

LocalTrans: A Multiscale Local Transformer Network for Cross-Resolution Homography Estimation

Ruizhi Shao^{1,4,*}, Gaochang Wu^{2,*}, Yuemei Zhou¹, Ying Fu³, Lu Fang¹, and Yebin Liu¹

¹Tsinghua University ²Northeastern University ³Beijing Institute of Technology ⁴Zhuohe Technology



Figure 1. We present a multiscale local transformer network, dubbed *LocalTrans*, for homography estimation. The proposed LocalTrans network achieves accurate homography estimation on challenging real-captured cross-resolution cases under resolution gap up to $10\times$. By warping each of the local high-resolution image on a global low-resolution image using the estimated homography matrix, we achieve artifact-free stitching on this challenging case, substantially outperforms feature-based homography estimation in [42].

Abstract

Cross-resolution image alignment is a key problem in multiscale gigapixel photography, which requires to estimate homography matrix using images with large resolution gap. Existing deep homography methods concatenate the input images or features, neglecting the explicit formulation of correspondences between them, which leads to degraded accuracy in cross-resolution challenges. In this paper, we consider the cross-resolution homography estimation as a multimodal problem, and propose a local transformer network embedded within a multiscale structure to explicitly learn correspondences between the multimodal inputs, namely, input images with different resolutions. The proposed local transformer adopts a local attention map specifically for each position in the feature. By combining the local transformer with the multiscale structure, the network is able to capture long-short range correspondences efficiently and accurately. Experiments on both the MS-COCO dataset and the real-captured cross-resolution dataset show that the proposed network outperforms existing state-of-the-art feature-based and deep-learning-based homography estimation methods, and is able to accurately align images under $10\times$ resolution gap.

* Equal contribution

1. Introduction

The rapidly development of multiscale gigapixel photography [5, 42, 45] brings large-scale, long-term and immersive visual experience. It synthesizes a single ultra-high-resolution image through aligning plenty of high-resolution local-views with a low-resolution global-view. In multiscale gigapixel photography, the large resolution gap between two views, namely cross-resolution, puts forward a new challenge to traditional homography estimation task. Homography estimation is defined as the estimation of the projection mapping between two views on the same plane in 3D space, which usually consists of three steps: feature extraction using SIFT [18] or SURF [2], correspondence matching, and homography matrix estimation based on the RANSAC [9] or a direct linear transform. It relies on dense features with the same resolution to achieve an accurate estimation, thus, usually fail in solving the cross-resolution problem.

Inspired by the success of deep learning, deep homography methods based on convolutional neural network are studied to deal with challenging scenes. The pioneer deep homography method proposed by DeTone *et al.* [8] implements the estimation of homography matrix with a typical VGG-net [30], which extracts correspondences from

the concatenated image pair. Based on this pioneer work, Le *et al.* [16] propose a multiscale strategy to progressively estimate the homography via network cascade. However, since the input views are concatenated and downsampled together, simply applying the multiscale strategy cannot solve the cross-resolution problem. A recent approach by Zhang *et al.* [44] proposed to extract features from input images separately with shared convolution layers. While the network directly concatenates the features in the following layer, which can be equivalent to concatenating the input images at the very beginning.

In this paper, we present a novel multiscale local transformer network, which we dubbed *LocalTrans*, to solve the cross-resolution problem in homography estimation. The transformer structure [33] has made a great success in learning the interaction between multimodal inputs [14, 26, 40] in the field of natural language processing and visual question answering. We therefore take a look at the cross-resolution problem through the lens of “multimodal”, and employ the transformer structure to explicitly capture correspondences through the correlation of the cross-resolution images in the feature space.

However, the vanilla transformer structure introduced in [33] brings high GPU memory and computational costs due to the outer product between high-dimensional matrices. To achieve a fast and accurate homography estimation, we introduce a local transformer and embed it within a multiscale structure. More specifically, we design a local convolution-based operation in the proposed local transformer, which applies a specific kernel to each position of the high-level feature to efficiently capture a local attention. Then the local transformer is deployed in each level of the multiscale structure, enabling the network to capture correspondences with a long-short range attention. The combination of the local transformer and the multiscale structure is significantly faster than the global attention mechanism in the vanilla transformer [33]. But most importantly, the proposed LocalTrans network shows a superiority to the vanilla transformer with the same backbone in the homography estimation task.

Benefiting from the combination of the local transformer layer and the multiscale structure, the proposed LocalTrans network outperforms the state-of-the-art homography estimation methods in terms of PSNR and corner error on the MS-COCO dataset [17]. Moreover, we demonstrate that the LocalTrans network highlights a superior performance on challenging real-captured cross-resolution cases under resolution gap up to 10 \times , and further apply it to multiscale gigapixel photography, see Fig. 1. The main contributions are summarized as

- We propose to solve the cross-resolution problem in homography estimation using the transformer structure by explicitly capturing the correspondences be-

tween the inputs.

- We design a novel local transformer layer embedded within multiscale structure, which is able to capture correspondences with a long-short range attention. Experiments demonstrate that the proposed structure outperforms the global attention mechanism.
- The proposed local transformer has significantly faster speed and lower GPU memory cost compared with the vanilla transformer structure, achieving real-time homography estimation at 60fps (please see Table 1).

2. Related Work

In this section, we review topics on homography estimation, cross-resolution image alignment and attention mechanism that are the most relevant to our work.

Feature-based Homography Estimation. Methods in this category utilize feature points extracted from the image pair to obtain a set of feature correspondences. Then a homography is estimated based on the direct linear transform or the robust fitting algorithms such as RANSAC [9]. The accuracy of homography estimation depends on the quality of the detected image features. Traditional feature detectors, such as SIFT [18], SURF [2], ORB [29], are able to detect reasonable keypoints, which are robust to lighting, blur and perspective distortion. Recently, several deep learning-based feature extraction methods are also developed, e.g., LFFNet [25] and ASLFeat [21], and reach a higher matching accuracy. Despite existing feature-based homography estimation methods can be robust to illuminance changes and foggy inputs, they often fail in cross-resolution cases.

Deep Homography Estimation. Deep learning-based homography estimation is first proposed by [8] using VGG-net [30] as backbone, which is more robust compared to traditional feature-based methods. To improve the generalization capacity on real data, Nguyen *et al.* [24] proposed an unsupervised learning method by minimizing a pixel-wise intensity error metric instead of the regression loss of homography matrix. To address the potential large motion problem, Le *et al.* [16] proposed to use a cascade strategy to estimate the motion mask and the homography matrix in a coarse-to-fine manner. The above methods implicitly estimate the correlation between the two views by concatenating the images along the channel dimension. Alternatively, Zhang *et al.* [44] proposed to use a feature extractor with shared weights to extract image features separately, and directly concatenate the features in the following homography estimator network. However, this architecture is equivalent to concatenating the input images at the beginning of the network, and thus, fail to address the cross-resolution problem, please refer to the comparison in Fig. 7 and Table.2.

Cross-Resolution Image Alignment. Cross-resolution image alignment is an open problem in cross-scale

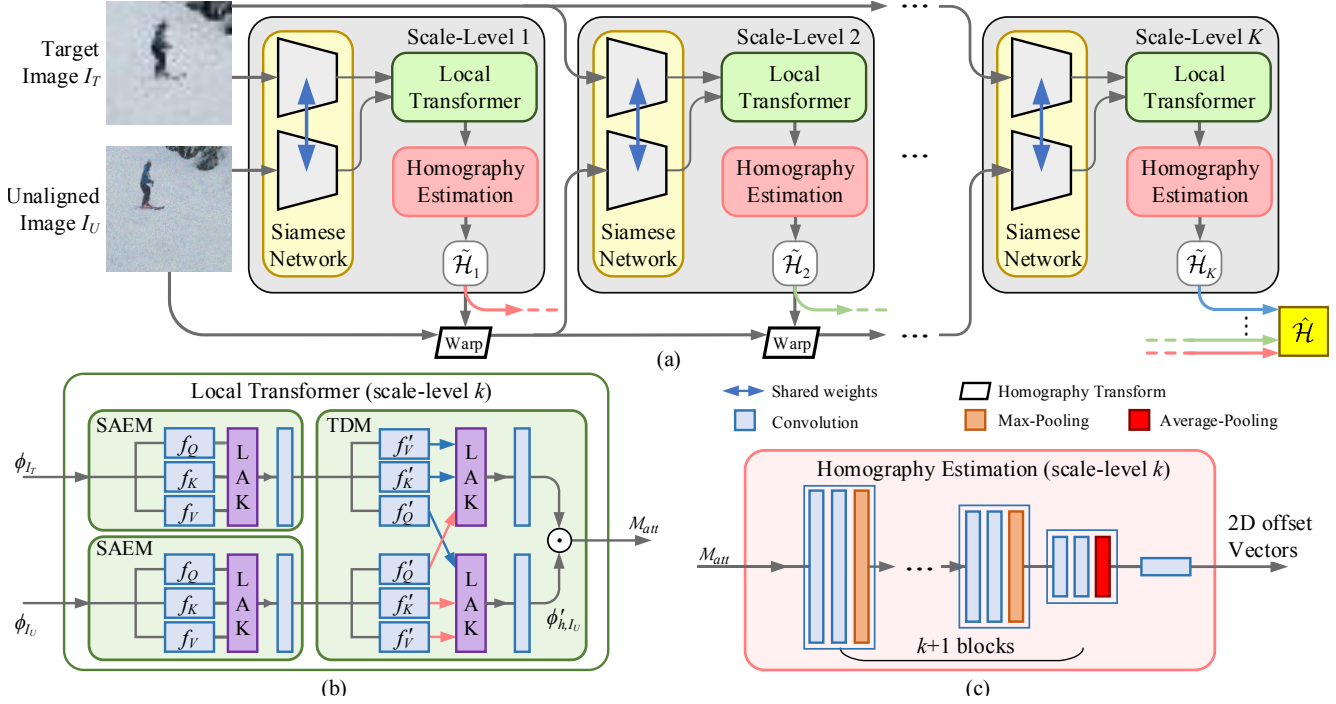


Figure 2. Architecture of the proposed LocalTrans network for homography estimation. (a) Overall structure of the LocalTrans; (b) Architecture of the local transformer that captures correspondences in different scales via a local self-attention encoder module (SAEM) and a local transformer decoder module (TDM); (c) Architecture of the homography estimation module that adopts local attention maps and high-level feature as input to estimate homography matrices from coarse-to-fine.

stereo [50], hybrid light field imaging [4, 48], multiscale gigapixel videography [42], etc. Commonly a pixel-to-pixel warping field is estimated and applied for registration between images with different resolutions. For example, Zhao *et al.* [47] presented a disparity estimation and refinement method for reconstructing high resolution light field in a hybrid light field imaging system. Zheng *et al.* [49] further proposed a deep learning-based optical flow estimation and image fusion method in a coarse-to-fine manner. To synthesize a multiscale gigapixel video, Yuan *et al.* [42] proposed an iterative feature matching and warping method to perform global and mesh-based homography estimations. In this paper, we investigate the cross-resolution problem in the homography estimation task by using transformer structure to pay attention to the correspondences across different resolution inputs.

Attention Mechanism. Attention was built to imitate the mechanism of human perception that mainly focuses on the salient part [13, 28, 7]. Vaswani *et al.* [33] indicated that the global attention mechanism is able to solve the long term dependency problem even without the backbone of a convolutional or a recurrent network. Wang *et al.* [36] introduced a self-attention to capture long-range dependencies (i.e., correspondences) by using matrix multiplication between reshaped feature maps. Alternative to the above global attention, many researches also investigated local attention approaches that focus on short-range dependencies

[10, 22, 31, 38, 34]. For example, Sperber *et al.* [31] introduces a soft Gaussian bias and a hard mask that is non-zero in a local region to control the context range attended by the network. Woo *et al.* [37] introduced a spatial attention module by using a convolution layer to extract inter-spatial relationship in a feature map.

Recently, numerous literatures show that the attention mechanism is efficient to obtain correlation across multi-modal inputs, e.g., for visual question answering [3, 15, 41], video description [11] and texture transferring [39]. We therefore apply the attention mechanism to explicitly capture the correspondences across the input images, especially cross-resolution images, for homography estimation.

3. Methodology

3.1. Overview

In this paper, we introduce a novel deep homography estimation network with a multiscale local transformer structure, dubbed *LocalTrans*. The proposed LocalTrans network first applies a deep siamese network (i.e., an image encoder with shared weights) to extract features ϕ_{I_T} and ϕ_{I_U} from the target image I_T and the unaligned image I_U , respectively. Two convolution layers followed by a 2×2 max-pooling layer constitute a basic block in the deep siamese network. Therefore, we construct features with different scales in the multiscale structure by controlling the




Figure 3. Visualization of the attention map. Top: self-attention map $\mathcal{M}_{s,I_U}^{(k)}$; Bottom: cross-attention map $\mathcal{M}_{I_U \rightarrow I_T}^{(k)}$.

block number of the deep siamese network in each scale-level. In concrete details, we use $K - k + 1$ blocks to construct feature maps $\phi_I^{(k)}$ of shape $H_k \times W_k$ in scale level k ($k = [1, 2, \dots, K]$), where $H_k = \frac{H}{2^{K-k+1}}$, $W_k = \frac{W}{2^{K-k+1}}$, and H and W are the height and width of the input images.

Different from existing deep homography methods that simply concatenate the images or features, we then explicitly formulate correspondences between the features ϕ_{I_T} and ϕ_{I_U} by using the transformer structure (Sec. 3.2). In each scale-level, a homography estimation module (Sec. 3.3) is adopted to estimate a homography matrix \mathcal{H}_k based on the attention map and feature maps. Then the unaligned image I_U is warped according to the homography matrix \mathcal{H}_k and then fed into the next scale-level (please refer to Fig. 2 (a)).

3.2. Multiscale Local Transformer Network

In this section, we introduce the proposed local transformer incorporating a multiscale structure. A straightforward option to achieve a long-short range perception of correspondence between features is to implement the transformer structure [33] after the siamese network. However, the vanilla transformer structure brings high GPU memory and computational costs when processing high-dimensional features. To accelerate the transformer, we replace the global attention in the vanilla transformer structure by designing a novel local attention kernel (Sec. 3.2.2) and combining it with a multiscale structure. Despite the designed local transformer kernel only captures correspondences in a limited range in the low scale-level, the multiscale structure enables the network to perceive the correspondence in a long-short range manner.

3.2.1 Transformer Structure

The detailed architecture of the transformer structure in each scale-level is shown in Fig. 2(b). The inputs of the transformer are two features $\phi_{I_T}^{(k)}$ and $\phi_{I_U}^{(k)}$ output by the deep siamese network. Two modules, Self-Attention Encoder Module (SAEM) and Transformer Decoder Module (TDM), are employed to exploit internal relations within

the feature maps via self-attention and to capture the correspondences across the two features from the multimodal inputs via the cross-attention, respectively.

Self-Attention Encoder Module (SAEM). The SAEM first applies three 1×1 convolution layers $f_Q(\phi_I)$, $f_K(\phi_I)$ and $f_V(\phi_I)$ without activation function to encode the input image feature ϕ_I ($\phi_{I_T}^{(k)}$ or $\phi_{I_U}^{(k)}$) into features $\phi_{Q,I}$, $\phi_{K,I}$, $\phi_{V,I}$ of shape $C \times H_k \times W_k$, where C denotes the channel number. Then the self-attention result in the SAEM is computed as follows

$$\mathcal{M}_s = \sigma\left(\frac{\phi_{Q,I} \odot \phi_{K,I}}{\sqrt{C}}\right), \quad (1)$$

$$\phi_{h,I} = \mathcal{M}_s \otimes \phi_{V,I},$$

where σ denotes the softmax function, and \odot and \otimes denote the operations in the designed local transformer structure, which will be described in Sec. 3.2.2. The tensor \mathcal{M}_s is usually interpreted as a self-attention map.

Since $\phi_{Q,I}$, $\phi_{K,I}$ and $\phi_{V,I}$ are derived from the same input ϕ_I , the self-attention mechanism encourages to enhance the edges and corners in the input feature $\phi(I)$. As shown in an example in Fig. 3 (top), the network pays more attention to the edge with the same feature of the center pixel, which is more prominent in attention maps of higher scale-levels (e.g., $\mathcal{M}_s^{(2)}$ and $\mathcal{M}_s^{(3)}$). The final high-level feature $\phi_{s,I}$ output by the SAEM is generated by encoding the self-attention result $\phi_{h,I}$ with a 1×1 convolution layer.

Transformer Decoder Module (TDM). In this module, two iterations of cross-attention are used. In the first iteration, we first adopt three 1×1 convolution layers $f'_Q(\cdot)$, $f'_K(\cdot)$, $f'_V(\cdot)$ without activation function to encode the high-level features $\phi_{s,I}$ into features $\phi'_{Q,I}$, $\phi'_{K,I}$ and $\phi'_{V,I}$ as those in the SAEM. But different with the self-attention in the SAEM (Eqn. 1), we apply a cross-attention mechanism between features from target image I_T and unaligned image I_U , denoted as

$$\mathcal{M}_{I_U \rightarrow I_T} = \sigma\left(\frac{\phi'_{Q,I_U} \odot \phi'_{K,I_T}}{\sqrt{C}}\right),$$

$$\phi'_{h,I_T} = \mathcal{M}_{I_U \rightarrow I_T} \otimes \phi'_{V,I_T}, \quad (2)$$

$$\mathcal{M}_{I_T \rightarrow I_U} = \sigma\left(\frac{\phi'_{Q,I_T} \odot \phi'_{K,I_U}}{\sqrt{C}}\right),$$

$$\phi'_{h,I_U} = \mathcal{M}_{I_T \rightarrow I_U} \otimes \phi'_{V,I_U},$$

where $\mathcal{M}_{I_U \rightarrow I_T}$ and $\mathcal{M}_{I_T \rightarrow I_U}$ are usually interpreted as cross-attention map, and ϕ'_{h,I_T} and ϕ'_{h,I_U} are attention-aware features.

In the second iteration, features ϕ'_{h,I_T} and ϕ'_{h,I_U} are first encoded into ϕ'_{s,I_T} and ϕ'_{s,I_U} using two 1×1 convolution layers, as shown in Fig. 2(b). Then we compute the attention map \mathcal{M}_{att} using ϕ'_{s,I_T} and ϕ'_{s,I_U} as inputs

$$\mathcal{M}_{att} = \phi'_{s,I_T} \odot \phi'_{s,I_U}. \quad (3)$$

The attention map \mathcal{M}_{att} is served to estimate the homography matrix, which will be introduced in Sec. 3.3.

The attention mechanism described in Eqn. 2 enables an interaction between target image I_T and unaligned image I_U in the feature space, and therefore, captures correspondence information more explicitly than simply feeding the image pair into the network through channel-wise concatenation [30, 8, 24, 16]. Besides, the network with the transformer structure is more robust when the input two images have considerably large resolution differences. Fig. 3 (bottom) visualizes the cross-attention map $\mathcal{M}_{I_U \rightarrow I_T}$ for input images under $4\times$ resolution gap. In this cross-resolution case, more attention values (weights) are assigned to the same edge as that in the self-attention map \mathcal{M}_{s, I_U} .

3.2.2 Local Attention Kernel

To accelerate the transformer structure in Sec. 3.2, we propose a Local Attention Kernel (LAK) that captures correspondences in a local range, which is inspired by conventional 2D deconvolution (also known as transposed convolution) and convolution. The main difference is that the proposed local transformer applies variable slice in the attention map as the convolution kernel while the traditional 2D convolution or deconvolution adopts a fixed kernel in the feedforward path. In the following, we will introduce the proposed LAK by decomposing the attention mechanism in Eqn. 1 and 2 into two steps, as shown in Fig. 4.

Local attention map generation. Local attention map describes the correspondences between features ϕ_Q and ϕ_K in a local range, i.e., a squared window. Consider $\phi_Q \in \mathbb{R}^{C \times H_k \times W_k}$, $\phi_K \in \mathbb{R}^{C \times H_k \times W_k}$ and the radius of LAK is r . Then for a certain element at position $\mathbf{x} = (x, y)$ in ϕ_Q will first query the relationship with elements in a local range of $\mathcal{N}(\mathbf{x})$ with radius r in ϕ_K . Suppose the position of an element in ϕ_K is $\mathbf{u} \in \mathcal{N}(\mathbf{x})$, then a local correspondence map \mathcal{M}' could be described as

$$\mathcal{M}'(\mathbf{x}, \mathbf{u}) = \phi_Q^T(\mathbf{x})\phi_K(\mathbf{u}),$$

where $\phi_Q^T(\mathbf{x}) \in \mathbb{R}^{1 \times C}$ and $\phi_K(\mathbf{u}) \in \mathbb{R}^{C \times 1}$. The above formulation also explains the operation $\phi_Q \odot \phi_K$ in Eqn. 1, Eqn. 2 and Eq. 3. And the final local attention map \mathcal{M} is

$$\mathcal{M} = \sigma\left(\frac{\mathcal{M}'}{\sqrt{C}}\right).$$

The above equation shows that the local attention map \mathcal{M} is a 4D tensor with shape $H_k \times W_k \times (2r+1) \times (2r+1)$, which records the correspondence of each position in ϕ_Q with that in ϕ_K in a local range. For instance, an element $\mathcal{M}(x, y, u, v)$ records the correspondence between point $\phi_Q(x, y)$ and $\phi_K(x+u, y+v)$, $u, v \in [-r, r]$.

Local attention convolution. This operation uses the 4D local attention map \mathcal{M} and the feature ϕ_V to obtain




Figure 4. The process of Local Attention Kernel (LAK). In the first step, a 2D slice of the local attention map is generated for each element of feature ϕ_Q in position \mathbf{x} . In the next step, we regard the 2D slice of the local attention map as a convolution kernel for a patch centered at position \mathbf{x} in feature ϕ_V .

high-level feature ϕ_h . Consider a 2D slice of \mathcal{M} at a certain position $\mathbf{x} = (x, y)$, denoted as $\omega_{\mathbf{x}}$. The feature ϕ_h is then obtained by performing the convolution between the feature ϕ_V and the 2D slice $\omega_{\mathbf{x}}$

$$\phi_h(\mathbf{x}) = \omega_{\mathbf{x}} * \phi_V(\mathbf{x}) = \sum_{\mathbf{u} \in \mathcal{N}(\mathbf{x})} \mathcal{M}(\mathbf{x}, \mathbf{u})\phi_V(\mathbf{x} + \mathbf{u}),$$

where $\mathcal{M}(\mathbf{x}, \mathbf{u}) \in \mathbb{R}^{1 \times 1}$ and $\phi_V(\mathbf{x} + \mathbf{u}) \in \mathbb{R}^{C \times 1}$. The above formulation also explains the operation $(\cdot) \otimes \phi_V$ in Eqn. 1 and Eqn. 2.

Discussions. It should be noted that there is a clear difference between the proposed LAK and existing approach with local attention [10, 31, 38, 34, 37, 35]. Existing approaches typically implement local attention by setting an attention bias (a Gaussian bias [38] or a hard mask that is non-zero in a local region [31]) or convolution layers to perform channel squeezing [34, 37, 35]. In our proposed LocalTrans, the convolution kernel in the LAK, i.e., the 2D slice $\omega_{\mathbf{x}}$, varies with the position \mathbf{x} , while the kernel in a conventional convolution is fixed in every position. We therefore term the operation as local attention convolution. The most related local attention structure was proposed in [22], which computes a local weights within a small window and produce feature vector through weighted average. We generalize this concept to 2D space in combination with 2D convolution, making it more suitable for visual tasks.

Comparing with the global attention in the vanilla transformer structure, the proposed LAK effectively reduces the computational complexity from $\mathcal{O}(H_k^2 \cdot W_k^2 \cdot C)$ to $\mathcal{O}(H_k \cdot W_k \cdot (2r+1)^2 \cdot C)$, and the memory usage of the attention map from $H_k^2 \cdot W_k^2$ to $H_k \cdot W_k \cdot (2r+1)^2$. For a certain scale-level $k \in K$, we set the radius of the LAK $r = k + 1$ to encourage the local transformer to notice a longer range of correspondence in the higher levels.

To demonstrate the effectiveness of the proposed local attention, we compare the local transformer against the vanilla transformer with the same deep siamese network and the homography estimation module, as shown in Fig. 8. The result shows that the proposed LAK with even 1 scale is superior to the global attention in the vanilla trans-

former structure. In conclusion, the proposed local transformer kernel not only shows high computational efficiency but also has superior performance to the vanilla transformer structure with global perception of correspondences.

3.3. Homography Estimation Module

In each scale-level, the homography estimation module applies the attention map \mathcal{M}_{att} in Eqn. 3 as input. A 8-dimensional vector is obtained by several convolution-pooling blocks to estimate the final homography matrix, as shown in Fig. 2(c).

More specifically, in a certain scale-level k , since the attention map \mathcal{M}_{att} is a 4D tensor, we first reshape it from $[\frac{H}{2^{K-k+1}}, \frac{W}{2^{K-k+1}}, 2k+3, 2k+3]$ to $[(2k+3)^2, \frac{H}{2^{K-k+1}}, \frac{W}{2^{K-k+1}}]$ (similar to the processing of cost volume in [32]) before feeding to the homography estimation module. Then we feed the feature into $k+1$ convolution blocks, where each block contains two 3×3 convolution layers and one max-pooling layer. In the last block, the max-pooling layer is replaced by an average-pooling layer for generating a feature of shape $C \times 1 \times 1$. After a 1D convolution layer (kernel size 1), the final output becomes a 8-dimensional vector, denoting the 2D offsets of 4 corner points $\hat{\mathbf{c}}_1^{(k)}, \dots, \hat{\mathbf{c}}_4^{(k)}$ in scale-level k . We can simply obtain the homography matrix $\tilde{\mathcal{H}}_k$ from the offset vectors and warp the the unaligned image $I_U^{(k)}$ using homography transform to obtain the input for the next level, i.e.,

$$I_U^{(k+1)} = \text{Warp}(I_U^{(k)}, \tilde{\mathcal{H}}_k).$$

Note that $\tilde{\mathcal{H}}$ in each scale-level represents the homography matrix of the full resolution. Thus, the final homography matrix $\tilde{\mathcal{H}}$ is computed directly by accumulating the estimated homography matrix $\tilde{\mathcal{H}}_k$ in each level as follows

$$\tilde{\mathcal{H}} = \tilde{\mathcal{H}}_1 \times \tilde{\mathcal{H}}_2 \times \dots \times \tilde{\mathcal{H}}_K.$$

3.4. Implementation Details

In our experiment, we set the number of scale-level $K = 3$. Except the 1×1 convolution layers $f_Q, f_K, f_V, f'_Q, f'_K$ and f'_V in the local transformer and the 1D convolution layer at the end of the homography estimation module, every other 2D convolution layer has a 3×3 kernel followed by a batch normalization layer [12] and a ReLU activation. More details of the network specification are listed in the supplementary file. For the local transformer layer, we implement the LAK including local attention map generation and local attention convolution in CUDA. To make them differentiable, we also deduce a backward propagation and package it to PyTorch autograd function [27].

For the training objective, we use the L1 norm of the corner error as the loss function $\mathcal{L} = \frac{1}{4} \sum_{i=1}^4 \|\mathbf{c} - \hat{\mathbf{c}}_i\|_1$, where \mathbf{c}_i and $\hat{\mathbf{c}}_i$ are corner point i transformed by the ground-truth homography and the estimated homography, respectively.

We only use the MS-COCO dataset [17] for the network training. And we follow the same data processing schemes in [8, 6] to generate image pairs. Moreover, we also add Gaussian noise and randomly adjust brightness, saturation and contrast to increase the robustness of the network.

4. Experiments

We compare the proposed LocalTrans network with both feature-based and deep learning-based homography estimation methods. We evaluate the proposed LocalTrans network in two different settings, common data (Sec. 4.1) in the MS-COCO dataset [17] as that in most deep homography estimation methods, and cross-resolution setting where the target image has lower resolution.

We have two kinds of datasets for the cross-resolution setting. The first is synthesized cross-resolution data (Sec. 4.2), in which the target images are downsampled using bicubic interpolation with factors $4\times$ and $8\times$. The second is optical zoom-in cross-resolution data (Sec. 4.3), which we apply multiscale gigapixel dataset from Yuan *et al.* [42] and cross-resolution stereo dataset from Zhou *et al.* [50]. In the cross-resolution setting, the proposed LocalTrans network as well as the baseline networks are re-trained.

4.1. Data in Common Setting

We compare our model on the MS-COCO dataset [17] in common setting with the following baseline methods, AffNet [23], LFNNet [25], DHN[8], UDHN by Zhang *et al.* [44], MHN [16], PFNet [43], PWC [32], SIFT+ContextDesc+RANSAC [19], SIFT+GeoDesc+RANSAC [20], SIFT+MAGSAC [1], SIFT+RANSAC [18]. Fig. 5 shows that the proposed LocalTrans outperforms feature-based homography estimation methods [18, 32, 43, 25, 19, 20, 1] and state-of-the-art deep learning-based methods [8, 23, 44, 16] in common setting. We also compare the proposed LocalTran network with a similar multiscale structure-based deep homography network MHN [16] with different numbers of scales on the MS-COCO dataset, as shown in Fig. 6. The result shows that the proposed network with the local transformer outperforms the MHN with different scale-levels. Moreover, our network on with 2 scales performs even better than the MHN with 3 scales. The experiment empirically validates that the proposed local transformer is able to capture correspondences more accurately than simply stacking the images [16] or feature maps [44] as input.

Ablation study. To verify the efficiency of local transformer structure, we replace the proposed LAK with the vanilla (global) transformer in [33] while keeping the rest architecture of the network unchanged. This experiment is performed on an Intel(R) Xeon CPU E5-2699 V4 with 16GB memory and a NVIDIA RTX 2080 GPU. The comparison in Table 1 shows that the proposed LAK has higher computational efficiency in terms of both running time and




Figure 5. Evaluations in common setting on the MS-COCO.




Figure 6. Comparison with MHN [16] on the MS-COCO dataset with different numbers of scales.

	Global Transformer		Local Transformer	
	Memory	Speed	Memory	Speed
128 × 128 (1 scale)	52.0M	152fps	48.9M	249fps
128 × 128 (2 scales)	75.0M	104fps	50.7M	203fps
128 × 128 (3 scales)	476.9M	16.4fps	67.7M	132fps
256 × 256 (1 scale)	208.4M	129fps	196.9M	213fps
256 × 256 (2 scales)	313.9M	52.9fps	204.8M	173fps
256 × 256 (3 scales)	2434M	4.31fps	287.0M	87.7fps

Table 1. Per-image memory consuming and speed comparison between global and local transformer using different sizes of inputs.

GPU memory cost than the vanilla transformer. Moreover, results in Fig. 8 also validate the superior performance compared with a single scale global transformer network. Please refer to the supplementary for more ablation studies.

4.2. Synthesized Cross-Resolution Data

We compare our model on synthesized cross-resolution dataset with 6 baseline methods, a conventional feature-based method, SIFT+RANSAC [18], three deep learning-based methods, DHN [8], UDHN [44] and MHN [16], as well as two state-of-the-art Reference-based Super-Resolution (RefSR) methods, SRNTT [46] and TTSR [39].

Fig. 7 shows the quantitative comparison on the MS-COCO dataset under 4× and 8× resolution gaps. The proposed LocalTrans network demonstrates superior performance on cross-resolution cases comparing with the con-




Figure 7. Evaluation on the MS-COCO dataset under 4× and 8× resolution gaps.




Figure 8. Comparison between the proposed LAK in different scale-levels and the vanilla (global) transformer structure on the MS-COCO dataset [17].

Method	4× ↓		8× ↓	
	PSNR	SSIM	PSNR	SSIM
SIFT+RANSAC [18]	19.32	0.802	11.93	0.64
DHN [8]	19.64	0.818	17.04	0.762
UDHN [44]	21.44	0.864	18.78	0.834
MHN [16]	25.42	0.951	20.24	0.860
SRNTT [46]	27.06	0.901	-	-
TTSR [39]	27.89	0.915	-	-
LocalTrans	30.17	0.981	24.12	0.930

Table 2. Numerical comparison (PSNR/SSIM) in different cross-scale settings on the MS-COCO datasets.

ventional feature-based method, SIFT+RANSAC [18], and two deep learning-based methods, DHN [8] and MHN [16]. The numerical comparison in Table 2 shows that the proposed LocalTrans network significantly outperforms the deep learning-based methods, DHN [8] and MHN [16], and the RefSR methods, SRNTT [46] and TTSR [39], which validates the superiority of the proposed local transformer structure in solving the cross-resolution problem.

4.3. Optical Zoom-in Cross-Resolution Data

In this experiment, 4 baseline methods, SIFT+RANSAC [18], DHN [8], UDHN by Zhang *et al.* [44] and MHN [16]

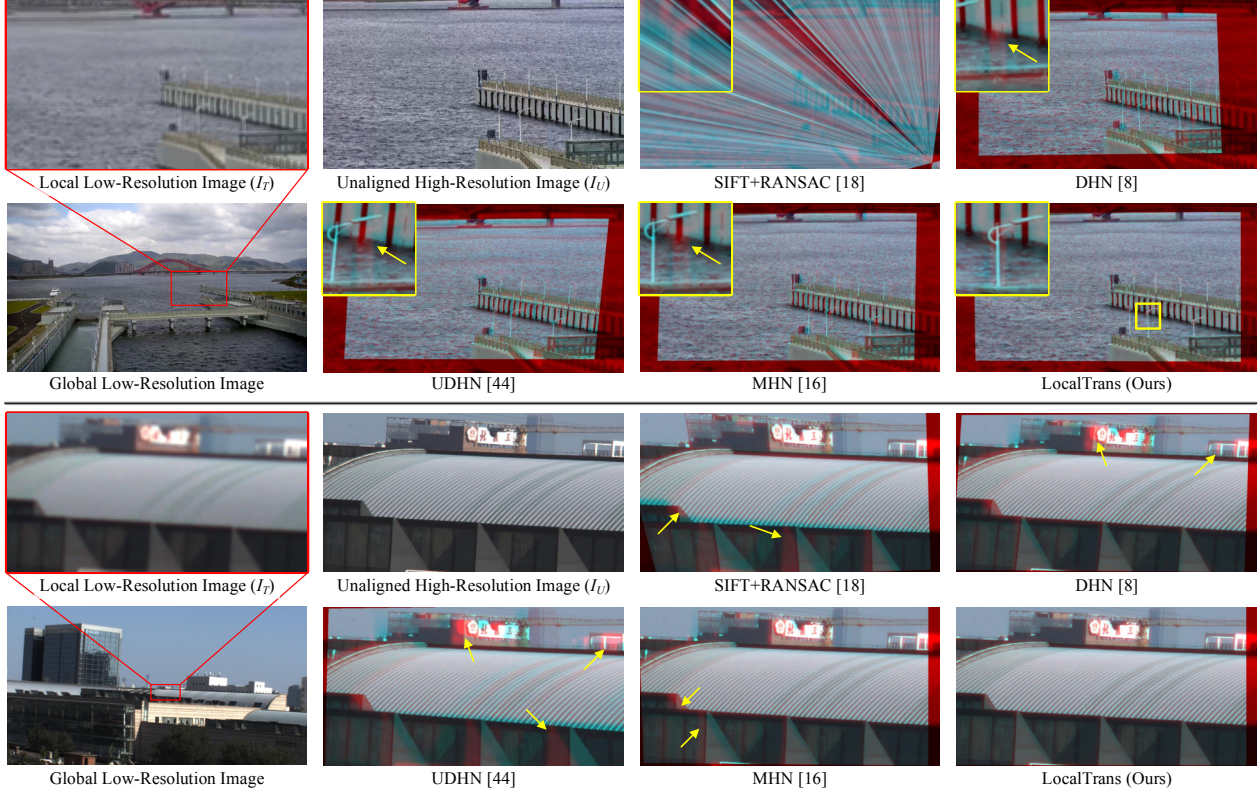


Figure 9. Visual evaluation on the multiscale gigapixel dataset [42] (top, $6\times$) and the cross-resolution stereo dataset [50] (bottom, $10\times$). We mix the GB channels of the aligned image and the R channel of the target image. The misaligned pixels appear as red or green ghosts.

are compared. On the multiscale gigapixel dataset [42], we then apply the local patches in a 5×5 grid of the global low-resolution image to estimate homography matrices separately. To ensure the spatial smoothness between neighbouring patches, we calculate the four corner points of each patch and take an average among its neighbors. The final result is obtained by warping each local high-resolution image to the corresponding grid as in [42].

Since there is no groundtruth for qualitative evaluation, we only demonstrate the visual comparison on the datasets [42, 50], as shown in Fig. 9. The resolution gaps between the local target images and the unaligned images are $6\times$ in the multiscale gigapixel dataset [42] (top of Fig. 9) and $10\times$ in the cross-resolution stereo dataset [50] (bottom of Fig. 9). The results show that the conventional SIFT+RANSAC [18] fails to estimate reasonable homography matrix in the first case, and the deep learning-based methods DHN [8], UDHN [44] and MHN [16] appear different degrees of missing alignments (please zoom-in for details). The proposed LocalTrans demonstrates the best visual results on the optical zoom-in cross-resolution datasets [42, 50], which has more complicated degradation model than synthesized cross-resolution data. More visual results on the optical zoom-in cross-resolution dataset are provided in the supplementary.

5. Conclusions

In this paper, we proposed a novel multiscale local transformer network, termed LocalTrans, for addressing the cross-resolution problem in homography estimation. We consider the cross-resolution images as some kind of multimodal input, and employ the transformer structure to explicitly capture correspondences between two modalities (cross-resolution images) in the feature space. To accelerate the transformer, we design a Local Attention Kernel (LAK) that generates a local attention map specifically for each position in the feature. By embedding the LAK within a multiscale structure, the proposed LocalTrans is able to capture correspondences in a long-short range manner. The proposed LocalTrans network outperforms state-of-the-art methods on the MS-COCO dataset and highlights a superior performance on the challenging real-captured cross-resolution dataset under resolution gap up to $10\times$.

We believe our LocalTrans gives a new opportunity to learn robust and accurate interactions between cross-resolution inputs, and would be further applied to various applications, such as reference-based super-resolution, cross-resolution stereo matching and hybrid light field imaging.

Acknowledgements This work is supported by the National Key Research and Development Program of China No.2018YFB2100500 and the NSFC No.62125107.

References

- [1] Daniel Barath, Jiri Matas, and Jana Noskova. Magsac: Marginalizing sample consensus. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10197–10205, 2019. 6
- [2] Herbert Bay, Tinne Tuytelaars, and Luc Van Gool. Surf: Speeded up robust features. In *European conference on computer vision*, pages 404–417. Springer, 2006. 1, 2
- [3] Hedi Ben-younes, Remi Cadene, Matthieu Cord, and Nicolas Thome. Mutan: Multimodal tucker fusion for visual question answering. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 2631–2639, 2017. 3
- [4] Vivek Boominathan, Kaushik Mitra, and Ashok Veeraraghavan. Improving resolution and depth-of-field of light field cameras using a hybrid imaging system. In *IEEE International Conference on Computational Photography*, 2014. 3
- [5] David J. Brady, Michael E. Gehm, Ronald A. Stack, Daniel L. Marks, David S. Kittle, Dathon R. Golish, Esteban Vera, and Steven D. Feller. Multiscale gigapixel photography. *Nature*, 486(7403):386–389, 2012. 1
- [6] Che-Han Chang, Chun-Nan Chou, and Edward Y Chang. Clkn: Cascaded lucas-kanade networks for image alignment. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2213–2221, 2017. 6
- [7] Maurizio Corbetta and Gordon L Shulman. Control of goal-directed and stimulus-driven attention in the brain. *Nature reviews neuroscience*, 3(3):201–215, 2002. 3
- [8] Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Deep image homography estimation. *arXiv preprint arXiv:1606.03798*, 2016. 1, 2, 5, 6, 7, 8
- [9] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981. 1, 2
- [10] Anmol Gulati, James Qin, Chung-Cheng Chiu, Niki Parmar, Yu Zhang, Jiahui Yu, Wei Han, Shibo Wang, Zhengdong Zhang, Yonghui Wu, and Ruoming Pang. Conformer: Convolution-augmented transformer for speech recognition. In *Interspeech 2020*, pages 5036–5040, 2020. 3, 5
- [11] Chiori Hori, Takaaki Hori, Teng-Yok Lee, Ziming Zhang, Bret Harsham, John R. Hershey, Tim K. Marks, and Kazuhiko Sumi. Attention-based multimodal fusion for video description. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 4203–4212, 2017. 3
- [12] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015. 6
- [13] Laurent Itti, Christof Koch, and Ernst Niebur. A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on pattern analysis and machine intelligence*, 20(11):1254–1259, 1998. 3
- [14] Yash Kant, Dhruv Batra, Peter Anderson, Alexander G. Schwing, Devi Parikh, Jiasen Lu, and Harsh Agrawal. Spatially aware multimodal transformers for textvqa. In *ECCV (9)*, pages 715–732, 2020. 2
- [15] Jin-Hwa Kim, Jaehyun Jun, and Byoung-Tak Zhang. Bilinear attention networks. In *Advances in Neural Information Processing Systems*, volume 31, pages 1564–1574, 2018. 3
- [16] Hoang Le, Feng Liu, Shu Zhang, and Aseem Agarwala. Deep homography estimation for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7652–7661, 2020. 2, 5, 6, 7, 8
- [17] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 2, 6, 7
- [18] David G. Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2):91–110, 2004. 1, 2, 6, 7, 8
- [19] Zixin Luo, Tianwei Shen, Lei Zhou, Jiahui Zhang, Yao Yao, Shiwei Li, Tian Fang, and Long Quan. Contextdesc: Local descriptor augmentation with cross-modality context. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2527–2536, 2019. 6
- [20] Zixin Luo, Tianwei Shen, Lei Zhou, Siyu Zhu, Runze Zhang, Yao Yao, Tian Fang, and Long Quan. Geodesc: Learning local descriptors by integrating geometry constraints. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 168–183, 2018. 6
- [21] Zixin Luo, Lei Zhou, Xuyang Bai, Hongkai Chen, Jiahui Zhang, Yao Yao, Shiwei Li, Tian Fang, and Long Quan. Aslfeat: Learning local features of accurate shape and localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6589–6598, 2020. 2
- [22] Minh-Thang Luong, Hieu Pham, and Christopher D. Manning. Effective approaches to attention-based neural machine translation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1412–1421, 2015. 3, 5
- [23] Dmytro Mishkin, Filip Radenovic, and Jiri Matas. Repeatability is not enough: Learning affine regions via discriminability. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 284–300, 2018. 6
- [24] Ty Nguyen, Steven W Chen, Shreyas S Shivakumar, Camillo Jose Taylor, and Vijay Kumar. Unsupervised deep homography: A fast and robust homography estimation model. *IEEE Robotics and Automation Letters*, 3(3):2346–2353, 2018. 2, 5
- [25] Yuki Ono, Eduard Trulls, Pascal Fua, and Kwang Moo Yi. Lf-net: learning local features from images. In *Advances in neural information processing systems*, pages 6234–6244, 2018. 2, 6
- [26] Georgios Paraskevopoulos, Srinivas Parthasarathy, Aparna Khare, and Shiva Sundaram. Multimodal and multiresolution speech recognition with transformers. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2381–2387, 2020. 2
- [27] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, and Soumith Chintala. Pytorch: An imperative style, high-

- performance deep learning library. In *NeurIPS*, pages 8026–8037, 2019. [6](#)
- [28] Ronald A Rensink. The dynamic representation of scenes. *Visual cognition*, 7(1-3):17–42, 2000. [3](#)
- [29] Ethan Rublee, Vincent Rabaud, Kurt Konolige, and Gary Bradski. Orb: An efficient alternative to sift or surf. In *2011 International conference on computer vision*, pages 2564–2571. Ieee, 2011. [2](#)
- [30] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. [1](#), [2](#), [5](#)
- [31] Matthias Sperber, Jan Niehues, Graham Neubig, Sebastian Stüker, and Alex Waibel. Self-attentional acoustic models. In *19th Annual Conference of the International Speech Communication, INTERSPEECH 2018; Hyderabad International Convention Centre (HICC)Hyderabad; India; 2 September 2018 through 6 September 2018. Ed.: C.C. Sekhar*, pages 3723–3727, 2018. [3](#), [5](#)
- [32] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8934–8943, 2018. [6](#)
- [33] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 5998–6008, 2017. [2](#), [3](#), [4](#), [6](#)
- [34] Fei Wang, Mengqing Jiang, Chen Qian, Shuo Yang, Cheng Li, Honggang Zhang, Xiaogang Wang, and Xiaoou Tang. Residual attention network for image classification. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6450–6458, 2017. [3](#), [5](#)
- [35] Qilong Wang, Banggu Wu, Pengfei Zhu, Peihua Li, Wangmeng Zuo, and Qinghua Hu. Eca-net: Efficient channel attention for deep convolutional neural networks. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11534–11542, 2020. [5](#)
- [36] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 7794–7803, 2018. [3](#)
- [37] Sanghyun Woo, Jongchan Park, Joon-Young Lee, and In So Kweon. Cbam: Convolutional block attention module. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 3–19, 2018. [3](#), [5](#)
- [38] Baosong Yang, Zhaopeng Tu, Derek F. Wong, Fandong Meng, Lidia S. Chao, and Tong Zhang. Modeling localness for self-attention networks. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4449–4458, 2018. [3](#), [5](#)
- [39] Fuzhi Yang, Huan Yang, Jianlong Fu, Hongtao Lu, and Bainig Guo. Learning texture transformer network for image super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5791–5800, 2020. [3](#), [7](#)
- [40] Jun Yu, Jing Li, Zhou Yu, and Qingming Huang. Multimodal transformer with multi-view visual representation for image captioning. *IEEE Transactions on Circuits and Systems for Video Technology*, 30(12):4467–4480, 2020. [2](#)
- [41] Zhou Yu, Jun Yu, Jianping Fan, and Dacheng Tao. Multi-modal factorized bilinear pooling with co-attention learning for visual question answering. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 1839–1848, 2017. [3](#)
- [42] Xiaoyun Yuan, Lu Fang, Qionghai Dai, David J Brady, and Yebin Liu. Multiscale gigapixel video: A cross resolution image matching and warping approach. In *2017 IEEE International Conference on Computational Photography (ICCP)*, pages 1–9. IEEE, 2017. [1](#), [3](#), [6](#), [8](#)
- [43] Rui Zeng, Simon Denman, Sridha Sridharan, and Clinton Fookes. Rethinking planar homography estimation using perspective fields. In *Asian Conference on Computer Vision*, pages 571–586. Springer, 2018. [6](#)
- [44] Jirong Zhang, Chuan Wang, Shuaicheng Liu, Lanpeng Jia, Nianjin Ye, Jue Wang, Ji Zhou, and Jian Sun. Content-aware unsupervised deep homography estimation. In *European Conference on Computer Vision*, 2020. [2](#), [6](#), [7](#), [8](#)
- [45] Jianing Zhang, Tianyi Zhu, Anke Zhang, Xiaoyun Yuan, Zihan Wang, Sebastian Beetschen, Lan Xu, Xing Lin, Qionghai Dai, and Lu Fang. Multiscale-vr: Multiscale gigapixel 3d panoramic videography for virtual reality. In *2020 IEEE International Conference on Computational Photography (ICCP)*, pages 1–12, 2020. [1](#)
- [46] Zhifei Zhang, Zhaowen Wang, Zhe Lin, and Hairong Qi. Image super-resolution by neural texture transfer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7982–7991, 2019. [7](#)
- [47] Mandan Zhao, Gaochang Wu, Yipeng Li, Xiangyang Hao, Lu Fang, and Yebin Liu. Cross-scale reference-based light field super-resolution. *IEEE Transactions on Computational Imaging*, 4(3):406–418, 2018. [3](#)
- [48] Haitian Zheng, Mengqi Ji, Lei Han, Ziwei Xu, Haoqian Wang, Yebin Liu, and Lu Fang. Learning cross-scale correspondence and patch-based synthesis for reference-based super-resolution. In *BMVC*, 2017. [3](#)
- [49] Haitian Zheng, Mengqi Ji, Haoqian Wang, Yebin Liu, and Lu Fang. Crossnet: An end-to-end reference-based super resolution network using cross-scale warping. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 88–104, 2018. [3](#)
- [50] Yuemei Zhou, Gaochang Wu, Ying Fu, Kun Li, and Yebin Liu. Cross-mpi: Cross-scale stereo for image super-resolution using multiplane images. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. [3](#), [6](#), [8](#)