

# Learning Second-Order Attentive Context for Efficient Correspondence Pruning

Xinyi Ye, Weiyue Zhao, Hao Lu, Zhiguo Cao\*

Key Laboratory of Image Processing and Intelligent Control, Ministry of Education School of Artificial Intelligence and Automation, Huazhong University of Science and Technology  
{xinyiye,zhaoweiyue,hlu,zgcao}@hust.edu.cn

## Abstract

Correspondence pruning aims to search consistent correspondences (inliers) from a set of putative correspondences. It is challenging because of the disorganized spatial distribution of numerous outliers, especially when putative correspondences are largely dominated by outliers. It is more challenging to ensure effectiveness while maintaining efficiency. In this paper, we propose an effective and efficient method for correspondence pruning. Inspired by the success of attentive context in correspondence problems, we first extend the attentive context to the first-order attentive context and then introduce the idea of attention in attention (ANA) to model second-order attentive context for correspondence pruning. Compared with first-order attention that focuses on feature-consistent context, second-order attention dedicates to attention weights itself and provides an additional source to encode consistent context from the attention map. For efficiency, we derive two approximate formulations for the naive implementation of second-order attention to optimize the cubic complexity to linear complexity, such that second-order attention can be used with negligible computational overheads. We further implement our formulations in a second-order context layer and then incorporate the layer in an ANA block. Extensive experiments demonstrate that our method is effective and efficient in pruning outliers, especially in high-outlier-ratio cases. Compared with the state-of-the-art correspondence pruning approach LMCNet, our method runs 14 times faster while maintaining a competitive accuracy.

## Introduction

Finding correspondences is a fundamental problem in many computer vision tasks, such as structure-from-motion (Schönberger and Frahm 2016), inpainting (Zhou et al. 2021), and simultaneous location and mapping (Mur-Artal, Montiel, and Tardós 2015). However, correspondences searched with existing detector-based descriptors (DeTone, Malisiewicz, and Rabinovich 2018; Mishchuk et al. 2017; Lowe 2004; Rublee et al. 2011; Ono et al. 2018) abound with outliers (Fig. 1, 1<sup>st</sup> column). To prevent outliers from disturbing downstream tasks, correspondence pruning

\*Corresponding author

Code will be available at <https://github.com/DIVE128/ANANet>  
Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

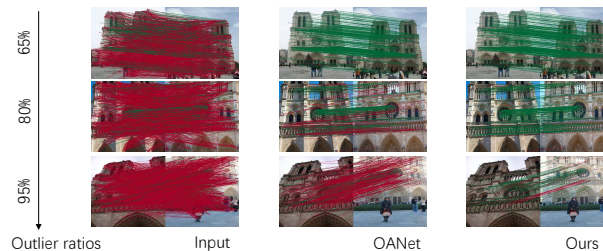


Figure 1: Imbalanced putative correspondences. Given a set of correspondences (1<sup>st</sup> column) with increased outlier ratios (from top to bottom, 65%, 80%, and 95%), OANet loses efficacy gradually and fails to reject mass outliers (red), while our method keeps (3<sup>rd</sup> column) a high accuracy.

(or consistency filtering) is by default used as a postprocessing procedure to identify consistent correspondences (inliers) (Bian et al. 2017; Ma et al. 2019; Yi et al. 2018).

Recently, deep models (Yi et al. 2018; Sun et al. 2020; Zhang et al. 2019) are proposed as a plug-in to prune outliers and report promising results. Unfortunately, as outlier ratios increase, they lose efficacy (Liu et al. 2021; Zhao et al. 2021) and even identify outliers as inliers (Fig. 1, 2<sup>nd</sup> column). Liu et al. (Liu et al. 2021) propose LMCNet that leverages the motion coherence to distinguish correspondences with discriminative eigenvectors. However, when the singular value decomposition deals with a large matrix, it can be computationally expensive: For correspondence pruning, the size of a matrix used to model point-to-point relation is often larger than  $10^3 \times 10^3$  and thus leads to one-order-of-magnitude extra inference time. This may be problematic when facing practical applications, especially the real-time ones. Here we explore how to tackle the correspondence pruning problem both effectively and efficiently.

In correspondence pruning, a fundamental problem is *how to distinguish inliers from outliers*. Previous works (Zhang et al. 2019; Sun et al. 2020) believe inliers are consistent in both local and global context so that local-global context is encoded to help to distinguish inliers from outliers. However, Zhao et al. (Zhao et al. 2019) points out local context would be unstable due to randomly distributed outliers. Global context suffers from the same affliction because inliers would be submerged in massive outliers. This results in

the problem of inconsistent context. The similar inputs are more likely to lead to similar outputs and it is perhaps one of the reasons why these methods lose their efficacy when massive outliers occur. A question of interest is *how to efficiently encode consistent context*.

ACNe (Sun et al. 2020) suggests that the attention mechanism can guide a network to focus on a subset of inlier features and thus can suppress outliers. ACNe uses this mechanism to generate the so-called attentive context that is consistent for the inliers. Following this spirit, we delve deep into attentive context and observe that the naive attentive context modeled by feature-based similarity, which could be interpreted as first-order attention, is not sufficient. Instead we find taking additional second-order attention as a supplementary is helpful. This finding is based on an observation that, inliers share similar attention weights, while outliers do not. We therefore introduce the idea of attention in attention to model second-order attentive context. The key idea is to introduce extra distinctive information from attention weights such that inliers can be grouped and distinguished.

To this end, we propose to jointly model first- and second-order attentive context, featured by feature- and attention-consistent context, respectively. In particular, we first encode with a first-order attention map generated by self-attention (Vaswani et al. 2017). We then seek attention-consistent context from a second-order similarity matrix generated by attention weights. For efficient end-to-end training, we discuss the most effective part of the naive implementation of second-order attention, optimize the cubic-complexity implementation to linear complexity, and propose a Second-Order Context (SOC) layer. By integrating the SOC layer into a permutation-invariant block, we further present an iterative Attention in Attention Network (ANA-Net) that can effectively address correspondence pruning. We show that this network is effective in pruning outliers, especially in high-outlier-ratio cases (Fig. 1, 3<sup>rd</sup> column).

We conduct extensive experiments to demonstrate the effectiveness and generalization of our approach on camera pose estimation and correspondence pruning. Thanks to the consistent context, ANA-Net achieves superior performance over most baselines and shows its potentials to address multi consistency problems. Meanwhile, ANA-Net maintains comparable performance against LMCNet while running 14 times faster than it.

**Contributions.** We introduce the idea of attention in attention that essentially models similarity between attention weights, which can also be interpreted as second-order attention. Technically, we present a formulation that models such second-order information and further derive two approximated formulations to reduce computational complexity. We show that our formulations can be implemented as a network layer and can be used to cooperate the first-order attention to address the problem of correspondence pruning with high accuracy, low computational cost and robust generalization.

## Related Work

**Deep Correspondence Pruning.** RANSAC (Fischler and Bolles 1981) and its variants (Chum, Werner, and Matas

2005; Cavalli et al. 2020; Barath et al. 2019; Barath and Matas 2018; Barath et al. 2020) always suffers from high-outlier-ratio cases, thus deep filters are critical to reject most outliers and guarantee a more manageable correspondence set for RANSAC. And pioneering works such as DSAC (Brachmann et al. 2017) and CNe (Yi et al. 2018) prove the superiority of neural networks for pruning correspondences. Since sparse correspondences are unordered and irregular, convolution operators (Krizhevsky, Sutskever, and Hinton 2012) cannot be adopted directly. Inspired by GNN (Scarselli et al. 2008), OANet (Zhang et al. 2019) proposes the generalized differentiable pooling (Ying et al. 2018) and unpooling to cluster correspondences. To improve the robustness of CNe (Yi et al. 2018), ACNe (Sun et al. 2020) introduces the attention mechanism to focus the normalization of the feature maps of inliers. According to the affine attributes, NM-Net (Zhao et al. 2019) mines compatibility-specific neighbors to aggregate features with a compatible Point-Net++ (Qi et al. 2017) architecture. Despite achieving satisfactory performance in correspondence pruning, they can still suffer from high-outlier-ratio cases.

**Consensus in Correspondences.** One cannot classify an isolated correspondence as inlier or outlier (Zhao et al. 2021). Consensus is the motivation for algorithms to distinguish outliers. There are generally two types of classical consensus in correspondences: global consensus and local consensus. RANSAC-related work is driven by the former consensus and identifies correct correspondences by judging whether correspondences conform to a task-specific geometric model such as epipolar geometry or homography (Liu et al. 2021). The latter is more intuitive. For instance, LPM (Ma et al. 2019) assumes local geometry around correct correspondences does not change freely. CFM (Chen, Lin, and Chen 2015) believes local structural or local textural variations estimated by a consistent pair of transformations should be similar. In this paper, the feature-consistent context generated with the first-order attention is an application of both types of consensus, while the attention-consistent context captured with second-order attention does not belong to either of them. It is dedicated to attention itself and we view it as a representation of attention consensus.

**Self-Attention Mechanism.** Inspired by the success of self-attention mechanism in NLP (Bahdanau, Cho, and Bengio 2014; Vaswani et al. 2017; Kim et al. 2017), many researchers notice its ability to focus on input of interest and to capture long-range context. Wang *et al.* (Wang et al. 2018) combines CNN-like architectures with self-attention and achieves non-local feature representation. Pan *et al.* (Pan et al. 2021) builds the cross context of view boundaries via the attention mechanism in image cropping. Wang *et al.* (Wang et al. 2022) models pixel-level attention for weak small object detection. Carion *et al.* (Carion et al. 2020) and Misra *et al.* (Misra, Girdhar, and Joulin 2021) adopt the attention mechanism to reason about the relation of objects in 2D and 3D detection. For correspondence pruning, ACNe (Sun et al. 2020) focuses on feature normalization w.r.t. a subset of inliers. Due to the non-local property of the self-attention mechanism, in this work we validate its use to deal with high-outlier-ratio cases in correspondence pruning.

## Learning Consistent Context

### Problem Formulation

Given an image pair  $(I^s, I^t)$ , a putative correspondence set  $C \in \mathbb{R}^{N \times 4}$  can be established via nearest neighbor feature matching. This set is our input to the correspondence pruning method. The goal of correspondence pruning is to predict a weighted vector  $w = [w_1, \dots, w_N] \in \mathbb{R}^{N \times 1}$ , where  $w_i \in [0, 1)$  indicates the inlier probability of  $i^{\text{th}}$  correspondence. The input correspondence is conventionally a concatenated 4D vector representing pairwise keypoints coordinates. However, the disorganized spatial distribution of numerous outliers makes it difficult to directly filter unreliable correspondences with a classical method such as RANSAC. To distinguish between inliers and outliers in the high-dimensional feature space, a few learnable methods (Yi et al. 2018; Liu et al. 2021; Zhao et al. 2019; Zhang et al. 2019) adopt an iterative strategy with multiple MLP layers to capture contextual information for better outlier rejection.

### First-Order Attention

Compared with the MLP-based techniques that integrate consistency information with restricted neighbor correspondences, ACNe (Sun et al. 2020) attempts to enlarge the receptive field with the attention mechanism. The attention used by ACNe can be thought of as first-order attention, because it essentially models the similarity between feature maps. By normalizing the feature map with an attentive weight function, ACNe could effectively cluster features and thus separate inliers from outliers. Inspired by (Sun et al. 2020), we consider optimizing the feature map for better classification by fusing attentive context. As shown in Fig. 2, each inlier shares the same motion coherence with certain correspondences. In other word, correspondences with similar motion trends are likely to focus on each other. Therefore, we instead use self-attention to guide features to encode consistent context for inliers. It is worth noting that, self-attention is also first-order attention, despite the fact that it is represented by a 2D map. Next we show how to encode correspondence features to obtain consistent context with a first-order attention map.

Given a first-order attention map  $A = [a_1, \dots, a_N] \in \mathbb{R}^{N \times N}$ , where  $\sum_j a_{ij} = 1$  denotes the feature similarity between the  $i^{\text{th}}$  correspondence and the correspondence set, we can acquire the feature-consistent context  $v_i$  by  $v_i = \sum_j a_{ij} x_j$  where  $x_j$  is the feature vector of the  $j^{\text{th}}$  correspondence. Essentially, feature-consistent context exploits the motion-consistent nature of inliers to filter outliers that have irregular motion distributions. Considering that the inliers have similar feature patterns, a question of interest is whether inliers possess the latent similarity of attention weights. To delve attention-consistent context among correspondences, we introduce second-order attention.

### Second-Order Attention

From Fig. 2, we notice that different inliers share mutual attentive correspondences (yellow), which suggests inliers may also share similar attention patterns. A key idea of our method is to explore attention-consistent context with

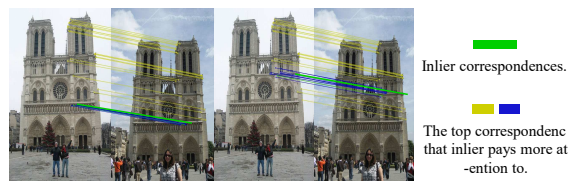


Figure 2: Motion coherence of inliers indicated by first-order attention. We select two inliers green (Left) and (Right) in the same image pair. To visualize the top correspondences that they care about, we color them according to the attention weights respectively. Specifically, yellow denotes the mutual correspondences shared by different inliers, and the rest are shown in blue. It suggests inliers pay more attention to each other with higher attention weights.

a second-order attention map as an explicit measurement of attentive consistency.

Given the first-order attention map  $A$ , we obtain the second-order attention map by  $W = A^T A \in \mathbb{R}^{N \times N}$ , where  $w_{ij} = \sum_k a_{ki} a_{kj}$  indicates the similarity of attention weights between the  $i^{\text{th}}$  and  $j^{\text{th}}$  correspondences. Our goal is to extract discriminative features from  $W$  as the attention-consistent context to separate the inliers from outliers. Since the matrix  $W$  shares the same spirit of spectral clustering, a natural idea is to follow the formulation of spectral clustering to extract features. In general, spectral clustering defines a degree matrix  $D = \text{diag}([d_i = \sum_j w_{ij}])$  of  $W$  and a Laplacian matrix  $L = D - W$ , to derive classifiable eigenvectors by applying singular value decomposition (SVD) on  $L$ . Due to the fact that spectral clustering is based on the theory of graph cuts, it cannot be directly applied to our problem since the given inconsistent motions among outliers. A huge time consumption is introduced as well because of the large matrix decomposition, which deviates from our goal. Yet, it still inspires us to extract discriminative context encoded in  $L$  via a differentiable operation.

As described in OANet (Zhang et al. 2019), the operation ought to be permutation-invariant such that the context is permutation-invariant, which is critical for correspondence pruning. Inspired by Diff-pool in OANet (Zhang et al. 2019), where input features are mapped to clusters, we utilize inner product like Diff-pool to map  $L \in \mathbb{R}^{N \times N}$  into a feature map with fixed dimension to be embedded into the network by the following proposition.

**Proposition 1 (Cubic Form.)** *Let  $l_i$  denote the  $i^{\text{th}}$  column vector of  $L$ . The distinctive information of the  $i^{\text{th}}$  correspondence is encoded by*

$$h_i = \sqrt{l_i^T l_i} = \sqrt{\left(\sum_{j, j \neq i} w_{ij}\right)^2 + \sum_{j, j \neq i} w_{ij}^2}, \quad (1)$$

where  $h_i$  is called the attention-consistent context.  $h_i$  requires  $N^2 + N + 1$  multiplications to be computed from the first-order attention map. For  $N$  correspondences, the computational complexity (ignoring the impact of summation) of the cubic form is  $\mathcal{O}(N^3)$ .

Essentially, Eq. (1) executes a cumulative operation on elements of  $\mathbf{W}$  except the self-correlation  $w_{ii}$ . There is an understanding derived from Fig. 2: inliers have more  $w$ 's of large values while outliers do not, which suggests the values of Eq. (1) for inliers will be different from outliers. In fact, Eq. (1) achieves the same thing like the SVD operation in spectral clustering: extracting classifiable features for inputs to categorize. With the priors from the observation in Fig. 2, it performs with relatively fewer computational overheads. However, the computational overheads of it are still heavy for correspondence pruning because of the cubic complexity. A question remains: *Is it possible to simplify the form while retaining distinguishability*. The answer is yes. Assume that  $y_i = 1$  in the Cauchy's Inequality  $(\sum_i x_i^2)(\sum_i y_i^2) \geq (\sum_i x_i y_i)^2$ , we have

$$N \sum_i x_i^2 \geq (\sum_i x_i)^2. \quad (2)$$

Meanwhile, since  $w_{ij} > 0$ , we have  $(\sum_{j,j \neq i} w_{ij})^2 \geq \sum_{j,j \neq i} w_{ij}^2$ . By inserting them into Eq. (2), one can derive the upper and lower bound of Eq. (1) as

$$\sqrt{2} \sum_{j,j \neq i} w_{ij} \geq h_i \geq \sqrt{\frac{N+1}{N}} \sum_{j,j \neq i} w_{ij}. \quad (3)$$

In this work, we use the upper bound as an approximation and obtain the quadratic formulation.

**Proposition 2 (Quadratic Form.)** *Let  $w_{ij}$  denote the similarity of attention weights between the  $i^{\text{th}}$  and  $j^{\text{th}}$  correspondences such that  $w_{ij} = \sum_k a_{ki} a_{kj}$ . Eq. (1) has a quadratic form, defined by*

$$h_i = \sqrt{2} \sum_{j,j \neq i} w_{ij} = \sqrt{2} (\sum_k a_{ki} - \sum_k a_{ki}^2). \quad (4)$$

Since  $h_i$  requires  $N+1$  multiplications to compute, the computational complexity of the quadratic form for all correspondences is  $\mathcal{O}(N^2)$ .

While the quadratic form has significantly reduced the computational complexity from  $\mathcal{O}(N^3)$  to  $\mathcal{O}(N^2)$ , we can further approximate the quadratic form to a linear form.

**Proposition 3 (Linear Form.)** *Let the upper bound of Eq. (4) be used for a second time according to Eq. (2). A linear approximation of Eq. (4) takes the form*

$$h_i = \sqrt{2} (\sum_k a_{ki} - \frac{1}{N} (\sum_k a_{ki})^2). \quad (5)$$

The computational complexity of the linear form is  $\mathcal{O}(N)$ .

We can expand  $\sum_k a_{ki}$  as:  $\sum_k a_{ki} = \sum_k (a_{ki} * 1) = \sum_k \sum_j a_{ki} a_{kj} = \sum_j \sum_k a_{ki} a_{kj} = \sum_j w_{ij} = w_{ii} + \sum_{j,j \neq i} w_{ij}$ , where  $\sum_{j,j \neq i} w_{ij}$  is the first term of Eq. (1), and  $\sum_k a_{ki}$  is the first term of Eqs. (4)&(5). Despite the non-tight upper bound, the distinguishability of Eqs. (4)&(5) is still reflected in the numerical differences like Eq. (1). They are thus suitable for encoding second-order attentive context without much accuracy loss. For end-to-end training, we present a Second-Order Context (SOC) layer to implement the formulations above.

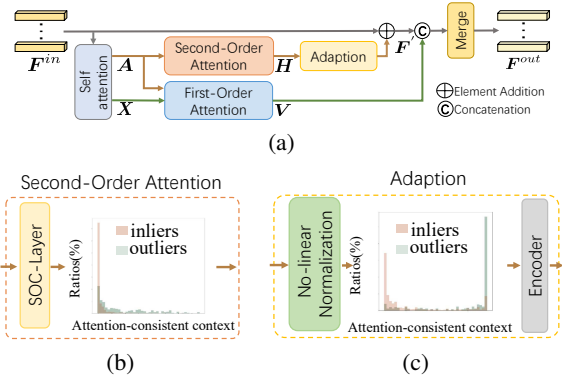


Figure 3: Attention in attention block (ANA-Block). (a) is the overview of ANA-Block. It includes two components: attention-consistent context enhancement and feature-consistent context enhancement. (b) illustrates the process to extract attention-consistent context with second-order attention. (c) shows how the attention-consistent context is embedded into the ANA-Block.

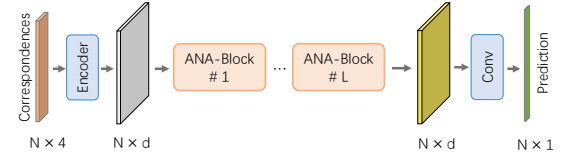


Figure 4: Attention in Attention Network (ANA-Net).  $N$  represents the number of putative correspondences and 4 denotes the concatenated 4D vector a pair of keypoints coordinates in correspondences.

### Attention in Attention Block

Here we present the Attention in Attention Block (ANA-Block), aiming to encode both feature- and attention-consistent context. The feature-consistent context can be obtained from any off-the-shelf first-order attention, and the attention-consistent context is extracted by a SOC layer implementing the second-order attention formulations. The overview is shown in Fig. 3(a).

Given a feature map  $\mathbf{F}^{in} = [\mathbf{f}_1^{in}, \dots, \mathbf{f}_N^{in}] \in \mathbb{R}^{N \times d}$ , the feature-consistent context  $\mathbf{V} = [\mathbf{v}_1, \dots, \mathbf{v}_N] \in \mathbb{R}^{N \times d}$  and the attention-consistent context  $\mathbf{H} = [h_1, \dots, h_N] \in \mathbb{R}^{N \times 1}$  can be obtained with the first-order and second-order attention, respectively. In particular,  $\mathbf{H}$  is obtained by the SOC layer, which defines a mapping function  $f_{SOC}$  such that  $\mathbf{H} = f_{SOC}(\mathbf{A})$ , where  $\mathbf{A}$  is the first-order attention map.  $f_{SOC}$  can be implemented by Eq. (1), Eq. (4), or Eq. (5). Since the attention-consistent context  $\mathbf{H}$  has different characteristics from the initial feature map, we choose to embed it into the feature map rather than concatenating it. In this way, the embedded feature map encodes a hybrid representation  $\mathbf{F}^{out} = [\mathbf{f}_1^{out}, \dots, \mathbf{f}_N^{out}] \in \mathbb{R}^{N \times d}$  for correspondence pruning, which is given by

$$\mathbf{f}_i^{out} = \text{merge}([\mathbf{f}_i^{in} + \psi(h_i)] \parallel \phi(\mathbf{v}_i)), \quad (6)$$

where  $[\cdot \parallel \cdot]$  denotes the concatenation operator,  $\psi(\cdot)$  is a

Dataset	Method	AUC@5°	AUC@10°	AUC@20°
YFCC.	Magsac	28.24	44.86	61.53
	LPM	10.48	18.91	29.26
	GMS	19.05	32.35	46.79
	CODE	16.99	30.23	43.85
	LMF	16.59	29.14	43.41
	CNe	25.18	41.52	57.36
	ACNe	28.90	47.15	63.95
	OANet	29.58	47.71	64.09
	LMCNet	<b>33.81</b>	<b>53.10</b>	<b>69.97</b>
	ANA-Net	<u>33.11</u>	<u>52.02</u>	68.77
SUN.	LPM	2.81	7.4	15.36
	GMS	4.36	11.08	21.68
	CODE	3.52	8.91	18.32
	LMF	3.39	8.91	18.14
	CNe	5.17	13.18	25.74
	ACNe	4.80	12.26	23.75
OANet	4.81	12.52	24.52	
LMCNet	<u>5.78</u>	<u>14.61</u>	<u>27.94</u>	
ANA-Net	<b>5.98</b>	<b>14.95</b>	<b>28.35</b>	

Table 1: Camera pose evaluation on YFCC100M and SUN3D. The best performance is in **bold**, and the second best is underlined.

Method	CNe	ACNe	OANet	LMCNet	ANA-Net	Ref.
runtimes	6.28	11.52	17.30	212.99	14.36	159.68

Table 2: Inference time (ms). Ref. indicates SIFT&NN.

preprocessing step (discussed below) before embedding attention-consistent context into an ANA-Block,  $\phi(\cdot)$  denotes a fusion function for feature-consistent context such as a shared MLP,  $\text{merge}(\cdot)$  is also a MLP that aggregates context and reduces the dimension from  $2d$  to  $d$ . Next we discuss this process in detail, which is featured by attention-consistent context enhancement and feature-consistent context enhancement, as shown in Fig. 3.

**Attention-Consistent Context Enhancement.** Here we explain  $\psi(\cdot)$  in detail. As shown in Fig. 3(b), we use a histogram to visualize attention-consistent context distribution for inliers and outliers, respectively. Since the distribution of inliers is dispersed, we design a non-linear normalization to cluster inliers when  $h_i \rightarrow +\infty$  using an  $\alpha$ -parameterized sigmoid function such that

$$\text{sigmoid}_\alpha(h_i) = \frac{1}{1 + \exp(-\alpha h_i)}, \quad (7)$$

where the learnable parameter  $\alpha$  is used to adjust the level of non-linearity. Moreover, since the dimensionality of normalized  $h_i$  does not equal to the original feature, we further map the normalized  $h_i$  through an encoder to match the dimensionality. The enhanced feature  $\mathbf{f}'_i \in \mathbb{R}^{N \times d}$  is given by

$$\mathbf{f}'_i = \mathbf{f}_i^{in} + \psi(h_i) = \mathbf{f}_i^{in} + \text{encoder}(\text{sigmoid}_\alpha(h_i)). \quad (8)$$

**Feature-Consistent Context Enhancement.** Here we embody  $\phi(\cdot)$ . To further integrate the feature-consistent context with  $\mathbf{f}'_i$ , we choose a MLP to implement  $\phi$  such that

$$\mathbf{f}_i^{out} = \text{merge}([\mathbf{f}'_i || \text{MLP}(\mathbf{v}_i)]). \quad (9)$$

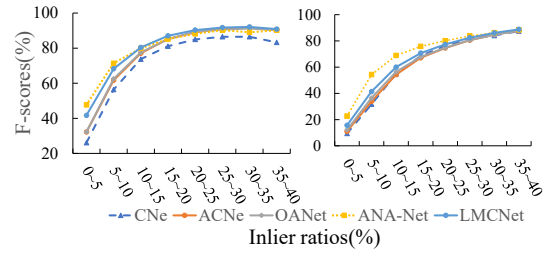


Figure 5: Performances with various inlier ratios. For different methods, we calculate F1-scores with various inlier ratios on both YFCC100M (left) and SUN3D (right).

Meanwhile, (Sun et al. 2020; Zhang et al. 2019) emphasize that the permutation-invariance plays a crucial role in tasks with unordered data, such as correspondence pruning. It should be noted that our ANA-Block designed with consistent context is permutation-invariant. The proof of permutation-invariance can be found in the supplementary.

## Implementation Details

In our implementation, we use the linear-form formulation for the SOC layer, and a multi-head structure in self attention (Vaswani et al. 2017) with 4 heads is used in ANA-Blocks to generate the first-order attention map. For feature dimensionality, we follow a common setting  $d = 128$  for intermediate feature maps. Besides, we apply an iterative strategy to construct the framework by stacking  $L = 5$  ANA-Blocks (Fig. 4). The correspondence set  $\mathcal{C}$  is normalized with camera intrinsic matrices if available, or is normalized into  $[-1, 1]$  according to the input image size in both training and testing.

**Loss Function.** We use a weighted binary cross-entropy loss to balance positive and negative samples following (Yi et al. 2018), which takes the form

$$\mathcal{L}(\mathcal{C}) = \frac{1}{N} \sum_{i=1}^N \tau_i \text{H}(y_i, \mathcal{S}(o_i)), \quad (10)$$

where  $o_i$  is the output of the  $i^{\text{th}}$  correspondence,  $y_i$  denotes the corresponding ground-truth label according to the epipolar distance and the threshold  $\sigma = 10^{-4}$ ,  $\mathcal{S}(\cdot)$  represents the logistic function used in the binary cross entropy  $\text{H}(\cdot)$ , and  $\tau_i$  is a weight that balances the relative importance between positive and negative samples.

## Experiments

### Datasets and Evaluation Protocols

**Datasets.** The outdoor YFCC100M (Thomee et al. 2016) dataset and the indoor SUN3D (Xiao, Owens, and Torralba 2013) dataset are used in camera pose evaluation. We apply the same data split following OANet (Zhang et al. 2019). The putative input correspondences are generated via nearest features matching of SIFT (Lowe 2004) and SuperPoint (DeTone, Malisiewicz, and Rabinovich 2018) with the maximum number of keypoints of 2000 and 1024, respectively. We consider correspondences with a small symmetric epipolar distance ( $< 10^{-4}$ ) as correct correspondences

Method	Precision	Recall	F1-Scores
CNe	66.53	35.67	43.98
ACNe	68.82	57.57	61.00
OANet	67.80	55.63	59.18
LMCNet	65.08	53.08	56.00
ANA-Net	<b>71.40</b>	<b>69.66</b>	<b>68.34</b>

Table 3: Results on MULTI. Best performance is in boldface.

Method	YFCC.	SUN.
CNe	18.03	5.34
ACNe	19.42	5.05
OANet	19.83	4.75
LMCNet	16.90	1.12
ANA-Net	<b>20.75</b>	<b>5.66</b>

Table 4: Generalization with SuperPoint (AUC@5°).

(inliers). An essential matrix is estimated with a robust estimator on the inliers predicted with evaluated methods. The essential matrix is decomposed into a rotation matrix and a translation vector to compare with the ground truth. We evaluate correspondence pruning in dynamic scenarios. The dynamic scenarios indicate there exist multiple consistencies between an image pair. We use a challenging dataset MULTI (Zhao et al. 2018) for evaluation, whose samples have at least two consistencies. This dataset provides coordinates of inliers, and the putative correspondences without ground truth. We generate labels for putative correspondences by searching in putative inliers with KDtree (Ramasubramanian and Paliwal 1992).

**Metrics.** We compute Area Under the Curve (AUC) to evaluate the pose accuracy at thresholds 5°, 10°, and 20°, which is also the metric used by (Sarlin et al. 2020; Liu et al. 2021) in camera pose estimation. We compute the precision, recall, and F1-scores with a symmetric epipolar distance threshold  $\sigma = 10^{-4}$  in camera pose estimation and with the computed labels in dynamic scenes.

### Camera Pose Estimation

**Baselines.** We consider LMF (Liu et al. 2021), CNe (Yi et al. 2018), ACNe (Sun et al. 2020), OANet (Zhang et al. 2019), and LMCNet (Liu et al. 2021) as baselines and include the results of classical correspondence pruning approaches (Barath et al. 2019; Ma et al. 2019; Bian et al. 2017; Lin et al. 2017). Learnable matchers, such as SuperGlue (Sarlin et al. 2020), act on an extended putative set compared to sparse correspondence filters. We thus do not include them as the baselines. For implementation, we follow official released pretrained models on YFCC100M with SIFT and our model is trained with the same setting. Considering the generalization performance, we evaluate all methods with the fixed models in different settings, *e.g.*, other different datasets and various descriptors.

**Results.** Quantitative results are shown in Table 1. We observe that, ANA-Net reports the second best performance (a comparable performance against the best performance) in YFCC100M and the best performance in SUN3D. Further,



Figure 6: Visualizations of camera pose estimation. Red/Green indicates incorrect/correct correspondences.

to explore the reason in performance differences, we calculate the F-scores with different inlier ratios. As shown in Fig. 5, ANA-Net outperforms baselines with +1.7% average improvement (0 ~ 20% inlier ratios and 80 ~ 100% outlier ratios) on YFCC100M and +8.5% average improvement on SUN3D. It indicates the better pruning behavior of ANA-Net in high-outlier-ratio cases. Fig. 6 shows the visual results with high outlier ratios. ANA-Net still correctly identifies inliers, while compared baselines identify a mass of outliers. Interestingly, as shown in the last column, ANA-Net can reject all correspondences when they are irregular.

We further evaluate the generalization ability in different settings and report results in Table 4. ANA-Net achieves leading performances in all settings, which indicates a clear advantage in generalization. Results on ScanNet(Dai et al. 2017) can be found in the supplementary. A similar tendency can be observed. We measure the inference time for 2048 correspondences over 4000 runs and report the average time in Table 2. We further take the time that generates putative correspondences (SIFT&NN) as reference time. Despite achieving superior performance, LMCNet takes more than 1.4 times as long as generating correspondences. Per to Table 1, Table 4, and Table 2, ANA-Net reports an approaching performance against LMCNet while running 14 times faster.

### The MULTI Dataset

Since MULTI (Zhao et al. 2018) only contains 45 image pairs, models trained in camera pose estimation are used here. In practice, we test them with the same form in camera pose estimation: predicting inlier probability from putative correspondences. We report results of precision, recall, and F1-scores in Table 3. In MULTI, ANA-Net exhibits the ability to capture diverse consistencies, implied by higher precision, recall, and F-scores. Contrary to the strict correspondence pruning in camera pose estimation, in this task ANA-Net is tolerant to consistent correspondences (even with different consistencies) and therefore could preserve more correspondences compared with other baselines.

### Ablation Study

Here we conduct ablation studies to understand how our method works and which part contributes to the most of performance improvement. We first compare different formulations discussed in “Second-Order Attention”. Results in Table 5(Left) suggest our approximations are effective and efficient. Our approximations reduce the complexity of modeling second-order context from  $\mathcal{O}(N^3)$  to  $\mathcal{O}(N)$  without

Form	AUC@5°	@10°	@20°	N=2048	4096	8192
Cub.	20.91	38.62	56.71	+5.35	+93.83	+914.15
Qua.	21.28	38.82	56.86	+4.72	+7.41	+26.85
Lin.	20.75	38.11	56.08	+4.07	+4.70	+8.89

Table 5: Left: Evaluation on different forms of second-order attention. We report the AUC on YFCC100M. Right: Times(ms). We measure extra(“+”) inference time(ms) introduced by second-order attention.

Method	SOC	LNN	Parm.(k)	+Parm.	AUC@5°	+ step	+ total
CNe	-	-	397	-	18.03	-	-
ACNe	-	-	410	+13	19.42	+1.39	+1.39
			524	+124	19.73	+0.31	
Ours	✓		662	+138	20.57	+0.84	+1.33
	✓	✓	≈662	+0.02	20.75	+0.18	

Table 6: Ablation studies on YFCC100M. ‘LNN’ indicates the learnable non-linear normalization. “+” is a comparison with the previous line.

much performance loss. We remark that the slightly better performance of the quadratic form over the cubic one is due to the optimization difficulty of the naive implementation. For the second-order attention, we measure with different N’s and results are shown in Table 5(Right). While the performance of the quadratic form seems better than the linear one, we recommend to use the linear form in practice, because the cost of its use almost comes for free.

With the linear form, we further decompose the ANA-Block and report results in Table 6. The baseline model only introduces the first-order attentive context, which can be viewed as an extension of ACNe, maintains better performance against ACNe. And all our components contribute to performance improvement. Specifically, the baseline brings a +0.31 improvement. The second-order attention (with vanilla sigmoid( $\cdot$ ) for non-linear normalization) brings a +0.8 improvement. It should be noted that its parameters increase is brought by the necessary encoder after the SOC layer as described in “attention-consistent context enhancement”, which dose not increase the capacity of the baseline. Besides, the proposed learnable non-linear normalization brings additional +0.2 improvement with only 20 extra parameters. Together, considering ACNe is the follow-up work of CNe, it improves CNe by 1.39. Instead our implemented first-order attention outperforms ACNe by +0.31 and the inclusion of the second-order attention improves ACNe by 1.33.

**How the ANA-Block works.** We collect  $k$ -nearest neighbors using KNN for inliers from each stacked blocks with different  $k$  and calculate the average inlier ratio. The number of blocks  $L = 5$ . Fig. 7 (Left) shows that, when more blocks are stacked, inliers tend to find additional inliers from their neighbors. The clustered inliers bodes well for subsequent inliers identification. It also embodies a inlier-clustering trend. Finally, we evaluate the performance with different  $L$ ’s and analyze the number of parameters. As shown in Fig. 7(Right), the repetitive use of ANA-Block

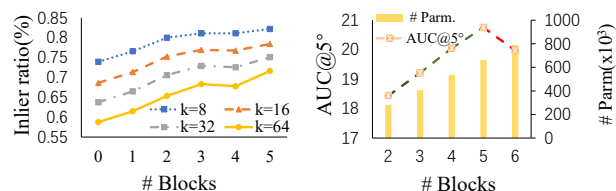


Figure 7: Left: Inlier ratios under different  $k$ ’s with changing # blocks. We collect intermediate features with different number of ANA blocks. The features are used to calculate inlier ratios in the neighbors of correct correspondences with different neighbor size  $k$ . Right: AUC and # parameters with changing # blocks.

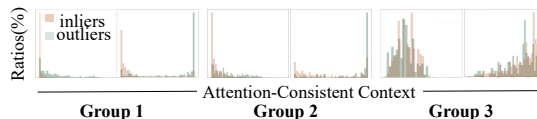


Figure 8: Conditions of applying attention-consistent context. Group 1: a clear bimodal distribution; Group 2: a weak bimodal distribution; Group 3: a confusing distribution.

yields an increase in AUC, while the AUC slightly drops when  $L = 6$ . Perhaps the distinction between inliers and outliers cannot be further enlarged with available information such that performance is saturated.

**How the second-order attention works.** Different from the conventional context considered in segmentation, attention-consistent context reflects a soft number of neighbors of attention weights, where inliers should own more neighbors, while outliers have less. This type of context can be visualized directly. As shown in Fig. 8, there are 3 typical situations when attention-consistent context is applied (each group visualizes an unnormalized and a normalized attention-consistent context): the group 1 shows an ideal situation with a bimodal distribution between inliers and outliers; the bimodal shape is slightly degraded in the group 2; and the group 3 shows a confusing case where the context may not be helpful. The confusing case can be caused by high outlier ratios. We remark that, the multi-head architecture proposed by transformer can somewhat alleviate this confusing case, because it is not likely that the confusing case appear in each head at the same time.

## Conclusion

We introduce the idea of attention in attention to model second-order attention. We present a naive formulation and two approximated formulations to capture attention-consistent context. We show that attention-consistent context can be complementary with feature-consistent context to address correspondence pruning problems effectively and efficiently. We believe the combination of the first- and second-order attention points out a new direction for the attention mechanism and shows potential to extend to additional tasks where attention mechanism matters.

## Acknowledgements

This work was supported in part by the National Natural Science Foundation of China (Grant No.U1913602) and by Huawei Technologies CO., LTD. (Grant No.YBN2020085023).

## References

- Bahdanau, D.; Cho, K.; and Bengio, Y. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Barath, D.; and Matas, J. 2018. Graph-cut RANSAC. In *Proc. IEEE Conf. Comput. Vis. Pattern Recogn.*, 6733–6741.
- Barath, D.; Noskova, J.; Ivashechkin, M.; and Matas, J. 2020. MAGSAC++, a fast, reliable and accurate robust estimator. In *Proc. IEEE Conf. Comput. Vis. Pattern Recogn.*, 1304–1312.
- Barath, D.; et al. 2019. MAGSAC: marginalizing sample consensus. In *Proc. IEEE Conf. Comput. Vis. Pattern Recogn.*, 10197–10205.
- Bian, J.; Lin, W.-Y.; Matsushita, Y.; Yeung, S.-K.; Nguyen, T.-D.; and Cheng, M.-M. 2017. Gms: Grid-based motion statistics for fast, ultra-robust feature correspondence. In *Proc. IEEE Conf. Comput. Vis. Pattern Recogn.*, 4181–4190.
- Brachmann, E.; Krull, A.; Nowozin, S.; Shotton, J.; Michel, F.; Gumhold, S.; and Rother, C. 2017. Dsac-differentiable ransac for camera localization. In *Proc. IEEE Conf. Comput. Vis. Pattern Recogn.*, 6684–6692.
- Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-end object detection with transformers. In *Proc. Eur. Conf. Comput. Vis.*, 213–229. Springer.
- Cavalli, L.; Larsson, V.; Oswald, M. R.; Sattler, T.; and Pollefeys, M. 2020. Adalam: Revisiting handcrafted outlier detection. *arXiv preprint arXiv:2006.04250*.
- Chen, H.-Y.; Lin, Y.-Y.; and Chen, B.-Y. 2015. Co-segmentation guided hough transform for robust feature matching. *IEEE Trans. Pattern Anal. Mach. Intell.*, 37(12): 2388–2401.
- Chum, O.; Werner, T.; and Matas, J. 2005. Two-view geometry estimation unaffected by a dominant plane. In *Proc. IEEE Conf. Comput. Vis. Pattern Recogn.*, volume 1, 772–779. IEEE.
- Dai, A.; Chang, A. X.; Savva, M.; Halber, M.; Funkhouser, T.; and Nießner, M. 2017. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proc. IEEE Conf. Comput. Vis. Pattern Recogn.*, 5828–5839.
- DeTone, D.; Malisiewicz, T.; and Rabinovich, A. 2018. Superpoint: Self-supervised interest point detection and description. In *Proc. IEEE Conf. Comput. Vis. Pattern Recogn. Workshop*, 224–236.
- Fischler, M. A.; and Bolles, R. C. 1981. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM*, 24(6): 381–395.
- Kim, Y.; Denton, C.; Hoang, L.; and Rush, A. M. 2017. Structured attention networks. *arXiv preprint arXiv:1702.00887*.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. ImageNet classification with deep convolutional neural networks. *Commun. ACM*, 60: 84–90.
- Lin, W.-Y.; Wang, F.; Cheng, M.-M.; Yeung, S.-K.; Torr, P. H.; Do, M. N.; and Lu, J. 2017. CODE: Coherence based decision boundaries for feature correspondence. *IEEE Trans. Pattern Anal. Mach. Intell.*, 40(1): 34–47.
- Liu, Y.; Liu, L.; Lin, C.; Dong, Z.; and Wang, W. 2021. Learnable Motion Coherence for Correspondence Pruning. In *Proc. IEEE Conf. Comput. Vis. Pattern Recogn.*, 3237–3246.
- Lowe, D. G. 2004. Distinctive image features from scale-invariant keypoints. *Int. J. Comput. Vis.*, 60(2): 91–110.
- Ma, J.; Zhao, J.; Jiang, J.; Zhou, H.; and Guo, X. 2019. Locality preserving matching. *Int. J. Comput. Vis.*, 127(5): 512–531.
- Mishchuk, A.; Mishkin, D.; Radenovic, F.; and Matas, J. 2017. Working hard to know your neighbor’s margins: Local descriptor learning loss. In *Proc. Adv. Neural Inf. Process. Syst.*, 4829–4840.
- Misra, I.; Girdhar, R.; and Joulin, A. 2021. An End-to-End Transformer Model for 3D Object Detection. In *Proc. IEEE Int. Conf. Comput. Vis.*, 2906–2917.
- Mur-Artal, R.; Montiel, J.; and Tardós, J. D. 2015. ORB-SLAM: A Versatile and Accurate Monocular SLAM System. *IEEE Trans. Robot.*, 31: 1147–1163.
- Ono, Y.; Trulls, E.; Fua, P.; and Yi, K. M. 2018. LF-Net: Learning Local Features from Images. In *Proc. Adv. Neural Inf. Process. Syst.*, 6237–6247.
- Pan, Z.; Cao, Z.; Wang, K.; Lu, H.; and Zhong, W. 2021. Transview: Inside, outside, and across the cropping view boundaries. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4218–4227.
- Qi, C. R.; Yi, L.; Su, H.; and Guibas, L. J. 2017. PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space. In *Proc. Adv. Neural Inf. Process. Syst.*, 5105–5114.
- Ramasubramanian, V.; and Paliwal, K. K. 1992. Fast k-dimensional tree algorithms for nearest neighbor search with application to vector quantization encoding. *IEEE Trans. Signal Process.*, 40(3): 518–531.
- Rublee, E.; Rabaud, V.; Konolige, K.; and Bradski, G. 2011. ORB: An efficient alternative to SIFT or SURF. In *Proc. IEEE Int. Conf. Comput. Vis.*, 2564–2571. Ieee. ISBN 1457711028.
- Sarlin, P.-E.; DeTone, D.; Malisiewicz, T.; and Rabinovich, A. 2020. Superglue: Learning feature matching with graph neural networks. In *Proc. IEEE Conf. Comput. Vis. Pattern Recogn.*, 4938–4947.
- Scarselli, F.; Gori, M.; Tsoi, A. C.; Hagenbuchner, M.; and Monfardini, G. 2008. The graph neural network model. *IEEE Trans. Neural Netw.*, 20(1): 61–80.

Schönberger, J. L.; and Frahm, J.-M. 2016. Structure-from-Motion Revisited. In *Proc. IEEE Conf. Comput. Vis. Pattern Recogn.*, 4104–4113.

Sun, W.; Jiang, W.; Trulls, E.; Tagliasacchi, A.; and Yi, K. M. 2020. ACNe: Attentive Context Normalization for Robust Permutation-Equivariant Learning. In *Proc. IEEE Conf. Comput. Vis. Pattern Recogn.*, 11283–11292.

Thomee, B.; Shamma, D.; Friedland, G.; Elizalde, B.; Ni, K. S.; Poland, D. N.; Borth, D.; and Li, L. 2016. YFCC100M: the new data in multimedia research. *Commun. ACM*, 59: 64–73.

Vaswani, A.; Shazeer, N. M.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is All you Need. In *Proc. Adv. Neural Inf. Process. Syst.*, 6000–6010.

Wang, K.; Du, S.; Liu, C.; and Cao, Z. 2022. Interior Attention-Aware Network for Infrared Small Target Detection. *IEEE Transactions on Geoscience and Remote Sensing*, 60: 1–13.

Wang, X.; Girshick, R.; Gupta, A.; and He, K. 2018. Non-local neural networks. In *Proc. IEEE Conf. Comput. Vis. Pattern Recogn.*, 7794–7803.

Xiao, J.; Owens, A.; and Torralba, A. 2013. SUN3D: A Database of Big Spaces Reconstructed Using SfM and Object Labels. *Proc. IEEE Int. Conf. Comput. Vis.*, 1625–1632.

Yi, K. M.; Trulls, E.; Ono, Y.; Lepetit, V.; Salzmann, M.; and Fua, P. 2018. Learning to Find Good Correspondences. In *Proc. IEEE Conf. Comput. Vis. Pattern Recogn.*, 2666–2674.

Ying, R.; You, J.; Morris, C.; Ren, X.; Hamilton, W. L.; and Leskovec, J. 2018. Hierarchical graph representation learning with differentiable pooling. *arXiv preprint arXiv:1806.08804*.

Zhang, J.; Sun, D.; Luo, Z.; Yao, A.; Zhou, L.; Shen, T.; Chen, Y.; Quan, L.; and Liao, H. 2019. Learning Two-View Correspondences and Geometry Using Order-Aware Network. In *Proc. IEEE Int. Conf. Comput. Vis.*, 5844–5853.

Zhao, C.; Cao, Z.; Li, C.; Li, X.; and Yang, J. 2019. NM-Net: Mining reliable neighbors for robust feature correspondences. In *Proc. IEEE Conf. Comput. Vis. Pattern Recogn.*, 215–224.

Zhao, C.; Ge, Y.; Zhu, F.; Zhao, R.; Li, H.; and Salzmann, M. 2021. Progressive Correspondence Pruning by Consensus Learning. In *Proc. IEEE Int. Conf. Comput. Vis.*, 6464–6473.

Zhao, C.; Yang, J.; Xiao, Y.; and Cao, Z. 2018. Scalable Multi-Consistency Feature Matching with Non-Cooperative Games. *Proc. IEEE Int. Conf. Image Process.*, 1258–1262.

Zhou, Y.; Barnes, C.; Shechtman, E.; and Amirghodsi, S. 2021. TransFill: Reference-guided Image Inpainting by Merging Multiple Color and Spatial Transformations. In *Proc. IEEE Conf. Comput. Vis. Pattern Recogn.*, 2266–2276.