To reduce the computational complexity of these dense attention layers, Linear Attention (Katharopoulos et al., 2020) is used instead of vanilla softmax attention (Vaswani et al., 2017). The resulting features are matched to obtain coarse correspondences. Each coarse correspondence is then refined by cropping 5×5 windows into the fine features and applying another attention-based module. Thus for each location of the semi-dense (factor of 1/8) source grid, a correspondent is predicted. In practice, a Mutual Nearest Neighbor (MNN) step is applied at the end of the coarse matching stage to remove outliers.

In (Tang et al., 2022), a QuadTree attention module is proposed to reduce the computational complexity of vanilla softmax attention from quadratic to linear while keeping its power, as opposed to Linear Attention (Katharopoulos et al., 2020) which was shown to underperform on local feature matching (Germain et al., 2022). The QuadTree attention module is used as a replacement for Linear Attention module in LoFTR. An architecture called ASpanFormer is introduced in (Chen et al., 2022) that employs the same refinement stage as LoFTR but a different coarse stage architecture. Instead of classical cross-attention layers, the coarse stage uses global-local cross-attention layers that have the ability to focus on regions around current potential correspondences. MatchFormer (Wang et al., 2022) proposes an




Figure 2: **Overview of the proposed Structured Attention-based image Matching (SAM) method.** (a) The matching architecture first extracts features from both source and target images at resolution $1/4$. Then it uses a set of learned latent vectors alongside the descriptors of the query locations, and performs an *input* structured cross-attention with the dense features of the target. The latent space is then processed through a succession of structured self-attention layers. An *output* structured cross-attention is applied to update the target features with the information from the latent space. Finally, the correspondence maps are obtained using a dot product. (b) Proposed structured attention layer. See text for details.

attention-based backbone, interleaving self and cross-attention layers, to progressively transform a tensor of size $2 \times H \times W \times 3$ into a tensor of size $2 \times \frac{H}{16} \times \frac{W}{16} \times C$. Efficient attention layers such as Spatial Efficient Attention (Wang et al., 2021; Xie et al., 2021) and Linear Attention (Katharopoulos et al., 2020) are employed. A feature pyramid network-like decoder (Lin et al., 2017) is used to output a fine tensor of size $2 \times \frac{H}{2} \times \frac{W}{2} \times C_f$ and a coarse tensor of size $2 \times \frac{H}{8} \times \frac{W}{8} \times C_c$, and the matching is performed as in LoFTR. TopicFM (Giang et al., 2023) follows the same global architecture as LoFTR with a different coarse stage. Here, topic distributions (latent semantic instances) are inferred from coarse CNN features using cross-attention layers with topic embeddings. These topic distributions are used to augment the coarse CNN features with self and cross-attention layers. In (Mao et al., 2022), a method called 3DG-STFM proposes to train the LoFTR architecture with a student-teacher method. The teacher is first trained on RGB-D image pairs. The teacher model then guides the student model to learn RGB-induced depth information.

3 Structured attention-based image matching architecture

In the following section, we present our novel Structured Attention-based image Matching architecture (SAM), whose architecture is illustrated in Figure 2.

3.1 Background and notations

Image matching - Given a set of 2D query locations $\{\mathbf{p}_{s,i}\}_{i=1\dots L}$ in a source image I_s , we seek to find their

2D correspondent locations $\{\hat{\mathbf{p}}_{t,i}\}_{i=1\dots L}$ in a target image I_t :

$$\{\hat{\mathbf{p}}_{t,i}\}_{i=1\dots L} = \mathcal{M}(I_s, I_t, \{\mathbf{p}_{s,i}\}_{i=1\dots L}). \quad (1)$$

Here \mathcal{M} is the image matching method. In the case of semi-dense detector-free methods, the query locations $\{\mathbf{p}_{s,i}\}_{i=1\dots L}$ are defined as the source grid locations using a stride of 8. As we will see, SAM is more flexible and can process any set of query locations. However, for a fair comparison, in the experiments all the methods (including SAM) will use source grid locations using a stride of 8 as query locations.

Vanilla softmax attention - A cross-attention operation, with H heads, between a D -dimensional query vector \mathbf{x} and a set of D -dimensional vectors $\{\mathbf{y}_n\}_{n=1\dots N}$ can be written as follows:

$$\sum_{h=1}^H \sum_{n=1}^N \underbrace{1 \times 1}_{s_{h,n}} \underbrace{D \times D_H}_{\mathbf{W}_{o,h}} \underbrace{D_H \times D}_{\mathbf{W}_{v,h}} \mathbf{y}_n, \quad (2)$$

$$\text{where } s_{h,n} = \frac{\exp(\mathbf{x}^\top \underbrace{\mathbf{W}_{q,h}^\top \mathbf{W}_{k,h}}_{1 \times D \times D_H \times D_H \times D \times D} \mathbf{y}_n^\top)}{\sum_{m=1}^N \exp(\mathbf{x}^\top \mathbf{W}_{q,h}^\top \mathbf{W}_{k,h} \mathbf{y}_m)}. \quad (3)$$

In practice, all the matrices \mathbf{W}_{\cdot} are learned. Usually $D_H = \frac{D}{H}$. The output is a linear combination (with coefficients $s_{h,n}$) of the linearly transformed set of vectors $\{\mathbf{y}_n\}_{n=1\dots N}$. This operation allows to extract the relevant information in $\{\mathbf{y}_n\}_{n=1\dots N}$ from the point of view of the query \mathbf{x} . The query \mathbf{x} is not an element of that linear combination. As a consequence, a residual connection is often added, usually followed by a layer normalization and a 2-layer MLP (with residual connection) (Vaswani et al., 2017).




Figure 3: **Visualization of the learned latent vectors of SAM.** The average query map is obtained by averaging 64 correspondence maps of 64 query locations (Red crosses) while the average latent map is obtained by averaging the 128 correspondence maps of the 128 learned latent vectors. We observe that the average query map is mainly activated around the correspondents, whereas these regions are less activated in the average latent map.

In practice, a cross-attention layer between a set of query vectors $\{\mathbf{x}_n\}_{n=1\dots Q}$ and $\{\mathbf{y}_n\}_{n=1\dots N}$ is performed in parallel, which is often computationally demanding since, for each head, the coefficients are stored in a $Q \times N$ matrix. A self-attention layer is a cross-attention layer where $\{\mathbf{y}_n\}_{n=1\dots N} = \{\mathbf{x}_n\}_{n=1\dots Q}$.

3.2 Feature extraction stage

Our method takes as input a source image I_s ($H_s \times W_s \times 3$), a target image I_t ($H_t \times W_t \times 3$) and a set of 2D query locations $\{\mathbf{p}_{s,i}\}_{i=1\dots L}$. The first stage of SAM is a classical feature extraction stage. From the source image I_s ($H_s \times W_s \times 3$) and the target image I_t ($H_t \times W_t \times 3$), dense visual source features F_s ($\frac{H_s}{4} \times \frac{W_s}{4} \times 128$) and target features F_t ($\frac{H_t}{4} \times \frac{W_t}{4} \times 128$) are extracted using a siamese CNN backbone. The Positional Encodings (PE) of the source and target are computed, using an MLP (Sarlin et al., 2020), and concatenated with the visual features of the source and target to obtain two tensors H_s ($\frac{H_s}{4} \times \frac{W_s}{4} \times 256$) and H_t ($\frac{H_t}{4} \times \frac{W_t}{4} \times 256$). For each 2D query point $\mathbf{p}_{s,i}$, a descriptor $\mathbf{h}_{s,i}$ of size 256 is extracted from H_s . In practice, we use integer query locations thus there is no need for interpolation here. Technical details concerning this stage are provided in the appendix.

3.3 Latent space attention-based stage

In order to allow for the descriptors $\{\mathbf{h}_{s,i}\}_{i=1\dots L}$ (at training-time we use $L = 1024$) to communicate and adjust with respect to H_t , we draw inspiration from Perceiver (Jaegle et al., 2021; Jaegle et al., 2022), and consider a set of $N = M + L$ latent vectors, composed of M learned latent vectors $\{\mathbf{m}_i\}_{i=1\dots M}$ and the L descriptors $\{\mathbf{h}_{s,i}\}_{i=1\dots L}$. These latent vectors $\{\{\mathbf{m}_i\}_{i=1\dots M}, \{\mathbf{h}_{s,i}\}_{i=1\dots L}\}$ are used as queries in an

input cross-attention layer to extract the relevant information from H_t , and finally obtain an updated set of latent vectors $\{\{\mathbf{m}_i^{(0)}\}_{i=1\dots M}, \{\mathbf{h}_{s,i}^{(0)}\}_{i=1\dots L}\}$. On the one hand, the outputs $\{\mathbf{h}_{s,i}^{(0)}\}_{i=1\dots L}$ contain the information relevant to find their respective correspondents within H_t . On the other hand, the outputs $\{\mathbf{m}_i^{(0)}\}_{i=1\dots M}$ extracted a general representation of H_t since they are not aware of the query 2D locations. In practice, we set M to 128. Afterward, a series of S self-attention layers are applied to the latent vectors to get $\{\{\mathbf{m}_i^{(S)}\}_{i=1\dots M}, \{\mathbf{h}_{s,i}^{(S)}\}_{i=1\dots L}\}$. In these layers, all latent vectors can communicate and adjust with respect to each other. For instance, the set $\{\mathbf{m}_i^{(S)}\}_{i=1\dots M}$ can be used to disambiguate certain correspondences. Then, H_t is used as a query in an *output* cross-attention layer to extract the relevant information from the latent vectors. The resulting tensor is written H_t^{out} . For each updated descriptor $\mathbf{h}_{s,i}^{(S)}$, a correspondence map $C_{t,i}$ (of size $\frac{H_t}{4} \times \frac{W_t}{4}$) is obtained by computing the dot product between $\mathbf{h}_{s,i}^{(S)}$ and H_t^{out} . Finally, for each 2D query location $\mathbf{p}_{s,i}$, the predicted 2D correspondent location $\hat{\mathbf{p}}_{t,i}$ is defined as the argmax of $C_{t,i}$.

In Figure 3, we propose a visualization of the learned latent vectors of SAM. For visualization purposes, we used $L = 64$. Thus the average query map is obtained by averaging the 64 correspondence maps of the 64 query locations, while the average latent map is obtained by averaging the 128 correspondence maps of the 128 learned latent vectors. We observe that the average query map is mainly activated around the correspondents, whereas these regions are less activated in the average latent map.

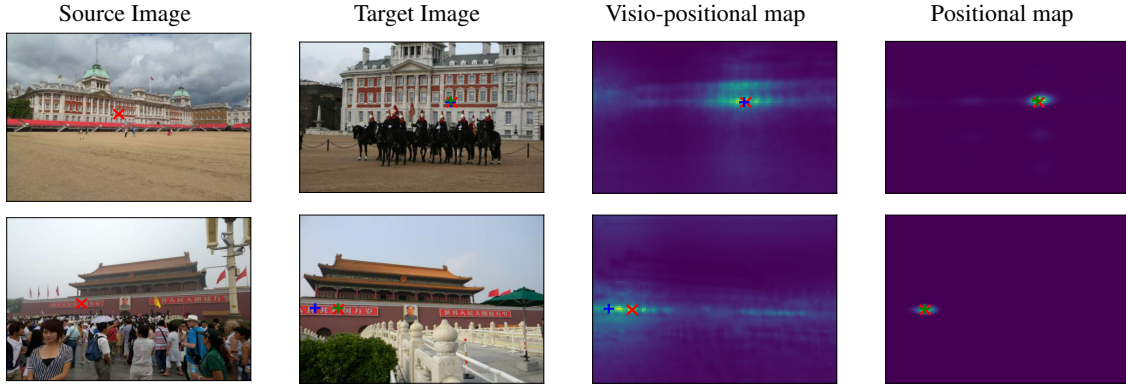


Figure 4: **Visualization - Structured attention.** The visio-positional and positional maps are computed before the output cross-attention. **Red** crosses represent the ground-truth correspondences. **Blue** and **Green** crosses are the maxima of the visio-positional and positional maps, respectively. One can see that the visio-positional maps are highly multimodal (*i.e.*, sensitive to repetitive structures) while the positional maps are almost unimodal.

3.4 Structured attention

At the end of the feature extraction stage, the upper half of each vector is visual features, while the lower half is positional encodings. In each attention layer (input cross-attention, self-attention and output cross-attention), we found it important to structure the linear transformation matrices $W_{o,h}$ and $W_{v,h}$ to constrain the lower half of each output vector to only contain a linear transformation of the positional encodings:

$$\underbrace{W_{o,h}}_{D \times D_H} = \begin{bmatrix} \underbrace{0}_{\frac{D}{2} \times \frac{D_H}{2}} & \underbrace{W_{o,h,up}}_{\frac{D}{2} \times D_H} \\ \underbrace{W_{o,h,low}}_{\frac{D}{2} \times \frac{D_H}{2}} & \end{bmatrix}. \quad (4)$$

The matrices $W_{v,h}$ and the fully connected layers within the MLP (at the end of each attention module) are structured the exact same way. Consequently, throughout the network, the lower half of each latent vector only contains transformations of positional encodings, but no visual feature.

In Figure 4 we propose a visualization of two different correspondence maps *before* the output cross-attention.

The first map is produced using the first 128 dimensions of the feature representation and contains high-level visio-positional features. The second map is produced using the 128 last dimensions of the feature representation and contains only positional encodings. By building these distinct correspondence maps we can observe that while the visio-positional representation is sensitive to repetitive structures, the purely positional representation tends to be only activated in the areas neighbouring the match.

3.5 Loss

At training time, we use a cross-entropy (CE) loss function (Germain et al., 2020) on each correspondence map $C_{t,i}$ in order to maximise its score at the ground truth location $\mathbf{p}_{t,i}$:

$$\text{CE}(C_{t,i}, \mathbf{p}_{t,i}) = -\ln \left(\frac{\exp C_{t,i}(\mathbf{p}_{t,i})}{\sum_{\mathbf{q}} \exp C_{t,i}(\mathbf{q})} \right). \quad (5)$$

3.6 Refinement

The previously described architecture produces coarse correspondence maps of resolution 1/4. Thus the predicted correspondents $\{\hat{\mathbf{p}}_{t,i}\}_{i=1\dots L}$ need to be refined. To do so, we simply use a second siamese CNN that outputs dense source and target features at full resolution. For each 2D query location $\mathbf{p}_{s,i}$, a correspondence map is computed on a window of size 11 centered around the coarse prediction. The predicted 2D correspondent location $\hat{\mathbf{p}}_{t,i}$ is defined as the argmax of this correspondence map. This refinement network is trained separately using the same cross-entropy loss eq. 5. Technical details are provided in the appendix.

4 Experiments

In these experiments, we focus on evaluating six SDF networks (LoFTR (Sun et al., 2021), QuadTree (Tang et al., 2022), ASpanFormer (Chen et al., 2022), 3DG-STFM (Mao et al., 2022), MatchFormer (Wang et al., 2022), TopicFM (Giang et al., 2023)) and our proposed architecture SAM. In all the following Tables, best and second best results are respectively bold and underlined.

Table 1: Evaluation on MegaDepth1500 (Sarlin et al., 2020). We report Matching Accuracy (6) for several thresholds η , computed on the set of all the semi-dense query locations (source grid with stride 8) with available ground truth correspondents (MA), and on a subset containing only query locations that are within a textured region of the source image (MA_{text}). Concerning the pose estimation metrics, we report the classical AUC at 5, 10 and 20 degrees. The proposed SAM method outperforms SDF methods in terms of pose estimation, while SDF methods are significantly better in terms of MA. However, when uniform regions are ignored (MA_{text}), SAM often surpasses SDF methods. These results highlight a strong correlation between the ability to establish precise correspondences in textured regions and the accuracy of the resulting estimated pose.

Method	Matching Accuracy (MA) (6) \uparrow						Matching Accuracy on textured regions (MA_{text}) (6) \uparrow						Pose estimation on 1/8 grid (AUC) \uparrow		
	$\eta=1$		$\eta=2$		$\eta=3$		$\eta=5$		$\eta=10$		$\eta=20$		$@5^\circ$	$@10^\circ$	$@20^\circ$
	$\eta=1$	$\eta=2$	$\eta=3$	$\eta=5$	$\eta=10$	$\eta=20$	$\eta=1$	$\eta=2$	$\eta=3$	$\eta=5$	$\eta=10$	$\eta=20$			
LoFTR (Sun et al., 2021)	49.7	73.2	81.6	87.4	90.5	91.8	55.3	75.2	81.6	86.9	89.9	91.9	52.8	69.2	82.0
MatchFormer (Wang et al., 2022)	51.1	73.4	81.0	86.9	89.5	90.9	56.5	75.6	81.8	87.1	89.6	90.9	52.9	69.7	82.0
TopicFM (Giang et al., 2023)	51.4	75.4	83.7	89.9	92.9	93.5	59.8	77.6	84.8	90.4	92.9	93.7	54.1	70.1	81.6
3DG-STFM(Mao et al., 2022)	51.6	73.7	80.7	86.4	89.0	90.7	57.0	75.8	81.8	86.8	88.8	90.5	52.6	68.5	80.0
ASpanFormer (Chen et al., 2022)	52.0	76.2	84.5	90.7	93.7	94.8	<u>62.2</u>	<u>80.3</u>	<u>85.9</u>	91.0	93.7	<u>94.7</u>	<u>55.3</u>	<u>71.5</u>	<u>83.1</u>
LoFTR+QuadTree (Tang et al., 2022)	<u>51.6</u>	<u>75.9</u>	<u>84.1</u>	<u>90.2</u>	<u>93.1</u>	<u>94.0</u>	61.7	79.9	85.5	90.5	93.3	94.1	54.6	70.5	82.2
SAM (ours)	48.5	70.4	78.0	83.0	85.4	86.4	67.9	83.8	87.3	<u>90.6</u>	<u>93.6</u>	95.2	55.8	72.8	84.2

Table 2: Evaluation on HPatches (Balntas et al., 2017). We report Matching Accuracy (6) for several thresholds η , computed on the set of all the semi-dense query locations (source grid with stride 8) with available ground truth correspondents (MA), and on a subset containing only query locations that are within a textured region of the source image (MA_{text}). Concerning the homography estimation metrics, we report the classical AUC at 3, 5 and 10 pixels. The proposed SAM method is on par with SDF methods in terms of homography estimation, while SDF methods are significantly better in terms of MA. However, when uniform regions are ignored (MA_{text}), SAM matches SDF performances. These results highlight a strong correlation between the ability to establish precise correspondences in textured regions and the accuracy of the resulting estimated homography.

Method	Matching Accuracy (MA) (6) \uparrow			Matching Accuracy on textured regions (MA_{text}) (6) \uparrow			Homography estimation (AUC) \uparrow		
	$\eta=3$	$\eta=5$	$\eta=10$	$\eta=3$	$\eta=5$	$\eta=10$	@3px	@5px	@10px
LoFTR (Sun et al., 2021)	66.8	74.3	77.3	67.6	75.3	78.4	65.9	75.6	84.6
MatchFormer (Wang et al., 2022)	66.2	74.9	78.2	67.7	75.8	79.1	65.0	73.1	81.2
TopicFM (Giang et al., 2023)	<u>72.7</u>	<u>85.0</u>	<u>87.5</u>	74.0	<u>86.0</u>	<u>88.5</u>	<u>67.3</u>	77.0	<u>85.7</u>
3DG-STFM(Mao et al., 2022)	64.9	75.1	78.2	66.2	74.3	77.6	64.7	73.1	81.0
ASpanFormer (Chen et al., 2022)	76.2	86.2	88.7	<u>73.9</u>	85.8	88.4	67.4	<u>76.9</u>	85.6
LoFTR+QuadTree (Tang et al., 2022)	70.2	83.1	85.9	73.5	84.3	86.9	67.1	76.1	85.3
SAM (ours)	62.4	70.9	74.2	73.4	86.6	89.3	67.1	<u>76.9</u>	85.9

Table 3: Evaluation on ETH3D (Schöps et al., 2019) for different frame interval sampling rates r . We report Matching Accuracy (6) for several thresholds η , computed on the set of all the semi-dense query locations (source grid with stride 8) with available ground truth correspondents (MA). For ETH3D, ground truth correspondents are based on structure from motion tracks. Consequently, the MA already ignores untextured regions of the source images which explains why SAM is able to outperform SDF methods.

Method	Matching Accuracy (MA) (6) \uparrow														
	$r = 3$					$r = 7$					$r = 15$				
	$\eta=1$	$\eta=2$	$\eta=3$	$\eta=5$	$\eta=10$	$\eta=1$	$\eta=2$	$\eta=3$	$\eta=5$	$\eta=10$	$\eta=1$	$\eta=2$	$\eta=3$	$\eta=5$	$\eta=10$
LoFTR (Sun et al., 2021)	44.8	76.5	88.4	97.0	99.4	39.7	73.1	87.6	95.9	98.5	33.3	66.2	84.8	92.5	96.3
MatchFormer (Wang et al., 2022)	45.5	77.1	89.2	97.2	99.7	40.4	73.8	87.8	96.6	99.0	34.2	66.7	84.9	93.5	97.0
TopicFM (Giang et al., 2023)	45.1	76.9	89.0	97.2	99.6	39.9	73.5	87.9	96.4	99.0	33.8	66.4	85.0	92.8	96.5
3DG-STFM(Mao et al., 2022)	43.9	76.3	88.0	96.9	99.3	39.3	72.7	87.4	95.5	98.3	32.4	65.7	84.7	92.0	96.0
ASpanFormer (Chen et al., 2022)	45.8	<u>77.6</u>	<u>89.6</u>	<u>97.8</u>	99.8	40.6	73.8	88.1	96.8	99.0	34.3	66.8	85.3	93.9	97.3
LoFTR+QuadTree (Tang et al., 2022)	<u>45.9</u>	77.5	89.5	97.8	99.7	<u>40.8</u>	<u>74.0</u>	<u>88.3</u>	<u>97.0</u>	<u>99.2</u>	<u>34.5</u>	<u>66.8</u>	<u>85.4</u>	<u>94.0</u>	<u>97.3</u>
SAM (ours)	53.4	79.9	91.5	98.0	99.8	48.6	78.6	91.7	98.2	99.4	40.1	70.2	87.8	95.4	97.8

Implementations - For each network, we employ the code and weights (trained on MegaDepth) made available by the authors. Concerning SAM, we train it similarly to the SDF networks on MegaDepth (Li and Snavely, 2018), during 100 hours, on four GeForce GTX 1080 Ti (11GB) GPUs (see the appendix for details). Our code will be made available.

Query locations - Concerning SDF methods, the query locations are defined as the source grid locations using a stride of 8. Thus, in order to be able to compare the performances of SAM against such methods, we use the exact same query locations. To do so (recall that SAM was trained with $L = 1024$ query locations), we simply shuffle the source grid locations (with stride 8) and feed them to SAM by batch of size 1024 (the CNN features are cached making the processing of each minibatch very efficient).

Evaluation criteria - Regarding the pose and homography estimation metrics, we use the classical AUC metrics for each dataset. Thus in each Table, the pose/homography results we obtained for SDF methods (we re-evaluated each method) are the same as the results published in the respective papers. Concerning SAM, the correspondences are classically filtered (similarly to what SDF methods do) before the pose/homography estimator using a simple Mutual Nearest Neighbor with a threshold of 5 pixels.

In order to evaluate the ability of each method to establish correspondences, we consider the Matching Accuracy (MA) as in (Truong et al., 2020), *i.e.*, the average on all images of the ratio of correct matches, for different pixel error thresholds (η):

$$\text{MA}(\eta) = \frac{1}{K} \sum_{k=1}^K \frac{\sum_{i=1}^{L_k} \left[\left\| \hat{\mathbf{p}}_{t_k,i} - \mathbf{p}_{t_k,i}^{\text{GT}} \right\|_2 < \eta \right]}{L_k}, \quad (6)$$

where K is the number of pairs of images, L_k is the number of ground truth correspondences available for the image pair $\#k$ and $[\cdot]$ is the Iverson bracket. We refer to this metric as MA and not "Percentage of Correct Keypoints" because we found this term misleading in cases where the underlying ground truth correspondences are not based on keypoints, as it is the case in MegaDepth and HPatches.

We also propose to introduce the Matching Accuracy in textured regions (MA_{text}) that consists in ignoring, in eq. (6), ground truth correspondences whose query location is within a low-contrast region of the source image, *i.e.*, an almost uniform region of the source image.

Note that SAM's MNN (and the MNN of SDF methods) is only used for the pose/homography estimation, *i.e.*, it is not used to compute matching accuracies, otherwise the set of correspondences would not be the same for each method.

4.1 Evaluation on MegaDepth

We consider the MegaDepth1500 (Sarlin et al., 2020) benchmark. We use the exact same settings as those used by SDF methods, such as an image resolution of 1200. The results are provided in Table 1. We report Matching Accuracy (6), for several thresholds η , computed on the set of all the semi-dense query locations (source grid with stride 8) with available ground truth correspondents (MA). The ground truth correspondents are obtained using the available depth maps and camera poses. Consequently, many query locations located in untextured regions have a ground truth correspondent. Thus, we also report the matching accuracy computed only on query locations that are within a textured region of the source image (MA_{text}). Concerning the pose estimation metrics, we report the classical AUC at 5, 10 and 20 degrees. The proposed SAM method outperforms SDF methods in terms of pose estimation, while SDF methods are significantly better in terms of MA. However, when uniform regions are ignored (MA_{text}), SAM often surpasses SDF methods. These results highlight a strong correlation between the ability to establish precise correspondences in textured regions and the accuracy of the resulting estimated pose.

In Figure 5, we report qualitative results that visually illustrate the previous findings.

4.2 Evaluation on HPatches

We evaluate the different architectures on HPatches (Balntas et al., 2017) (see Table 2). We use the exact same settings as those used by SDF methods. We report Matching Accuracy (6) for several thresholds η , computed on the set of all the semi-dense query locations (source grid with stride 8) with available ground truth correspondents (MA). The ground truth correspondents are obtained using the available homography matrices. Consequently, many query locations located in untextured regions have a ground truth correspondent. Thus, we also report the matching accuracy computed only on query locations that are within a textured region of the source image (MA_{text}). Concerning the homography estimation metrics, we report the classical AUC at 3, 5 and 10 pixels. The proposed SAM method is on par with SDF methods in terms of homography estimation, while SDF methods are significantly better in terms of MA. However, when uniform regions are ignored (MA_{text}), SAM matches SDF performances. These results highlight a strong correlation between the ability to establish precise correspondences in textured regions and the accuracy

Table 4: Ablation study of proposed SAM method (MegaDepth validation set).

Method	MA _{text} ↑					
	η=1	η=2	η=5	η=10	η=20	η=50
Siamese CNN	0.029	0.112	0.436	0.581	0.622	0.687
+ Input CA and SA (x16)	0.137	0.321	0.671	0.734	0.767	0.822
+ Learned LV and output CA	0.132	0.462	0.823	0.871	0.898	0.935
+ PE concatenated	0.121	0.419	0.796	0.868	0.902	0.939
+ Structured AM	<u>0.140</u>	<u>0.487</u>	<u>0.857</u>	<u>0.902</u>	0.922	0.947
+ Refinement (Full model)	0.673	0.791	0.870	0.902	<u>0.921</u>	<u>0.946</u>

of the resulting estimated homography.

In Figure 6, we report qualitative results that visually illustrate the previous findings.

4.3 Evaluation on ETH3D

We evaluate the different networks on several sequences of the ETH3D dataset (Schöps et al., 2019) as proposed in (Truong et al., 2020). Different frame interval sampling rates r are considered. As the rate r increases, the overlap between the image pairs reduces, hence making the matching problem more difficult. The results are provided in Table 3. We report Matching Accuracy (6) for several thresholds η , computed on the set of all the semi-dense query locations (source grid with stride 8) with available ground truth correspondents (MA). For ETH3D, ground truth correspondents are based on structure from motion tracks. Consequently, the MA ignores untextured regions of the source images which explains why SAM is able to outperform SDF methods in terms of MA.

In Figure 7, we report qualitative results that illustrate the accuracy of the proposed SAM method.

4.4 Ablation study

In Table 4, we propose an ablation study of SAM to evaluate the impact of each part of our architecture. This study is performed on MegaDepth validation scenes. Starting from a standard siamese CNN, we show that a significant gain in performance can simply be obtained with the input cross-attention layer (here the PE is added and not concatenated) and self-attention layers. We then add learned latent vectors (LV) in the latent space and use an output cross-attention which again significantly improves the performance. Concatenating the positional encoding information instead of adding it to the visual features reduces the matching accuracy at $\eta = 2$ and $\eta = 5$. However, combining it with structured attention leads to a significant improvement in terms of MA. Finally, as expected, the refinement step improves the matching accuracy for small pixel error thresholds.

5 Conclusion

We proposed a novel Structured Attention-based image Matching architecture (SAM). The flexibility of this novel architecture allowed us to fairly compare it against SDF methods, *i.e.*, in all the experiments we used the same query locations (source grid with stride 8). The experiments highlighted a counter-intuitive result on two datasets (MegaDepth and HPatches): on the one hand SAM either outperforms or is on par with SDF methods in terms of pose/homography estimation metrics, but on the other hand SDF approaches are significantly better than SAM in terms of Matching Accuracy (MA). Here the MA is computed on all the semi-dense locations of the source image with available ground truth correspondent, which includes both textured and uniform regions. We proposed to limit the computation of the matching accuracy to textured regions, and showed that in this case SAM often surpasses SDF methods. These findings highlighted a strong correlation between the ability to establish precise correspondences in textured regions and the accuracy of the resulting estimated pose/homography. We also evaluated the aforementioned methods on ETH3D which confirmed, on a third dataset, that SAM has a strong ability to establish correspondences in textures regions. We finally performed an ablation study of SAM to demonstrate that each part of the architecture is important to obtain such a strong matching capacity.

REFERENCES

- Arnold, E., Wynn, J., Vicente, S., Garcia-Hernando, G., Monszpart, Á., Prisacariu, V. A., Turmukhambetov, D., and Brachmann, E. (2022). Map-free visual relocalization: Metric pose relative to a single image. In *ECCV*.
- Baltas, V., Lenc, K., Vedaldi, A., and Mikolajczyk, K. (2017). Hpatches: A benchmark and evaluation of handcrafted and learned local descriptors. In *CVPR*.
- Chen, H., Luo, Z., Zhou, L., Tian, Y., Zhen, M., Fang, T., McKinnon, D., Tsin, Y., and Quan, L. (2022). Aspanformer: Detector-free image matching with adaptive span transformer. In *ECCV*.
- Dusmanu, M., Rocco, I., Pajdla, T., Pollefeys, M., Sivic, J., Torii, A., and Sattler, T. (2019). D2-net: A trainable CNN for joint description and detection of local features. In *CVPR*.
- Edstedt, J., Athanasiadis, I., Wadenbäck, M., and Felsberg, M. (2023). Dkm: Dense kernelized feature matching for geometry estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17765–17775.
- Germain, H., Bourmaud, G., and Lepetit, V. (2020).




Figure 5: **Qualitative results on MegaDepth1500.** For each image pair: (top row) Visualization of established correspondences used to compute the MA, (bottom row) Visualization of established correspondences used to compute the MA_{text}. Line colors indicate the distance in pixels with respect to the ground truth correspondent.

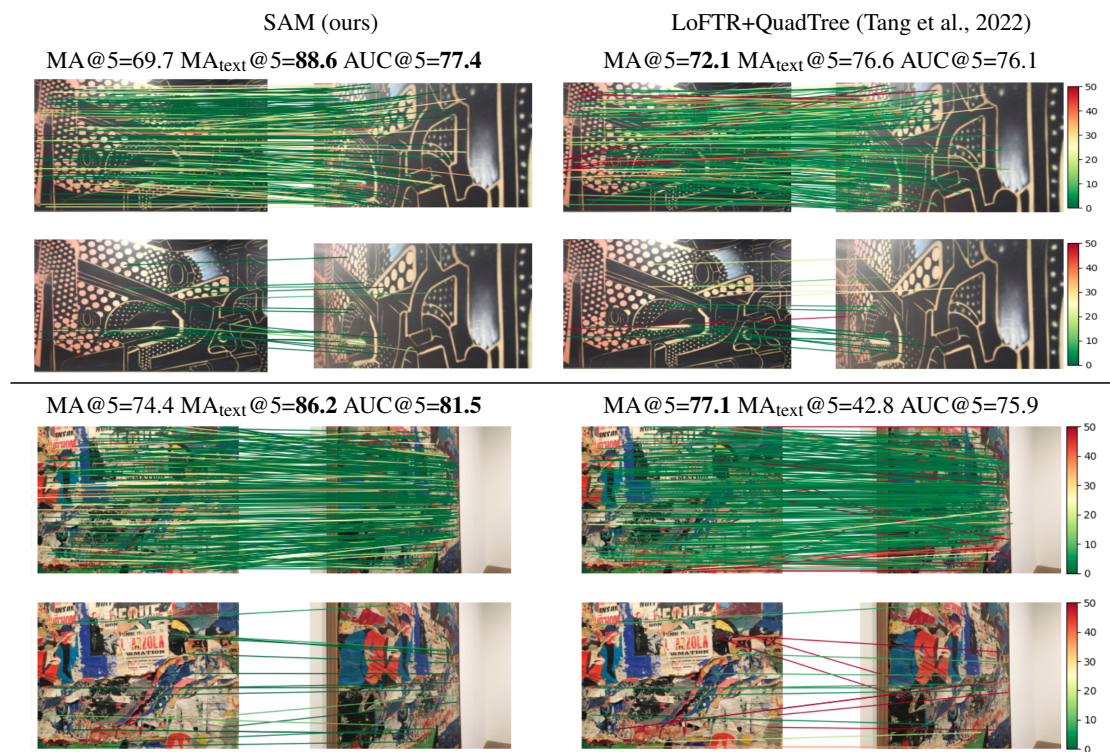


Figure 6: **Qualitative results on HPatches.** Line colors indicate the distance in pixels to the ground truth correspondent.

- S2DNet: learning image features for accurate sparse-to-dense matching. In *ECCV*.
- Germain, H., Lepetit, V., and Bourmaud, G. (2021). Neural reprojection error: Merging feature learning and camera pose estimation. In *CVPR*.
- Germain, H., Lepetit, V., and Bourmaud, G. (2022). Visual correspondence hallucination. In *ICLR*.
- Giang, K. T., Song, S., and Jo, S. (2023). TopicFM: Robust and interpretable feature matching with topic-assisted. In *AAAI*.
- Heinly, J., Schönberger, J. L., Dunn, E., and Frahm, J.-M. (2015). Reconstructing the World* in Six Days *(as Captured by the Yahoo 100 Million Image Dataset). In *CVPR*.
- Jaegle, A., Borgeaud, S., Alayrac, J.-B., Doersch, C., Ionescu, C., Ding, D., Koppula, S., Zoran, D., Brock, A., Shelhamer, E., Henaff, O. J., Botvinick, M., Zisserman, A., Vinyals, O., and Carreira, J. (2022). Perceiver IO: A general architecture for structured inputs & outputs. In *ICLR*.
- Jaegle, A., Gimeno, F., Brock, A., Vinyals, O., Zisserman, A., and Carreira, J. (2021). Perceiver: General perception with iterative attention. In *ICML*.
- Jin, Y., Mishkin, D., Mishchuk, A., Matas, J., Fua, P., Yi, K. M., and Trulls, E. (2020). Image Matching across Wide Baselines: From Paper to Practice. *IJCV*.
- Katharopoulos, A., Vyas, A., Pappas, N., and Fleuret, F. (2020). Transformers are RNNs: Fast autoregressive transformers with linear attention. In *ICML*.
- Li, X., Han, K., Li, S., and Prisacariu, V. (2020). Dual-resolution correspondence networks. *NeurIPS*.
- Li, Z. and Snavely, N. (2018). Megadepth: Learning single-view depth prediction from internet photos. In *CVPR*.
- Lin, T.-Y., Dollár, P., Girshick, R., He, K., Hariharan, B., and Belongie, S. (2017). Feature pyramid networks for object detection. In *CVPR*.
- Lowe, D. G. (1999). Object recognition from local scale-invariant features. In *ICCV*.
- Mao, R., Bai, C., An, Y., Zhu, F., and Lu, C. (2022). 3D-STFM: 3D geometric guided student-teacher feature matching. In *ECCV*.
- Ni, J., Li, Y., Huang, Z., Li, H., Bao, H., Cui, Z., and Zhang, G. (2023). PATS: Patch area transportation with subdivision for local feature matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17776–17786.
- Ono, Y., Trulls, E., Fua, P., and Yi, K. M. (2018). LF-Net: learning local features from images. *NeurIPS*.
- Revaud, J., De Souza, C., Humenberger, M., and Weinzaepfel, P. (2019). R2D2: reliable and repeatable detector and descriptor. *NeurIPS*.
- Rocco, I., Arandjelović, R., and Sivic, J. (2020a). Efficient neighbourhood consensus networks via submanifold sparse convolutions. In *ECCV*.
- Rocco, I., Cimpoi, M., Arandjelović, R., Torii, A., Pajdla, T., and Sivic, J. (2018). Neighbourhood consensus networks. *NeurIPS*.

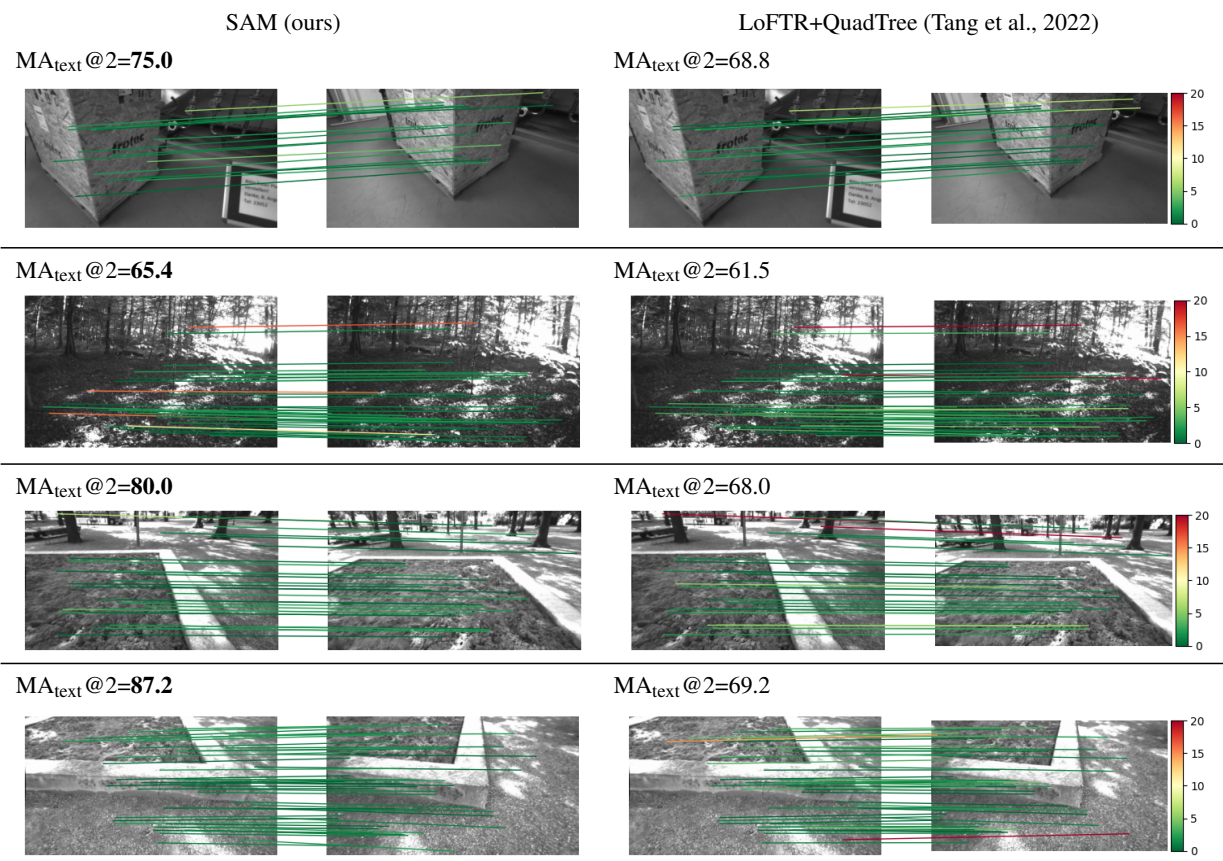


Figure 7: **Qualitative results on ETH3D.** Line colors indicate the distance in pixels to the ground truth correspondent.

- Rocco, I., Cimpoi, M., Arandjelović, R., Torii, A., Pajdla, T., and Sivic, J. (2020b). NCNet: neighbourhood consensus networks for estimating image correspondences. *PAMI*.
- Sarlin, P.-E., DeTone, D., Malisiewicz, T., and Rabinovich, A. (2020). Superglue: Learning feature matching with graph neural networks. In *CVPR*.
- Sarlin, P.-E., Dusmanu, M., Schönberger, J. L., Speciale, P., Gruber, L., Larsson, V., Miksik, O., and Pollefeys, M. (2022). LaMAR: Benchmarking Localization and Mapping for Augmented Reality. In *ECCV*.
- Sattler, T., Torii, A., Sivic, J., Pollefeys, M., Taira, H., Okutomi, M., and Pajdla, T. (2017). Are large-scale 3D models really necessary for accurate visual localization? In *CVPR*.
- Schönberger, J. L. and Frahm, J.-M. (2016). Structure-from-motion revisited. In *CVPR*.
- Schöps, T., Sattler, T., and Pollefeys, M. (2019). BAD SLAM: Bundle adjusted direct RGB-D SLAM. In *CVPR*.
- Strasdat, H., Davison, A. J., Montiel, J. M., and Konolige, K. (2011). Double window optimisation for constant time visual slam. In *ICCV*.
- Sun, J., Shen, Z., Wang, Y., Bao, H., and Zhou, X. (2021). LoFTR: detector-free local feature matching with transformers. In *CVPR*.
- Svärm, L., Enqvist, O., Kahl, F., and Oskarsson, M. (2017). City-scale localization for cameras with known vertical direction. *PAMI*.
- Taira, H., Okutomi, M., Sattler, T., Cimpoi, M., Pollefeys, M., Sivic, J., Pajdla, T., and Torii, A. (2018). In-Loc: Indoor visual localization with dense matching and view synthesis. In *CVPR*.
- Tang, S., Zhang, J., Zhu, S., and Tan, P. (2022). Quadtree attention for vision transformers. In *ICLR*.
- Truong, P., Danelljan, M., and Timofte, R. (2020). Glunet: Global-local universal network for dense flow and correspondences. In *CVPR*.
- Truong, P., Danelljan, M., Van Gool, L., and Timofte, R. (2021). Learning accurate dense correspondences and when to trust them. In *CVPR*.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. (2017). Attention is all you need. *NeurIPS*.
- Wang, Q., Zhang, J., Yang, K., Peng, K., and Stiefelhagen, R. (2022). Matchformer: Interleaving attention in transformers for feature matching. In *ACCV*.
- Wang, W., Xie, E., Li, X., Fan, D.-P., Song, K., Liang, D., Lu, T., Luo, P., and Shao, L. (2021). Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In *ICCV*.
- Xie, E., Wang, W., Yu, Z., Anandkumar, A., Alvarez, J. M.,

and Luo, P. (2021). Segformer: Simple and efficient design for semantic segmentation with transformers. *NeurIPS*.

Yi, K. M., Trulls, E., Lepetit, V., and Fua, P. (2016). LIFT: Learned invariant feature transform. In *ECCV*.

Zhang, Z., Sattler, T., and Scaramuzza, D. (2021). Reference pose generation for long-term visual localization via learned features and view synthesis. *IJCV*.

Zhou, Q., Sattler, T., and Leal-Taixe, L. (2021). Patch2Pix: Epipolar-guided pixel-level correspondences. In *CVPR*.

Zhu, S. and Liu, X. (2023). Pmatch: Paired masked image modeling for dense geometric matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21909–21918.

APPENDIX

SAM architecture. Each attention layer in SAM is a classical series of softmax attention, normalization, MLP and normalization layers with residual connections, except for the "output" cross-attention where we removed the MLP, the last normalization and the residual connection.

We empirically choose a combination of 16 self-attention layers and 128 vectors in the learned latent space. We found this combination to be a good trade-off between performance and computational cost.

Backbone description. We use a modified ResNet-18 as our feature extraction backbone. We extract the feature map after the last ResBlock in which we removed the last ReLU. The features are reduced to $\frac{1}{4}$ of the image resolution ($3 \times H \times W$) by applying a stride of 2 in the first and third layer. The feature representation built by the backbone is of size $128 \times \frac{1}{4}H \times \frac{1}{4}W$.

Refinement description. The previous CNN backbone is used for refinement. We simply added a FPN (Feature Pyramid Network) module to extract the feature maps after the first convolutional layer and each ResNet layer. It leaves us with a feature map of size $128 \times H \times W$ on which the refinement stage is performed.

Structured attention. In our implementation, $D = 256$, we use 8 heads, thus $D_H = 32$.

Training details. The first part of the model to be trained is the refinement model (CNN+FPN). It is trained during 50 hours on 4 Nvidia GTX1080 Ti GPUs using Adam optimizer, a constant learning rate




Figure 8: **Training (left) and Validation (right) loss for refinement ResNet-18+FPN.** x-axis represent the number of mini-batches seen by the model and y-axis the cross-entropy value.




Figure 9: **Training (left) and Validation (right) loss for SAM.** x-axis represent the number of mini-batches seen by the model and y-axis the cross-entropy value.

of 10^{-3} and a batch size of 1 (Figure 8). We minimize the cross-entropy of the full resolution correspondence maps produced using the FPN.

The CNN backbone of SAM is initialized with the weights of the previously trained network. We train SAM for 100 hours. The same setup with 4 Nvidia GTX1080 Ti GPUs, Adam optimizer and a batch size of 1 is used. Concerning the learning rate schedule, we use a linear warm-up of 5000 steps (from 0 to 10^{-4}) and then an exponential decay (towards 10^{-5}) (Figure 9). We minimize the cross-entropy of the $\frac{1}{4}$ resolution correspondence maps produced by SAM.