

TransforMatcher: Match-to-Match Attention for Semantic Correspondence

Seungwook Kim Juhong Min Minsu Cho

Pohang University of Science and Technology (POSTECH), South Korea

<http://cvlab.postech.ac.kr/research/TransforMatcher>

Abstract

Establishing correspondences between images remains a challenging task, especially under large appearance changes due to different viewpoints or intra-class variations. In this work, we introduce a strong semantic image matching learner, dubbed *TransforMatcher*, which builds on the success of transformer networks in vision domains. Unlike existing convolution- or attention-based schemes for correspondence, *TransforMatcher* performs global match-to-match attention for precise match localization and dynamic refinement. To handle a large number of matches in a dense correlation map, we develop a light-weight attention architecture to consider the global match-to-match interactions. We also propose to utilize a multi-channel correlation map for refinement, treating the multi-level scores as features instead of a single score to fully exploit the richer layer-wise semantics. In experiments, *TransforMatcher* sets a new state of the art on *SPair-71k* while performing on par with existing SOTA methods on the *PF-PASCAL* dataset.

1. Introduction

Establishing correspondences between images is a fundamental task in computer vision, and is used for a wide range of problems including 3D reconstruction, visual localization and object recognition [11]. With the recent advances of deep neural networks, many learning-based keypoint extractors and feature descriptors were introduced [7, 10, 41, 51, 53], showing significantly improved performances over their traditional counterparts [1, 6, 32, 33]. More recently, dense feature matching methods - which use all extracted features for matching - have shown impressive performances despite higher computation complexities [29, 34, 45]. However, establishing reliable correspondences between images under the presence of intra-class variations *i.e.*, different instances of the same category, remains a critical challenge for semantic visual correspondence [3, 12–14, 16–18, 20, 30, 30, 34, 36, 38, 42, 43, 45, 53].

The idea of applying high-dimensional convolutional

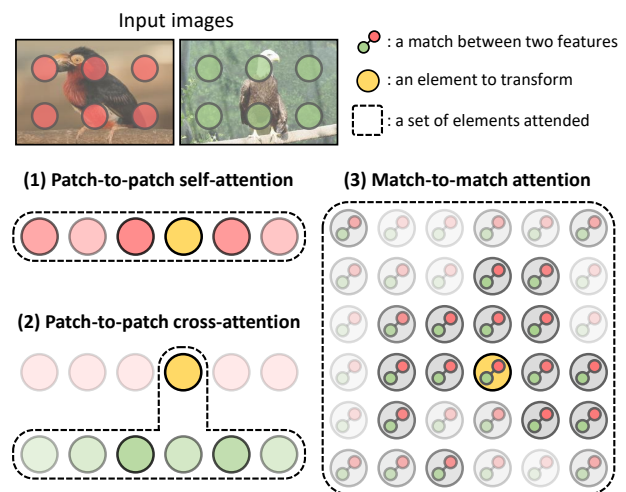


Figure 1. **Patch-to-patch vs. Match-to-match attention.** Patch-to-patch attention considers each position in a 2D feature map as an individual element, while match-to-match attention considers every match in pair-wise correlations as an individual element.

layers on the 4D feature correlation map was first proposed in NCNet [45], which proposes that unique matches will support the nearby ambiguous matches. Among the various methods proposed for establishing semantic correspondences, NCNet and its follow-up methods have shown impressive results [16, 27, 34, 44, 45]. These methods evidence that considering the match-to-match consensus by utilizing the full set of dense correspondences represented by the 4D correlation map is effective in establishing robust and accurate semantic correspondences. However, the convolution-based methods suffer from inherent limitations of *local* and *static* transformations; performing the same local transformation over all spatial positions of the input.

While convolutional neural networks have been the de-facto standard for visual correspondence, transformer networks have recently shown promising results in the computer vision domain. The success of transformer networks can be largely attributed to their *dynamic* feature transform unlike stationary convolutional layers, and the *non-local* interactions between input elements which enable easy scal-




Figure 2. **Conceptual difference between recent methods and ours.** Convolution-based matching methods [16, 34, 35, 45] (left), Cost Aggregation Transformers [3] (middle), and ours (right).

ability to attend to global contexts. For example, ViT [9] attains excellent results compared to convolutional baselines on the task of image recognition with fewer training computational resources; Segmenter [47] outperforms convolution-based methods by modeling global context already at the first layer and throughout the network. These pioneering work show that transformer layers are attractive alternatives to convolutional layers in vision models.

Inspired by the effectiveness of match-to-match consensus consideration and transformer networks, we propose a novel semantic matching pipeline, dubbed *TransforMatcher*. Specifically, we introduce match-to-match attention, a self-attention based mechanism to consider the *global* match-to-match interactions by leveraging the 4D correlation maps computed from features of images to match. Considering the global match-wise interactions allows to capture long-range relevance across matches, and incorporates geometric consistency between distant matches in a dynamic manner especially under challenging appearance variations. This is achieved by considering each spatial entry of the 4D correlation map (*i.e.* a match) as an individual element for attention, which differs from LoFTR [49] or CoTR [19] which consider the patch-to-patch relations within or across 2D feature maps through self- or cross-attention. Figure 1 visualizes the comparison between patch-to-patch and match-to-match attention.

Our contributions can be summarized as follows:

- We propose the TransforMatcher, a novel image matching pipeline built on transformer networks for dynamic match-to-match interactions at a global scale,
- To the best of our knowledge, we are the first to model the *global* interactions between the full set of dense correspondences using a self-attention mechanism within feasible computational constraints,
- We leverage multi-level correlation scores to be used as features, improving over using a single score,
- We demonstrate state-of-the-art or on-par performances on standard benchmarks of category-level matching - SPair-71k and PF-PASCAL.

2. Related work

Category-level matching using convolutional networks.

Category-level matching, a.k.a semantic matching aims to find corresponding elements between images of different instances in the same category. Traditional approaches to category-level matching use hand-crafted descriptors to obtain matches between images [2, 50]. Recent approaches [18, 27, 38] build on the success of deep learning to extract learned features from convolutional neural networks, usually pretrained on the ImageNet classification task [23]. An emerging trend is to exploit high-dimensional convolution on the correlation map obtained from features of images to match, considering the local match-to-match consensus to refine the correlation map [24, 26, 34, 45].

While these work have proven the efficacy of utilizing correlation maps for *local* match-to-match consensus in discovering reliable matches, we propose that exploiting the *global* match-to-match interactions further enables to capture long-range relevance between matches, which is crucial for image pairs with challenging appearance variations. We therefore impose efficient match-to-match attention on the 4D correlation map, exploiting a lightweight attention scheme to easily scale to use the global context.

Image matching using transformer networks.

Following the success of transformer networks in computer vision [9, 31, 52, 54, 57], recent instance-level matching methods propose to use transformer networks. On a conceptual level, SuperGlue [46] employs an attention-like mechanism on a set of sparse keypoints and their descriptors. LoFTR [49] extends this idea to dense 2D feature maps of the images to match, leveraging self- and cross-attention layers between the feature maps to generate strong features for matching. CoTR [19] concatenates the feature maps of images to match along the spatial dimension, which is used as input to the transformer networks together with the query point to output the target point. Note that these methods are actually performing patch-to-patch attention, not leveraging the match-to-match interactions between feature maps.

The work of CATs [3] does employ the transformer networks to model global consensus on the 4D correlation map for the task of semantic correspondence. However, they differ from our work in the following aspects: (1) We use every match on the correlation map as the input element and multi-level scores as features to perform match-to-match attention to model fine-grained interaction, but CATs reshapes the 4D correlation map to 2D feature maps to perform patch-to-patch attention, modeling a comparatively coarse-grained interaction between elements. This is illustrated in Figure 2. (2) CATs additionally concatenates a transformed feature map to the reshaped correlation map, increasing the memory overhead of each transformer layer, making it infeasible to stack multiple layers.

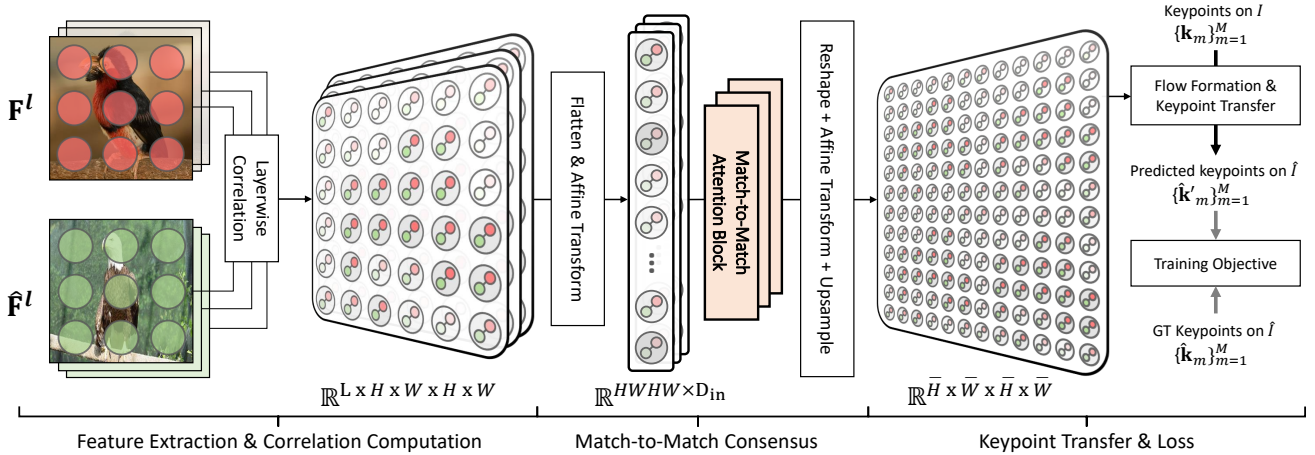


Figure 3. **Overview of TransforMatcher.** The feature maps extracted from an image pair are used to compute a multi-channel correlation map to be processed by our match-to-match attention module for refinement. We construct a dense flow field from the resulting correlation map, which can be used to transfer keypoints for training with keypoint pair annotation.

Efficient Transformers. Due to the quadratic complexity of conventional transformers [55], they are infeasible to model extremely long-range interactions. This motivates the use of efficient transformers with lower computational complexity for feasible computation overhead when handling long sequences. Reformer [22] reduces the complexity down to log-linear using locality-sensitive hashing and reversible residual layers. Linformer [56] approximates the self-attention mechanism using low-rank matrices for linear complexity. Instead of relying on sparsity or low-rankness, Performer [4] proposes positive orthogonal random features approach (FAVOR+) to achieve linear complexity as well. Recently, Fastformer [58] proposes an architecture which uses additive attention techniques only with element-wise products. We build on the success of additive attention to implement global match-to-match attention for its scalable complexity and efficacy.

3. Preliminaries: Transformer

Transformers [55] are built on multi-head self-attention (MHSA) which consists of multiple self-attention layers. Each self-attention layer takes input elements $\mathbf{X} \in \mathbb{R}^{T \times D_{in}}$ to form global self-attention matrices using linear projections of $\mathbf{W}_Q^{(h)}, \mathbf{W}_K^{(h)} \in \mathbb{R}^{D_{in} \times D_h}$ and $\mathbf{W}_V^{(h)} \in \mathbb{R}^{D_{in} \times D_v}$, capturing long-range dependencies between the elements:

$$\text{SA}^{(h)}(\mathbf{X}) = \sigma(\tau \mathbf{X} \mathbf{W}_Q^{(h)} (\mathbf{X} \mathbf{W}_K^{(h)})^\top) \mathbf{X} \mathbf{W}_V^{(h)} \quad (1)$$

$$= \sigma(\tau \mathbf{Q}^{(h)} \mathbf{K}^{(h)\top}) \mathbf{V}^{(h)}, \quad (2)$$

where (h) is the head index, τ is a scaling parameter, and $\sigma(\cdot)$ is row-wise softmax function. The MHSA layer with N_h heads aggregates the self-attention outputs by affine

transformation with $\mathbf{W}_O \in \mathbb{R}^{N_h D_v \times D_{out}}$ and $\mathbf{b}_O \in \mathbb{R}^{D_{out}}$:

$$\text{MHSA}(\mathbf{X}) = \text{concat}_{h \in [N_h]} [\text{SA}^{(h)}(\mathbf{X})] \mathbf{W}_O + \mathbf{b}_O. \quad (3)$$

It can be seen that the computational complexity of the transformer architecture is quadratic with respect to the sequence length T , being a fundamental bottleneck when handling long sequences ($T \gg D_h$). This bottleneck also pertains to our case of processing 4D correlation map, *i.e.*, a full set of pair-wise correlations between two 2D feature maps, as establishing match-to-match attention matrix in self-attention layer demands *quartic* memory with respect to the spatial size of the feature maps. In the next section, we provide an overview of our method as well as an efficient self-attention layer which implements global match-to-match interactions without quartic complexity.

4. TransforMatcher

We first provide an overview of our TransforMatcher pipeline. Given a pair of images to match, a feature extractor provides a set of intermediate feature map pairs which are used to construct a multi-channel correlation map. Due to multifarious match-wise interactions within the 4D global correlation map, we employ additive attention with linear complexity to perform match-to-match attention with feasible computation overhead. We refine the multi-channel correlation map with several match-to-match attention layers, considering the global context within the correlation map in a dynamic manner. The refined correlation map is used to construct a dense flow field, which can be used for keypoint transfer to supervise our pipeline with ground-truth keypoint pair annotations. Fig. 3 illustrates the overview architecture of our method.




Figure 4. **Match-to-match attention module.** The multi-channel correlation map is projected to query, key and value matrices, which are multiplied with rotary positional embeddings. The match-to-match attention module exploits additive addition mechanisms to aggregate query/key matrices to global vectors, which is used for element-wise product to induce global context awareness. The final output is projected to a single-width channel to be reshaped to a refined 4D correlation map.

4.1. Multi-channel correlation computation

We use the ImageNet-pretrained ResNet-101 [15] architecture as the feature extractor. We use all bottleneck layers of `conv4_x` and `conv5_x` to extract the features given an input pair of images $I, \hat{I} \in \mathbb{R}^{H \times W \times 3}$, and denote the set of intermediate feature pairs as $\{(\mathbf{F}^l, \hat{\mathbf{F}}^l)\}_{l=1}^L$.

A feature map pair extracted from the same bottleneck layer, $\mathbf{F}^l, \hat{\mathbf{F}}^l \in \mathbb{R}^{H_l \times W_l \times D_l}$, are used to construct a correlation map $\mathbf{C}^l \in \mathbb{R}^{H_l \times W_l \times H_l \times W_l}$ which represents the confidence score for all candidate correspondences between the two feature maps. Given a set of feature map pairs from different bottleneck layers $\{(\mathbf{F}^l, \hat{\mathbf{F}}^l)\}_{l=1}^L$, we compute the 4D correlation tensors for each pair as follows:

$$\mathbf{C}_{\mathbf{x}, \hat{\mathbf{x}}}^l = \text{ReLU}\left(\frac{\mathbf{F}_{\mathbf{x},:}^l \cdot \hat{\mathbf{F}}_{\hat{\mathbf{x}},:}^l}{\|\mathbf{F}_{\mathbf{x},:}^l\| \|\hat{\mathbf{F}}_{\hat{\mathbf{x}},:}^l\|}\right), \quad (4)$$

where $\mathbf{x}, \hat{\mathbf{x}} \in \mathbb{R}^2$ refer to 2-dimensional spatial positions of the feature maps corresponding to the image pair (I, \hat{I}) . The L correlation tensors are then stacked together along the channel dimension after bilinear interpolation to the size of $H \times W \times H \times W$, *i.e.*, $\frac{1}{16}$ the size of the input image resolutions, resulting in the final multi-channel correlation map $\mathbf{C} \in \mathbb{R}^{L \times H \times W \times H \times W}$.

This is unlike correlation maps used in prior work [45], which only have a single channel, *i.e.*, one similarity score value for each pair of positions between the source and target feature maps. By constructing a multi-channel correlation map, we treat the multi-level scores for each candidate match as *features* instead of a single *score*. This leverage of different correlation tensors across the bottleneck layers allows us to exploit the richer semantics in different levels of feature maps, unlike previous methods which disre-

gard the layer-wise similarities and semantics. Furthermore, having a non-single channel prior to the linear projection to query, key and value matrices is architecturally natural for a transformer-based architecture.

4.2. Match-to-match attention

Attention layer. We flatten the 4D correlation map to behave as the input sequence for the transformer module, *i.e.*, $\mathbb{R}^{L \times H \times W \times H \times W} \rightarrow \mathbb{R}^{L \times HWHW}$, considering the match at each spatial position as an element for attention. We then linearly embed the channel dimension of our flattened correlation map, *i.e.*, $\mathbf{X} = \mathbf{C}^T \mathbf{W}_{\text{in}}$, where \mathbf{C} refers to the correlation map, $\mathbf{W}_{\text{in}} \in \mathbb{R}^{L \times D_{\text{in}}}$ is the linear transformation matrix, and $\mathbf{X} \in \mathbb{R}^{HWHW \times D_{\text{in}}}$ is the input to the subsequent attention blocks. However, the quadratic complexity of conventional self-attention in transformers poses an infeasible computation overhead in our setting, as a flattened 4D tensor results in a significantly long 1D tensor.

Inspired by Fastformer [58], we aim to alleviate this bottleneck through the use of *additive* attention to effectively model long-range match-to-match interactions; instead of computing a quartic attention map (with respect to the spatial size of feature maps) which encodes all possible interactions between candidate matches $\mathbf{Q}\mathbf{K}^T \in \mathbb{R}^{T \times T}$ where $T = HWHW$, we form a compact representation of query-key interactions $\mathbf{H} \in \mathbb{R}^{T \times D_h}$ via additive attention which computes interactions between a global query representation and every key vector:

$$\mathbf{H}_{i,:}^{(h)} = \mathbf{K}_{i,:}^{(h)} \odot \sum_{j=1}^T \mathbf{Q}_{j,:}^{(h)} \sigma(\tau \mathbf{w}_q \mathbf{Q}^{(h)\top})_j, \quad (5)$$

where $\mathbf{w}_q \in \mathbb{R}^{D_h}$ learns to transform the query vectors into a global vector. A similar additive attention mecha-

nism summarizes the context-aware key representations \mathbf{H} with a linear projection $\mathbf{w}_k \in \mathbb{R}^{D_h}$ to model its interaction with value vectors as follows:

$$\text{SA}_{\text{TM}}^{(h)}(\mathbf{X})_{i,:} = \mathbf{V}_{i,:}^{(h)} \odot \sum_{j=1}^T \mathbf{H}_{j,:}^{(h)} \sigma(\tau \mathbf{w}_k \mathbf{H}^{(h)\top})_j, \quad (6)$$

with the assumption of $D_h = D_v$. The output is transformed by an MLP followed by residual connection with \mathbf{Q} . Our proposed match-to-match attention layer reduces the time and memory complexity down to linear with respect to the input length: $\mathcal{O}(T^2 D_h) \rightarrow \mathcal{O}(T D_h)$.

Finally, to ensure that our attention layer can attend to parts of the flattened correlation map differently, we formulate our multi-head self-attention layer as follows:

$$\text{MHSA}_{\text{TM}}(\mathbf{X}) = \text{concat}_{h \in [N_h]} [\text{SA}_{\text{TM}}^{(h)}(\mathbf{X})] \mathbf{W}_O + \mathbf{b}_O. \quad (7)$$

where a linear transformation layer transforms the concatenated outputs of the multiple self-attention layers. We use the pre-LN approach, where the layer normalization is placed inside the residual blocks of the attention layers.

4D rotary positional embedding. In transformer-based networks, positional embedding models the dependency between elements at different positions in the sequence. While relative positional embedding has shown to outperform absolute positional embedding in modelling relation-aware interactions, it is not applicable to linear-complexity transformers as they do not explicitly compute the quadratic-complexity attention matrix. To this end, we employ rotary positional embedding (RoPE) [48] and extend it to be applicable on our 4D correlation map input.

RoPE aims to make the interaction of query and key (inner product for vanilla transformers) encode the position information only in the relative form. Their proposed attention matrix computation with RoPE in vanilla quadratic-complexity transformers can be formulated as follows:

$$\mathbf{Q}_{m,:}^{(h)} \mathbf{K}_{n,:}^{(h)\top} = (\mathbf{X}_{m,:} \mathbf{W}_Q^{(h)} \mathbf{R}_{(\Theta, m)}) (\mathbf{X}_{n,:} \mathbf{W}_K^{(h)} \mathbf{R}_{(\Theta, n)})^\top \quad (8)$$

$$= \mathbf{X}_{m,:} \mathbf{W}_Q^{(h)} \mathbf{R}_{(\Theta, n-m)} \mathbf{W}_K^{(h)\top} \mathbf{X}_{n,:}^\top, \quad (9)$$

where $\mathbf{R}_{(\Theta, *)} \in \mathbb{R}^{D_h \times D_h}$ is the rotary matrix which is for rotating the key or query vectors by amount of angle in multiples of their position indices to incorporate relative positional embedding. We guide the readers to the supplementary for detailed explanations.

RoPE can be applied to linear-complexity transformers as well [48]. In our work, we achieve this by using Eq. (5) to calculate global context-aware query-key interactions, but with $\mathbf{K} = \mathbf{X} \mathbf{W}_K \mathbf{R}_{(\Theta, *)}$ and $\mathbf{Q} = \mathbf{X} \mathbf{W}_Q \mathbf{R}_{(\Theta, *)}$.

Single-channel refined correlation computation. In a nutshell, our match-to-match module takes as input a noisy

4D correlation map to refine it using match-to-match interactions, outputting a refined correlation map for robust image matching. This process is repeated N times, providing a tensor in $\mathbb{R}^{L \times H W H W}$. The output from the final match-to-match attention module is linearly projected to a single channel dimension, and is reshaped back to 4D correlation map *i.e.* $\mathbb{R}^{L \times H W H W} \rightarrow \mathbb{R}^{H \times W \times H \times W}$, for reliable keypoint transfer. For precise transfer, we perform a 4-dimensional upsampling function on the 4D correlation map, and denote the tensor as $\mathbf{C}^{\text{out}} \in \mathbb{R}^{\bar{H} \times \bar{W} \times \bar{H} \times \bar{W}}$ where $\bar{H} = 2H$ and $\bar{W} = 2W$ which corresponds to $\frac{1}{8}$ the size of the original image. We illustrate the outline of our match-to-match attention module in Figure 4.

4.3. Flow field formation

The output correlation tensor \mathbf{C}^{out} can be transformed into a dense flow field by applying kernel soft-argmax [25]. We normalize the raw correlation outputs using softmax:

$$\mathbf{C}^{\text{norm}} = \frac{\exp(\mathbf{G}_{kl}^{\mathbf{P}} \mathbf{C}_{ijkl}^{\text{out}})}{\sum_{(k',l') \in \bar{H} \times \bar{W}} \exp(\mathbf{G}_{k'l'}^{\mathbf{P}} \mathbf{C}_{ij k'l'}^{\text{out}})}, \quad (10)$$

where $\mathbf{G}^{\mathbf{P}} \in \mathbb{R}^{\bar{H} \times \bar{W}}$ is a 2-dimensional Gaussian kernel centered on $\mathbf{p} = \arg \max_{k,l} \mathbf{C}_{i,j,k,l}^{\text{out}}$, which is applied to smooth the potentially irregular correlation values. The normalized correlation tensor \mathbf{C}^{norm} encodes a set of probability simplexes, which we use to transfer all the coordinates on the dense regular grid $\mathbf{P} \in \mathbb{R}^{\bar{H} \times \bar{W} \times 2}$ of source image I to obtain their corresponding coordinates $\hat{\mathbf{P}}' \in \mathbb{R}^{\bar{H} \times \bar{W} \times 2}$ on target image \hat{I} : $\hat{\mathbf{R}}'_{i,j} = \sum_{(k,l) \in \bar{H} \times \bar{W}} \mathbf{C}_{i,j,k,l}^{\text{norm}} \mathbf{P}_{k,l}$. We then can construct a dense flow field at sub-pixel level using the set of estimated matches $(\mathbf{P}, \hat{\mathbf{P}}')$.

4.4. Training objective

We assume that we are given a set of ground-truth coordinate pairs $\mathcal{M} = \{(\mathbf{k}_m, \hat{\mathbf{k}}_m)\}_{m=1}^M$ for each training image pair, where M is the number of annotated keypoint matches. We carry out keypoint transfer from the source to the target keypoints using the constructed dense flow field. For a given keypoint $\mathbf{k} = (x_k, y_k)$, we define a soft sampler $\mathbf{W}^{(k)} \in \mathbb{R}^{\bar{H} \times \bar{W}}$:

$$\mathbf{W}_{ij}^{(k)} = \frac{\max(0, \tau - \sqrt{(x_k - j)^2 + (y_k - i)^2})}{\sum_{i',j'} \max(0, \tau - \sqrt{(x_k - j')^2 + (y_k - i')^2})}, \quad (11)$$

where τ is a distance threshold, and $\sum_{ij} \mathbf{W}_{ij}^{(k)} = 1$. It can be seen that the soft sampler effectively samples each transferred keypoint $\hat{\mathbf{P}}'_{i,j}$ by assigning weights inversely proportional to the distance to \mathbf{k} . Using this soft sampler, we assign a match to the keypoint \mathbf{k} as $\hat{\mathbf{k}}' = \sum_{(i,j) \in \bar{H} \times \bar{W}} \hat{\mathbf{P}}'_{i,j} \mathbf{W}_{ij}^{(k)}$, being able to achieve up to sub-pixel-wise accurate keypoint matches. By applying this

Method	SPair-71k		PF-PASCAL		PF-WILLOW		time (ms)	memory (GB)	FLOPs (G)
	@ α_{bbox}		@ α_{img}		@ $\alpha_{\text{bbox-kp}}$	@ α_{bbox}			
	0.1 (F)	0.1 (T)	0.05 (F)	0.1 (F)	0.1 (T)	0.1 (T)			
NC-Net [45]	20.1	26.4	54.3	78.9	67.0	-	222	1.2	44.9
DCC-Net [16]	-	26.7	55.6	82.3	73.8	-	567	2.7	47.1
DHPF [38]	27.7	28.5	56.1	82.1	74.1	80.2	58	1.6	2.0
PMD [28]	26.5	-	-	81.2	74.7	-	-	-	-
UCN [5]	-	17.7	-	75.1	-	-	-	-	-
HPF [36]	28.2	-	60.1	84.8	74.4	-	63	-	-
SCOT [30]	35.6	-	63.1	85.4	76.0	-	151	4.6	6.2
SCNet [14]	-	-	36.2	72.2	-	70.4	>1000	-	-
DHPF [38]	37.3	27.4	75.7	90.7	71.0	77.6	58	1.6	2.0
DHPF† [38]	39.4	-	-	-	-	-	58	1.6	2.0
NC-Net* [45]	-	-	-	81.9	-	-	222	1.2	44.9
DCC-Net* [16]	-	-	-	83.7	-	-	567	2.7	47.1
ANC-Net [27]	-	28.7	-	86.1	-	-	216	0.9	44.9
PMD [28]	37.4	-	-	90.7	75.6	-	-	-	-
CHMNet [34]	46.3	<u>30.1</u>	80.1	91.6	69.6	<u>79.4</u>	54	1.6	19.6
PMNC [26]	<u>50.4</u>	-	82.4	90.6	-	-	-	-	-
MMNet [59]	40.9	-	77.6	89.1	-	-	86	-	-
CATs [3]	43.5	-	-	-	-	-	45	1.6	28.4
CATs† [3]	49.9	27.1	75.4	92.6	69.0	79.2	45	1.6	28.4
TransforMatcher (ours)	50.2	30.5	78.9	90.5	66.7	75.1	54	1.6	33.5
TransforMatcher† (ours)	53.7	<u>30.1</u>	<u>80.8</u>	<u>91.8</u>	65.3	76.0	54	1.6	33.5

Table 1. **Performance on standard benchmarks of semantic matching.** Higher PCK is better. All the results reported in the table uses pretrained ResNet-101 model as the feature extractor. Methods in the first group were trained with weak supervision (image pair annotations), while those in the second group were trained with strong supervision (sparse keypoint match annotations). Models with * are retrained using keypoint annotations from ANC-Net [27]. † indicates the use of data augmentation during training. Numbers in bold indicate the best performance, followed by the underlined numbers. Some results are from [34].

keypoint transfer method on the source keypoints, we obtain the predicted keypoint pairs on image $\hat{I} : \{(\mathbf{k}_m, \hat{\mathbf{k}}'_m)\}_{m=1}^M$ by assigning a match $\hat{\mathbf{k}}'_m$ to each keypoint \mathbf{k}_m in the source image. We formulate our training objective to minimize the average Euclidean distance between the predicted target keypoints and the ground-truth target keypoints as follows:

$$\mathcal{L} = \frac{1}{M} \sum_{m=1}^M \|\hat{\mathbf{k}}_m - \hat{\mathbf{k}}'_m\|_2^2. \quad (12)$$

5. Experiments

We evaluate our method on the semantic correspondence task, which aims to match semantically similar parts between images of the same category but different instances.

Datasets. We report our results on standard benchmark datasets of semantic correspondence: SPair-71k [37], PF-PASCAL [13], and PF-WILLOW [12]. The SPair-71k dataset has diverse variations in viewpoint and scale, with 53,340 / 5,384 / 12,234 image pairs for training, validation, and testing, respectively. The PF-PASCAL and PF-WILLOW datasets are taken from four categories of the

PASCAL VOC dataset, having small viewpoint and scale variations. The PF-PASCAL dataset contains 2,940 / 308 / 299 image pairs for training, validation and testing, respectively. The PF-WILLOW dataset contains 900 image pairs for testing only. The SPair-71k dataset is significantly larger than the other two datasets, and has more accurate and richer annotations regarding different levels of difficulty in occlusion, truncation, viewpoint and illumination. Being the most challenging dataset, the results on SPair-71k are less saturated in comparison.

Implementation details. Following recent methods [3,34], we employ the ResNet-101 model pre-trained on the ImageNet classification task [23] as the feature extraction network. Note that the `conv4_x` and `conv5_x` layers in ResNet-101 have 23 and 3 bottleneck layers respectively, from which we extract feature maps to compute 26 layer-wise correlations maps for each image pair. We set the spatial size of the input image to 240×240 , resulting in $H = W = 15$ for feature maps used for correlation computation, and $\bar{H} = \bar{W} = 30$. Each of our match-to-match attention layers have 8 heads for multi-head self attention ($N_h = 8$), with head dimension of 4 ($D_h = D_v = 4$).

Augmentation	Positional Embedding	SPair-71k		PF-PASCAL	
		@ α_{bbox}		@ α_{img}	
		0.05	0.1	0.05	0.1
✓	Absolute [39]	29.9	48.7	74.5	89.4
	Absolute [39]	26.6	48.9	79.4	91.8
	Rotary [48]	<u>30.5</u>	<u>50.2</u>	78.9	90.4
✓	Rotary [48]	32.4	53.7	80.8	91.8

Table 2. **Ablation on augmentation and positional embedding.** The results show that using data augmentation and rotary positional embedding gives the best results.

The overall pipeline of our method is implemented using PyTorch [40], and is optimized using the Adam [21] optimizer with a constant learning rate of 1e-3. We finetune the feature extractor network at a lower learning rate of 1e-5.

Evaluation metric. We use the percentage of correct keypoints (PCK) for evaluation, which is the standard evaluation metric for category-level matching. Given a pair of ground-truth and predicted target keypoints $\mathcal{K} = \{(\hat{\mathbf{k}}_m, \hat{\mathbf{k}}'_m)\}_{m=1}^M$, PCK is measured by:

$$\text{PCK}(\mathcal{K}) = \frac{1}{M} \sum_{m=1}^M \mathbb{1}[\|\hat{\mathbf{k}}_m - \hat{\mathbf{k}}'_m\| \leq \alpha_\tau \cdot \max(w_\tau, h_\tau)], \quad (13)$$

where w_τ and h_τ are the width and height of either the entire image or the object bounding box, *i.e.*, $\tau \in \{\text{img}, \text{bbox-kp}, \text{bbox}\}$, and α_τ is a tolerance factor.

5.1. Results and analysis.




Figure 5. **Sample results on SPair-71k.** Source images are TPS-transformed [8] to target images using predicted correspondences.

Architecture	SPair-71k		time (ms)	memory (GB)	FLOPs (G)
	@ α_{bbox}				
	0.05	0.1			
Transformer [55]	-	-	Out-Of-Memory		
Linformer [56]	0.34	1.3	36	1.7	33.4
Performer [4]	28.2	48.8	88	1.6	35.9
Additive Attn.	26.6	48.9	54	1.6	33.5

Table 3. **Results of different transformer architectures.** Vanilla transformer could not be evaluated within memory capabilities. Additive attention yields the most favorable results.

For the SPair-71k dataset, we evaluate two versions for our model: a finetuned model (F) trained on SPair-71k, and a transferred model (T) trained on PF-PASCAL. On the PF-PASCAL and PF-WILLOW datasets, we follow the common evaluation protocol to train our network on the training split of PF-PASCAL and evaluate on the test splits of PF-PASCAL and PF-WILLOW. The quantitative results are illustrated in Table 1. Previous methods have been using two different schemes, *e.g.*, $\tau \in \{\text{bbox-kp}, \text{bbox}\}$, when computing the threshold for PF-WILLOW [35], so we report our results using both thresholds.

We show that TransforMatcher finetuned on SPair-71k sets a new state of the art. A notable observation is that TransforMatcher finetuned on SPair-71k *without* data augmentation outperforms CATs [3] trained *with* augmentation, proving the efficacy of our 4D match-to-match attention and multi-level correlation score features. Using data augmentations leads to improved PCK on both SPair-71k and PF-PASCAL datasets, but transformer-based models benefit more from augmentations as seen from the lower PCK increase in DHPF [3]. It is interesting that TransforMatcher trained without data augmentations transfer slightly better to SPair-71k and PF-WILLOW datasets than our model trained with data augmentations, albeit its lower PCK performance on PF-PASCAL. This potentially hints that while data augmentations do help TransforMatcher to learn better, it overfits more to the training data domain. TransforMatcher also exhibits state-of-the-art performance when transferred to the SPair-71k dataset, while being comparable on the PF-PASCAL dataset. However, TransforMatcher shows substandard results when transferred to the PF-WILLOW dataset, unlike the SPair-71k dataset. This evidences that the match-to-match interactions learned from the PF-PASCAL dataset is better transferable to the SPair-71k dataset, but is not as effective on the PF-WILLOW dataset. Figure 5 visualizes example qualitative results on SPair-71K using our model.

5.2. Ablation study and analysis

Effect of data augmentation during training. CATs [3] found that using data augmentation for category-level matching model is beneficial, especially for data-hungry

transformer-based architectures. We study the effect of applying data augmentation to our model as well, following the schemes used in CATs. The results in Table 2 show that using data augmentation indeed gives consistent improvements to the performance of our model.

Analysis on positional embedding. We investigate the effect of positional embedding used in our pipeline. As conventional relative positional embedding requires an explicit computation of the attention matrix, is not applicable to our transformer architecture with the linear-complexity additive attention. On the other hand, rotary positional embeddings can be seamlessly applied to our model as an alternative method to model relative positional embedding. The results in Table 2 show that using rotary positional embedding results in significant gains over absolute positional embedding, especially on the more challenging SPair-71k dataset.

Analysis on efficient transformer architecture. We try replacing our match-to-match attention architecture with other efficient transformer designs [4, 56], and also the vanilla transformer [55] design to compare the performances. We use absolute learnable positional embedding in this experiment. The results in Table 3 show that the additive attention architecture shows the most favorable results, with similarly high performance as Performer but with lower latency. We found that the Linformer architecture [56] failed to train, which we conjecture is due to the low head dimension of our network, and the reliance of Linformer on kernel approximations which could lead to inaccurate interactions between the position-sensitive matches. Training with vanilla Transformers was infeasible due to its large memory demands of the pair-wise attention matrices.

Analysis on nonlocality of match-to-match attention. For an in-depth analysis, we investigate how nonlocally our match-to-match attention layers operate in comparison to convolutional counterparts [34, 45]. We define the measure of nonlocality of an MHSA at layer l as the average of interactions between attention scores and relative offsets:

$$\Phi^l = \frac{1}{Z} \sum_{h \in [N_h]} \sum_{(\mathbf{q}, \mathbf{k}) \in \mathcal{X} \times \mathcal{X}} \mathbf{A}_{\mathbf{q}, \mathbf{k}}^{(h)} \|\mathbf{q} - \mathbf{k}\|^2, \quad (14)$$

where Z is normalization constant and \mathcal{X} is a set of spatial positions in \mathbb{C} . Figure 6 plots distributions of nonlocality values for high-dim convolutional layers and MHSA layers in TransforMatcher; convolutional layers *statically* transform with *fixed, local* receptive fields ($\Phi_{\text{conv}}^K < 8$) regardless of input contents. In contrast, TransforMatcher layers can *dynamically* transform input contents by *adaptively* deciding regions of attention for effective transformation with *global* receptive fields ($\Phi_{\text{TM}}^l \approx 12.5$). To verify the benefits of dynamic global match-to-match attention, we measure sample-wise nonlocality ($\Phi = \sum_{l=1}^L \Phi^l$) for each test image pair in the SPair-71k, assort them into 20 groups with increasing nonlocality, and visualize the proportion of




Figure 6. Nonlocality distributions of high-dim. conv kernels (left) and TransforMatcher’s attention layers (right).




Figure 7. Proportion of image pair difficulty w.r.t. nonlocality.

the difficulty levels for each group in Fig. 7. For all difficulty types, the proportion of hard/medium samples increase with increasing nonlocality. This trend is especially visible in types of truncation/occlusion; our model attends larger contexts to better perceive truncated/occluded parts. We guide the readers to the supplementary material for the implementation details of this analysis, together with additional analyses and qualitative results of TransforMatcher.

6. Conclusion

In this paper, we have proposed the TransforMatcher, an effective semantic matching learner. Our principal contribution is the match-to-match attention mechanism, which is, to the best of our knowledge, the first attempt to directly process a 4D input, *i.e.*, correlation map, with every spatial entry (match) as an element for attention using a transformer-based network with *global* receptive fields. This has been a challenging pursuit due to the quadratic complexity of vanilla transformers in modeling global-range interactions, which was addressed by additive attention with linear complexity. We further propose to treat multi-level correlation scores as features to better exploit the richer semantics in different levels of feature maps. The proposed model outperforms state of the arts on the SPair-71k dataset, while performing on par with the SOTA methods on the PF-PASCAL dataset. While the memory usage of TransforMatcher increases quadratically with respect to the number of pixels as in other dense matching methods, we anticipate this work will motivate the use of transformers with high-dimensional inputs in other domains.

Acknowledgement. This work was supported by Samsung Advanced Institute of Technology (SAIT) and also by the NRF grant (NRF-2021R1A2C3012728) and the IITP grants (No.2021-0-02068: AI Innovation Hub, No.2019-0-01906: Artificial Intelligence Graduate School Program at POSTECH) funded by the Korea government (MSIT).

References

- [1] Herbert Bay, Tinne Tuytelaars, and Luc Van Gool. Surf: Speeded up robust features. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2006. 1
- [2] Minsu Cho, Suha Kwak, Cordelia Schmid, and Jean Ponce. Unsupervised object discovery and localization in the wild: Part-based matching with bottom-up region proposals. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 2
- [3] Seokju Cho, Sunghwan Hong, Sangryul Jeon, Yunsung Lee, Kwanghoon Sohn, and Seungryong Kim. Semantic correspondence with transformers. *arXiv preprint arXiv:2106.02520*, 2021. 1, 2, 6, 7
- [4] Krzysztof Choromanski, Valerii Likhoshesterov, David Dohan, Xingyou Song, Andreea Gane, Tamas Sarlos, Peter Hawkins, Jared Davis, Afroz Mohiuddin, Lukasz Kaiser, et al. Rethinking attention with performers. *arXiv preprint arXiv:2009.14794*, 2020. 3, 7, 8
- [5] Christopher Choy, JunYoung Gwak, Silvio Savarese, and Manmohan Chandraker. Universal correspondence network. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2016. 6
- [6] Navneet Dalal and Bill Triggs. Histograms of oriented gradients for human detection. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2005. 1
- [7] Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superpoint: Self-supervised interest point detection and description. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 224–236, 2018. 1
- [8] Gianluca Donato and Serge Belongie. Approximate thin plate spline mappings. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2002. 7
- [9] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021. 2
- [10] Mihai Dusmanu, Ignacio Rocco, Tomas Pajdla, Marc Pollefeys, Josef Sivic, Akihiko Torii, and Torsten Sattler. D2-net: A trainable cnn for joint detection and description of local features. *arXiv preprint arXiv:1905.03561*, 2019. 1
- [11] David Forsyth and Jean Ponce. *Computer Vision: A Modern Approach. (Second edition)*. Prentice Hall, Nov. 2011. 1
- [12] Bumsu Ham, Minsu Cho, Cordelia Schmid, and Jean Ponce. Proposal flow. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 1, 6
- [13] Bumsu Ham, Minsu Cho, Cordelia Schmid, and Jean Ponce. Proposal flow: Semantic correspondences from object proposals. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2018. 1, 6
- [14] Kai Han, Rafael S Rezende, Bumsu Ham, Kwan-Yee K Wong, Minsu Cho, Cordelia Schmid, and Jean Ponce. Sc-net: Learning semantic correspondence. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2017. 1, 6
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 4
- [16] Shuaiyi Huang, Qiuyue Wang, Songyang Zhang, Shipeng Yan, and Xuming He. Dynamic context correspondence network for semantic alignment. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2019. 1, 2, 6
- [17] Sangryul Jeon, Seungryong Kim, Dongbo Min, and Kwanghoon Sohn. Parn: Pyramidal affine regression networks for dense semantic correspondence. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 1
- [18] Sangryul Jeon, Dongbo Min, Seungryong Kim, Jihwan Choe, and Kwanghoon Sohn. Guided semantic flow. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 1, 2
- [19] Wei Jiang, Eduard Trulls, Jan Hosang, Andrea Tagliasacchi, and Kwang Moo Yi. COTR: Correspondence Transformer for Matching Across Images. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2021. 2
- [20] Seungryong Kim, Stephen Lin, Sangryul Jeon, Dongbo Min, and Kwanghoon Sohn. Recurrent transformer networks for semantic correspondence. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018. 1
- [21] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2015. 7
- [22] Nikita Kitaev, Łukasz Kaiser, and Anselm Levskaya. Reformer: The efficient transformer. *arXiv preprint arXiv:2001.04451*, 2020. 3
- [23] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2012. 2, 6
- [24] Jongmin Lee, Yoonwoo Jeong, Seungwook Kim, Juhong Min, and Minsu Cho. Learning to distill convolutional features into compact local descriptors. In *2021 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 897–907, 2021. 2
- [25] Junghyup Lee, Dohyung Kim, Jean Ponce, and Bumsu Ham. Sfnet: Learning object-aware semantic correspondence. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 5
- [26] Jae Yong Lee, Joseph DeGol, Victor Fragoso, and Sudipta Sinha. Patchmatch-based neighborhood consensus for semantic correspondence. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2, 6
- [27] Shuda Li, Kai Han, Theo W. Costain, Henry Howard-Jenkins, and Victor Prisacariu. Correspondence networks with adaptive neighbourhood consensus. In *Proceedings of*

- the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 1, 2, 6
- [28] Xin Li, Deng-Ping Fan, Fan Yang, Ao Luo, Hong Cheng, and Zicheng Liu. Probabilistic model distillation for semantic correspondence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7505–7514, June 2021. 6
- [29] Xinghui Li, Kai Han, Shuda Li, and Victor Prisacariu. Dual-resolution correspondence networks. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2020. 1
- [30] Yanbin Liu, Linchao Zhu, Makoto Yamada, and Yi Yang. Semantic correspondence as an optimal transport problem. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 1, 6
- [31] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. *arXiv preprint arXiv:2103.14030*, 2021. 2
- [32] David G. Lowe. Object recognition from local scale-invariant features. In *Proceedings of the International Conference on Computer Vision (ICCV)*, volume 2, pages 1150–1157 vol.2, 1999. 1
- [33] David G. Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision (IJCV)*, 2004. 1
- [34] Juhong Min and Minsu Cho. Convolutional hough matching networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2940–2950, June 2021. 1, 2, 6, 8
- [35] Juhong Min, SeungWook Kim, and Minsu Cho. Convolutional hough matching networks for robust and efficient visual correspondence. *arXiv preprint arXiv:2109.05221*, 2021. 2, 7
- [36] Juhong Min, Jongmin Lee, Jean Ponce, and Minsu Cho. Hyperpixel flow: Semantic correspondence with multi-layer neural features. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2019. 1, 6
- [37] Juhong Min, Jongmin Lee, Jean Ponce, and Minsu Cho. SPair-71k: A large-scale benchmark for semantic correspondence. *arXiv preprint arXiv:1908.10543*, 2019. 6
- [38] Juhong Min, Jongmin Lee, Jean Ponce, and Minsu Cho. Learning to compose hypercolumns for visual correspondence. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 1, 2, 6
- [39] Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. fairseq: A fast, extensible toolkit for sequence modeling. In *Proceedings of NAACL-HLT 2019: Demonstrations*, 2019. 7
- [40] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems (NeurIPS)*, pages 8024–8035. Curran Associates, Inc., 2019. 7
- [41] Jerome Revaud, Philippe Weinzaepfel, César De Souza, Noe Pion, Gabriela Csurka, Yohann Cabon, and Martin Humenberger. R2d2: Repeatable and reliable detector and descriptor. *arXiv preprint arXiv:1906.06195*, 2019. 1
- [42] Ignacio Rocco, Relja Arandjelovic, and Josef Sivic. Convolutional neural network architecture for geometric matching. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1
- [43] Ignacio Rocco, Relja Arandjelović, and Josef Sivic. End-to-end weakly-supervised semantic alignment. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1
- [44] Ignacio Rocco, Relja Arandjelović, and Josef Sivic. Efficient neighbourhood consensus networks via submanifold sparse convolutions. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 1
- [45] Ignacio Rocco, Mircea Cimpoi, Relja Arandjelović, Akihiko Torii, Tomas Pajdla, and Josef Sivic. Neighbourhood consensus networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018. 1, 2, 4, 6, 8
- [46] Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. SuperGlue: Learning feature matching with graph neural networks. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [47] Robin Strudel, Ricardo Garcia, Ivan Laptev, and Cordelia Schmid. Segmenter: Transformer for semantic segmentation. *arXiv preprint arXiv:2105.05633*, 2021. 2
- [48] Jianlin Su, Yu Lu, Shengfeng Pan, Bo Wen, and Yunfeng Liu. Roformer: Enhanced transformer with rotary position embedding. *arXiv preprint arXiv:2104.09864*, 2021. 5, 7
- [49] Jiaming Sun, Zehong Shen, Yuang Wang, Hujun Bao, and Xiaowei Zhou. LoFTR: Detector-free local feature matching with transformers. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [50] Tatsunori Tanai, Sudipta N Sinha, and Yoichi Sato. Joint recovery of dense correspondence and cosegmentation in two images. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2
- [51] Yurun Tian, Xin Yu, Bin Fan, Fuchao Wu, Huub Heijnen, and Vassileios Balntas. Sosnet: Second order similarity regularization for local descriptor learning. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11016–11025, 2019. 1
- [52] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Herve Jegou. Training data-efficient image transformers & distillation through attention. In *Proceedings of the International Conference on Learning Representations (ICLR)*, volume 139, pages 10347–10357, July 2021. 2
- [53] Prune Truong, Martin Danelljan, and Radu Timofte. GLU-Net: Global-local universal network for dense flow and correspondences. In *Proceedings of the IEEE International*

Conference on Computer Vision and Pattern Recognition (CVPR), 2020. [1](#)

- [54] Ashish Vaswani, Prajit Ramachandran, Aravind Srinivas, Niki Parmar, Blake Hechtman, and Jonathon Shlens. Scaling local self-attention for parameter efficient visual backbones. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12894–12904, June 2021. [2](#)
- [55] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017. [3](#), [7](#), [8](#)
- [56] Sinong Wang, Belinda Z Li, Madian Khabsa, Han Fang, and Hao Ma. Linformer: Self-attention with linear complexity. *arXiv preprint arXiv:2006.04768*, 2020. [3](#), [7](#), [8](#)
- [57] Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. *arXiv preprint arXiv:2102.12122*, 2021. [2](#)
- [58] Chuhan Wu, Fangzhao Wu, Tao Qi, and Yongfeng Huang. Fastformer: Additive attention can be all you need. *arXiv preprint arXiv:2108.09084*, 2021. [3](#), [4](#)
- [59] Dongyang Zhao, Ziyang Song, Zhenghao Ji, Gangming Zhao, Weifeng Ge, and Yizhou Yu. Multi-scale matching networks for semantic correspondence. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. [6](#)