SSORN: Self-Supervised Outlier Removal Network for Robust Homography Estimation

Yi Li*, Wenjie Pei*, Zhenyu He†

School of Computer Science and Technology, Harbin Institute of Technology, Shenzhen, China.

ly_res@163.com, wenjiecoder@outlook.com, zhenyuhe@hit.edu.cn.

Abstract—The traditional homography estimation pipeline consists of four main steps: feature detection, feature matching, outlier removal and transformation estimation. Recent deep learning models intend to address the homography estimation problem using a single convolutional network. While these models are trained in an end-to-end fashion to simplify the homography estimation problem, they lack the feature matching step and/or the outlier removal step, which are important steps in the traditional homography estimation pipeline. In this paper, we attempt to build a deep learning model that mimics all four steps in the traditional homography estimation pipeline. In particular, the feature matching step is implemented using the cost volume technique. To remove outliers in the cost volume, we treat this outlier removal problem as a denoising problem and propose a novel self-supervised loss to solve the problem. Extensive experiments on synthetic and real datasets demonstrate that the proposed model outperforms existing deep learning models.

Index Terms—Homography estimation, outlier removal, self-supervised, cost volume.

I. Introduction

Homography estimation is a fundamental task in many computer vision applications, such as Augmented Reality (AR) [1]–[3], image stitching [4]–[7], and Simultaneous Location And Mapping (SLAM) [8]–[10]. Although this task has been extensively studied in the past, designing a robust homography estimation method remains a challenging problem.

The traditional homography estimation pipeline mainly includes four steps: feature detection, feature matching, outlier removal and transformation estimation. Among these four steps, feature detection is generally considered to be the most important step. Therefore, a large number of popular hand-crafted features, such as SIFT [11], SURF [12] and ORB [13], have been proposed over the past few decades. However, the feature representation capabilities of these feature are limited, especially compared to deep learning features.

Currently, there are two research lines using deep convolutional networks (CNNs) for homography estimation. The first research line is to replace one of the four steps in the traditional homography estimation pipeline with CNNs or neural networks [14]–[18]. For example, SuperPoint [15] adopts a VGG-like architecture for interest point detection and description. SuperGlue [19] solves the matching problem based on a graph neural network. DSAC [16] uses a neural networks to mimic the RANSAC algorithm [17]. However,

optimizing each step individually does not necessarily improve the performance of the homography estimation task.

Another research line is to build an end-to-end model for homography estimation. DeTone *et al.* [20] propose HomographyNet, the first end-to-end model for homography estimation. It adopts a typical CNNs with 8 convolutional layers and 2 fully connected layers, where the convolutional layers are served as the feature detector/extractor and the fully connected layers are served as the homography estimator. However, it lacks the feature matching and outlier removal steps, which are essential in the tradition homography estimation pipeline. Li *et al.* [21] propose the SRHEN model, which explicitly implements the feature matching step through a correspondence layer. SRHEN performs much better than HomographyNet, which proves the importance of the feature matching step. However, the output of the correspondence layer in SRHEN contains many outliers. It is quite arduous for a homography estimator consisting of only fully connected layers to reject these outliers on its own.

In this paper, we propose a Self-Supervised Outlier Removal Network (SSORN) for robust homography estimation. The model is built upon a Siamese structure. It consists of four components, *i.e.* a feature extractor, a feature matching module, an outlier removal module, and a homography estimator, to mimic all four steps in the traditional homography estimation pipeline. The model first projects two input images into the same deep space and obtains two feature maps. Then, the feature matching module, which is based on the cost volume technique [22], computes a cost volume between the extracted feature maps. The cost volume generated by the feature matching module usually contains many outliers. To remove the outliers in the cost volume, we treat this outlier removal problem as a denoising problem and propose a novel self-supervised loss to train the outlier removal module. Finally, the homography estimator learns the mapping from a clean cost volume produced by the outlier removal module to the homography matrix. Table I summaries the difference between our model and existing deep learning models. Experiments on both synthetic and real dataset show that the proposed model outperforms existing deep learning models.

II. Related Work

In this section, we mainly focus on end-to-end homography estimation models. These models can be divided into supervised models and unsupervised models.

---

* Equal contribution. † Corresponding author.

TABLE I
DIFFERENCE BETWEEN OUR MODEL AND EXISTING DEEP LEARNING MODELS. OUR MODEL IS THE ONLY ONE THAT CONTAINS ALL FOUR STEPS IN THE TRADITIONAL HOMOGRAPHY ESTIMATION PIPELINE.

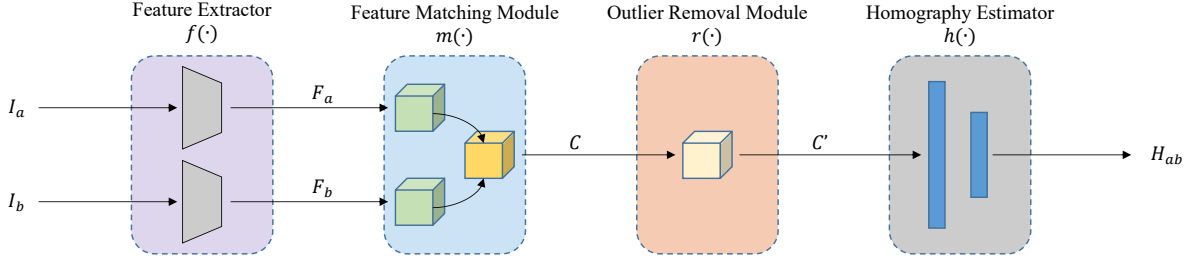|	Feature Extractor	Feature Matching	Outlier Removal	Homography Estimator
HomographyNet [20]	✓			✓
UnsupervisedNet [23]	✓			✓
Zhang et al. [24]	✓			✓
Koguciuk et al. [25]	✓			✓
SRHEN [21]	✓	✓		✓
Ours	✓	✓	✓	✓

Fig. 1. The overall architecture of the proposed model. The model is built upon a Siamese structure. It mimics the traditional homography estimation pipeline using four components: a feature extractor, a feature matching module, an outlier removal module and a homography estimator.

DeTone *et al.* [20] proposes the first supervised deep learning model, named HomographyNet, for homography estimation. Their model employs a VGG-like network structure with eight outputs on top (corresponding to eight parameters of a homography represented by the 4-point parameterization [26]). It takes a concatenated image as input and regresses the homography directly. They also propose a data generation rule for generating training and test datasets based on the COCO dataset [27] for training and evaluating homography estimation models. Japkowicz *et al.* [28] proposes a hierarchical model consisting of multiple DNN models with the same structure to handle large transformations. Each DNN model is trained using image pairs with different scale of transformations, enabling the entire model to perform homography estimation in a coarse-to-fine fashion. Li *et al.* [21] adopts a similar coarse-to-fine framework, but they use different feature maps at different stages and construct multi-level correspondence layers to mimic feature matching step. The above models are designed to solve the general homography estimation problem, while there are other models designed for specific scenarios, such as dynamic scenes [29], cross-resolution scenes [30] and scenes with multiple planes [6].

Nguyen *et al.* [23] treats the homography estimation problem as an unsupervised learning task. The predicted homography, as an intermediate output variable of their model, is used to warp one of the input images. Then, they compute the pixel-wise intensity loss between the warped image and another input image. By minimizing this loss, the model can be trained without ground truth labels. However, the pixel-level intensity loss is very sensitive to illumination changes. For this reason, Zhang *et al.* [24] proposes to compute the loss in the feature space rather than the intensity space. Koguciuk *et al.* [25] further extends this idea by computing the perceptual loss [31], which significantly improves the robustness of their model against illumination changes. Ye *et al.* [32] proposes

to use a homography flow, instead of the commonly used 4-point parameterization, as an intermediate representation of the homography. While their model produces good results on image pairs with small viewpoint changes, it cannot handle image pairs with large viewpoint changes.

III. SSORN

A. Network Architecture

Our model is built upon a Siamese structure. It takes two grayscale images I_a and I_b as input, and produces a homography matrix H_{ab} from I_a to I_b as output. The entire architecture is composed of four components: a feature extractor $f(\cdot)$, a feature matching module $m(\cdot)$, an outlier removal module $r(\cdot)$, and a homography estimator $h(\cdot)$. The resulting model is illustrated in Fig. 1.

Feature extractor. The backbone of the feature extractor follows the ResNet-34 structure. We use the output of *layer2* in the ResNet-34. For a pair of input images I_a and I_b of size $H \times W \times 1$, the feature extractor $f(\cdot)$ produces feature maps F_a and F_b of size $H/8 \times W/8 \times C$:

$$F_a = f(I_a), \quad F_b = f(I_b) \tag{1}$$

Feature matching module. Since homography is characterized by pixel-to-pixel correspondences rather than image features [33], feature matching (or establishing correspondences) is an important step in homography estimation. However, the way of establishing pixel-to-pixel correspondences in the traditional homography estimation pipeline is a non-differentiable operation. An alternative solution for feature matching in deep models is the cost volume technique, which has been widely used in optical flow estimation [22]. The feature matching module $m(\cdot)$ adopts the cost volume technique to bridge the gap between image features and homography. It computes a cost volume C between the extracted feature maps F_a and F_b:
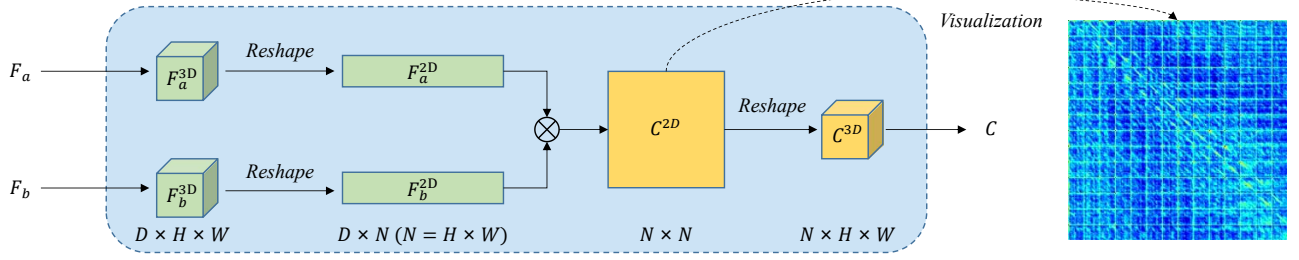
$$C = m(F_a, F_b) \tag{2}$$

Fig. 2. The feature matching module adopts the cost volume technique. It explicitly computes a cost volume between two feature maps. We show an example of the cost volume in 2D form on the right.
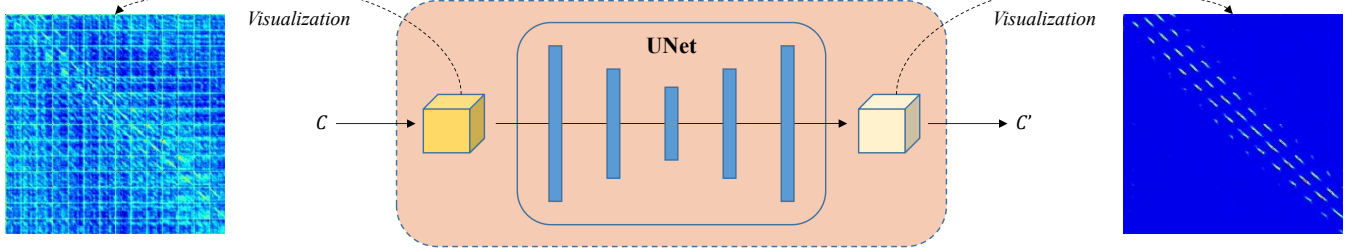

Fig. 3. The outlier removal model uses the UNet structure as the backbone. It successfully removes most of the outliers in the original cost volume C produced by the feature matching module, obtaining a clean cost volume C'.

Outlier removal module. Removing outliers is crucial for robust homography estimation. Traditional outlier removal algorithms like RANSAC [17] or more recent algorithms like DSAC [16] are designed to deal with outliers in pixel-to-pixel correspondences. Unfortunately, these algorithms are not able to handle outliers in a cost volume. In this paper, we treat the problem of removing outliers in a cost volume as a denoising problem and solve it using an outlier removal module. The outlier removal model $r(\cdot)$ takes the cost volume C as input and produces a new cost volume C' as output:

$$C' = r(C) \tag{3}$$

Homography estimator. The homography estimator simply consists of two fully connected layers. It learns the mapping function from the cost volume C' to the homography matrix H_{ab}:

$$H_{ab} = h(C') \tag{4}$$

B. Feature Matching Module

Unlike other components in our model, the feature matching module has no trainable parameters. It explicitly computes a cost volume C between two feature maps F_a and F_b. Intuitively, the cost volume can be thought of as a similarity matrix in 3D form. It stores dense feature matching costs of two sets of feature vectors. The process of computing the cost volume is shown in Fig. 2. For clarity, we use the superscripts "2D" and "3D" to denote 2D and 3D tensors, respectively.

To compute the 3D cost volume C^{3D} from F_a^{3D} to F_b^{3D}, we first reshape F_a^{3D} and F_b^{3D} into corresponding 2D tensors F_a^{2D} and F_b^{2D} respectively. Then, the matching cost $C^{2D}(i,j)$ between the $i-$th feature vector in F_a^{2D} and the $j-$th feature

vector in F_b^{2D} is implemented as the correlation between the feature vectors:

$$C^{2D}(i,j) = \frac{1}{D}(F_a^{2D}(i))^T \odot F_b^{2D}(j) \tag{5}$$

where $\odot$ denotes the dot product, T is the transpose operator, and D is the dimension of feature vectors. Accordingly, the full cost volume C^{2D} is defined as:

$$C^{2D} = \frac{1}{D}(F_a^{2D})^T \otimes F_b^{2D} \tag{6}$$

where $\otimes$ denotes the matrix product. Finally, the 2D cost volume C^{2D} will be reshaped into the corresponding 3D cost volume C^{3D}.

C. Outlier Removal Module

In the traditional view, outliers usually mean incorrect correspondences. In a similar spirit, we can treat those incorrect matching costs in a cost volume as outliers in the cost volume. Since a cost volume stores dense matching costs, it often contains a large number of outliers due to the existence of many similar image features. It is difficult for the homography estimator to predict an exact homography matrix from this cost volume. To this end, we propose the outlier removal model, which treats the outlier removal problem as a cost volume denoising problem. We believe that, for a specific homography matrix, there exists a specific pattern (*i.e.* a clean cost volume) corresponding to that homography matrix. The goal of the outlier removal module is to recover this pattern from the noisy cost volume produced by the feature matching module.

Apparently, the cost volume is different from natural images, thus the priors for the natural image denoising problem
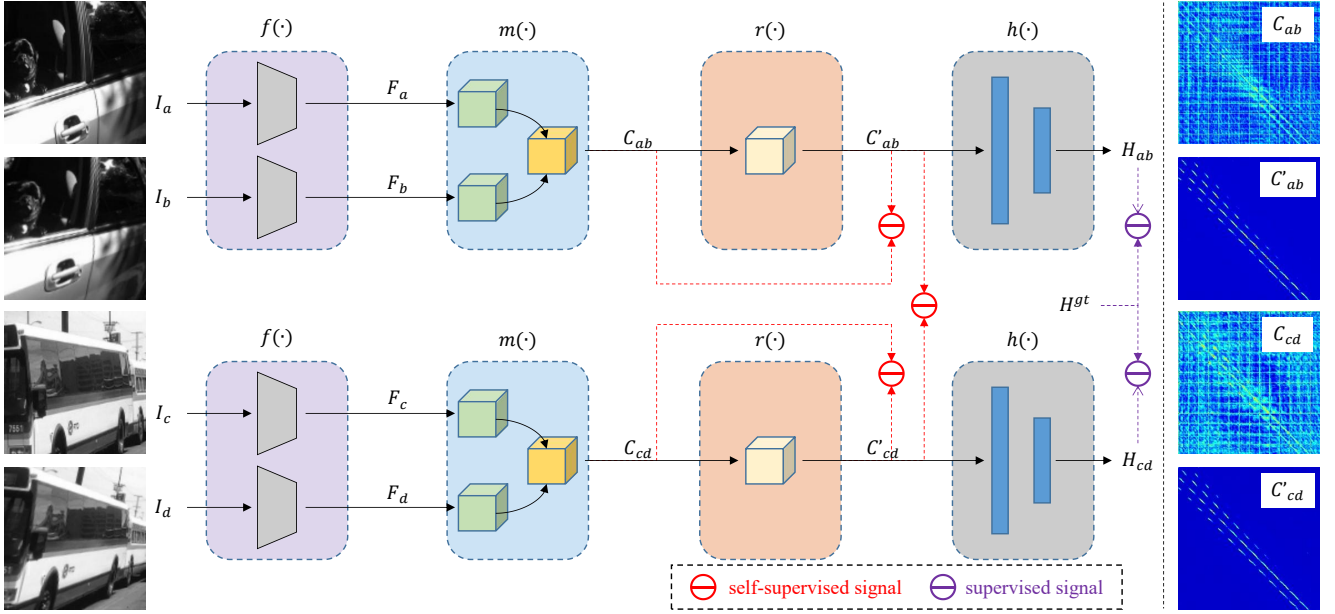
Fig. 4. The training process of the whole model. Our model requires two image pairs instead of one image pair for training. The two image pairs are related by the same ground truth homography matrix H^{gt}. The training loss consists of three self-supervised signals (Eq. 8) and two supervised signals (Eq. 9). For these two image pairs, the two cost volumes C_{ab} and C_{cd} directly produced by the feature matching module may look very different, but the two clean cost volumes C'_{ab} and C'_{cd} produced by the outlier removal module always look the same.

are not suitable for the cost volume denoising problem. To address the training problem, we propose a novel self-supervised loss function as described in Sec. III-D. Theoretically, any deep denoising model can be used as the outlier removal module. In practice, we adopt the UNet structure as the backbone of the outlier removal model. After training with the proposed self-supervised loss function, the outlier removal module is able to remove most of the outliers in the original cost volume C produced by the feature matching module, obtaining a clean cost volume C', as shown in Fig. 3. As expected, for a specific homography matrix, the outlier removal module successfully finds a specific pattern (*i.e.* the clean cost volume C') corresponding to that homography matrix.

D. Self-Supervised Training

Given a pair of images I_a and I_b and the ground truth homography matrix H^{gt}_{ab}, one simple way to train the model is to use the following supervised loss:

$$L = \|H_{ab} - H^{gt}_{ab}\|_2^2 \tag{7}$$

where $\|\cdot\|_2^2$ is L2 norm.

However, this way does not take full advantage of the outlier removal module, because the above loss function does not ensure that the outlier removal module learns the ability to remove outliers in the cost volume. To this end, we propose an additional self-supervised loss function to train the outlier removal module. We exploit the fact that two image pairs related by the same homography matrix should correspond to a similar clean cost volume. We force the outlier removal module to produce similar outputs for these two image pairs. Fig. 4 illustrates the training process of the whole model.

Formally, given two image pairs (I_a, I_b) and (I_c, I_d) related by the same ground truth homography matrix H^{gt}, the self-supervised loss for training the outlier removal module is designed as:

$$\begin{aligned} L_{ss} = {} & \lambda_1 \|C'_{ab} - C'_{cd}\|_1 \\ & + \lambda_2(\|C'_{ab} - C_{ab}\|_1 + \|C'_{cd} - C_{cd}\|_1) \end{aligned} \tag{8}$$

where $\|\cdot\|_1$ is L1 norm. The first half of the equation forces the outlier removal module to produce similar cost volumes, and the second half of the equation is used to avoid trivial solutions, *i.e.* $C'_{ab} = C'_{cd} = 0$.

To train the homography estimator, we still need the following supervised loss:

$$L_s = \|H_{ab} - H^{gt}\|_2^2 + \|H_{cd} - H^{gt}\|_2^2 \tag{9}$$

Combining the above two losses, we get the final loss for training the whole model:

$$L_f = L_s + L_{ss} \tag{10}$$

E. Discussion

If the idea that a specific homography corresponds to a specific clean cost volume is true, then we seem to be able to derive the grund truth cost volume directly from a ground truth homography, and thus allowing for simple supervised training. In theory, given a ground truth homography and the coordinate of a pixel on one image, we can always compute the coordinate of the unique corresponding pixel on another image, from which we can derive the ground truth cost volume. However, this derivation does not take into account the local self-similarity property of images, that is, adjacent pixels usually have similar (deep) features. In other words, a pixel
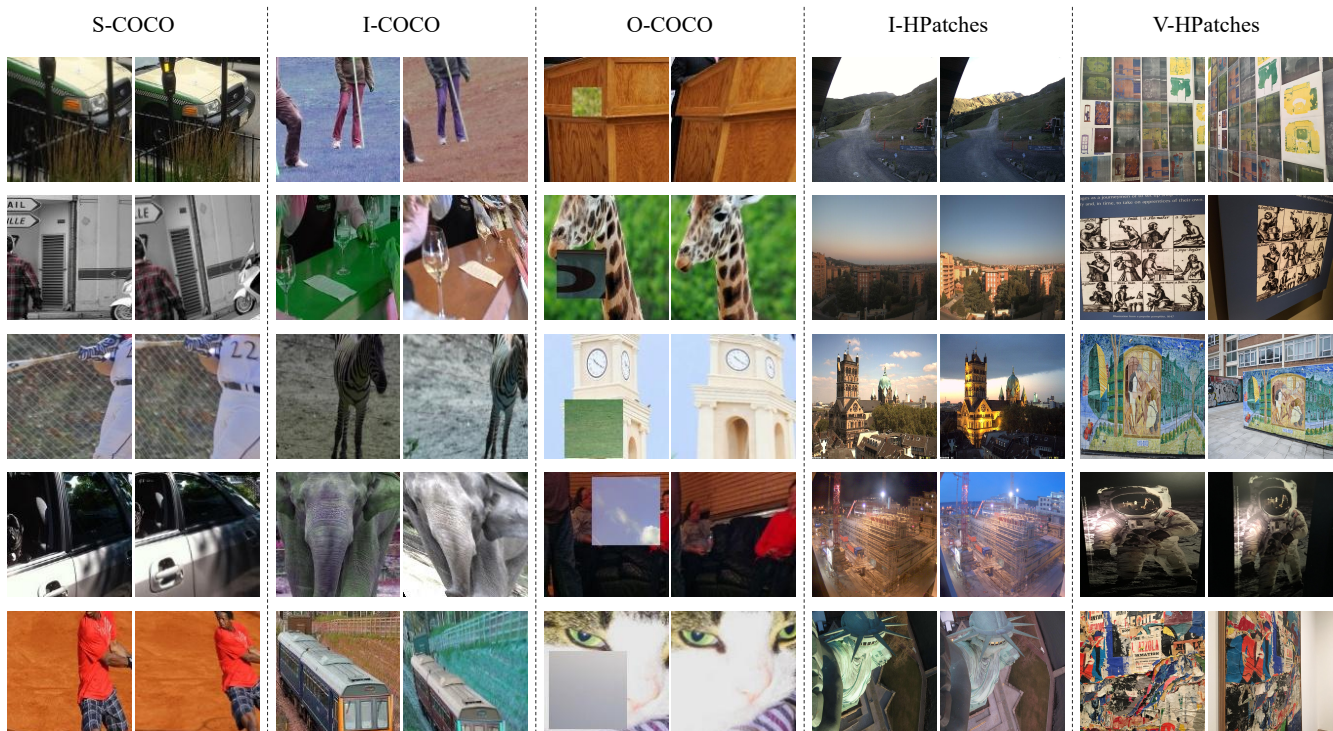
Fig. 5. Some examples from the datasets used in experiments. Each example consists of two images related by a specific homography matrix. The goal of homography estimation is to predict the homography matrix between the input image pair.

on one image often has multiple high-scoring matches on another image, not just one match (see the visualization of the clean cost volume in Fig. 3). It is not trivial to take the local self-similarity property of images into account when deriving the ground truth cost volume from a ground truth homography. This motivates us to propose the self-supervised training, which avoids the difficulty of directly deriving the ground truth cost volume.

IV. EXPERIMENTS

A. Datasets and Evaluation Metric

Although there exist several datasets [34]–[36] for homography estimation, these datasets are too small for training deep homography estimation models. To this end, Detone *et al.* [20] proposed a synthetic dataset S-COCO based on the COCO dataset [27]. S-COCO is generated by applying random projective transformations to the COCO dataset. By using the same data generation rule, we can easily generate two image pairs with the same homography matrix for training our model.

Based on S-COCO, we additionally generate two datasets, called I-COCO and O-COCO, for testing the robustness of different models to illumination and occlusion. In I-COCO, we apply the photometric distortion techniques in [37] to S-COCO. The illumination change is controlled by a parameter δ ranging from 8 to 32 with an interval of 8. In O-COCO, we simulate occlusion by overlaying a random image patch of size p on one of the input images, where p ranges from 0 to 80 with an interval of 10.

In addition to the above three synthetic datasets, we also conduct experiments on a real dataset, HPatches [35]. The

dataset contains two subsets: the illumination subset (*i.e.* I-HPatches) and the viewpoint subset (*i.e.* V-HPatches). As this dataset only contains 580 pairs of images, we only use it for testing.

Fig. 5 shows some examples from the above datasets. All images are resized to 256×256 for training and testing. To measure the performance of different models, we adopt the Mean Average Corner Error (MACE) [20] as the evaluation metric.

B. Experimental Setup

We compare our model with 6 recently proposed deep learning models (*i.e.* HomographyNet [20], UnsupervisedNet [23], SRHEN [21], Zhang's *et al.* [24] and Koguciuk's *et al.* [25] models) and two representative keypoint based methods (*i.e.* SIFT [11] and SuperPoint [15]).

To ensure that the performance difference between deep learning methods comes from better network structure and loss function design, rather than better CNN backbone, we implement these methods to use the same ResNet34 structure as the backbone. Using the same backbone and the same training settings, most deep learning models achieve very similar performance, but this is not the case for SRHEN and our model (see experiments in the following sections). For keypoint based methods, we compare two feature matchers: the nearest neighbor (NN) with Lowe's ratio test (ratio is set to 0.9) and SuperGlue (SG). MAGSAC++ [38] is used as the robust estimator.

We adopt the 4-point parameterization [26] to represent the homography matrix instead of 3×3 matrix form. The weight
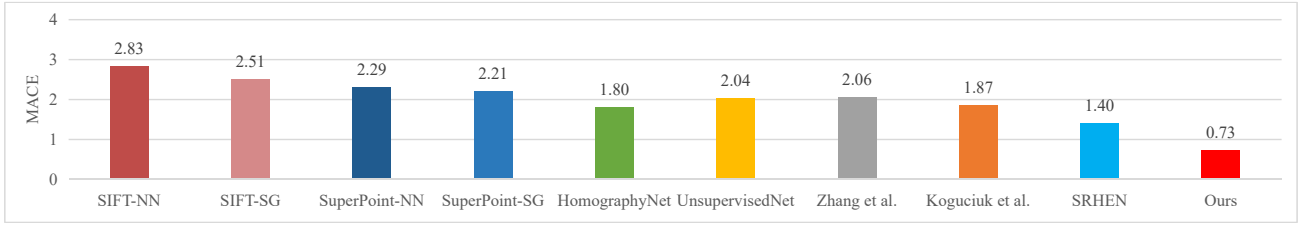
Fig. 6. MACE for different methods on the S-COCO dataset. Lower MACE means better performance. NN - nearest neighbor. SG - SuperGlue.
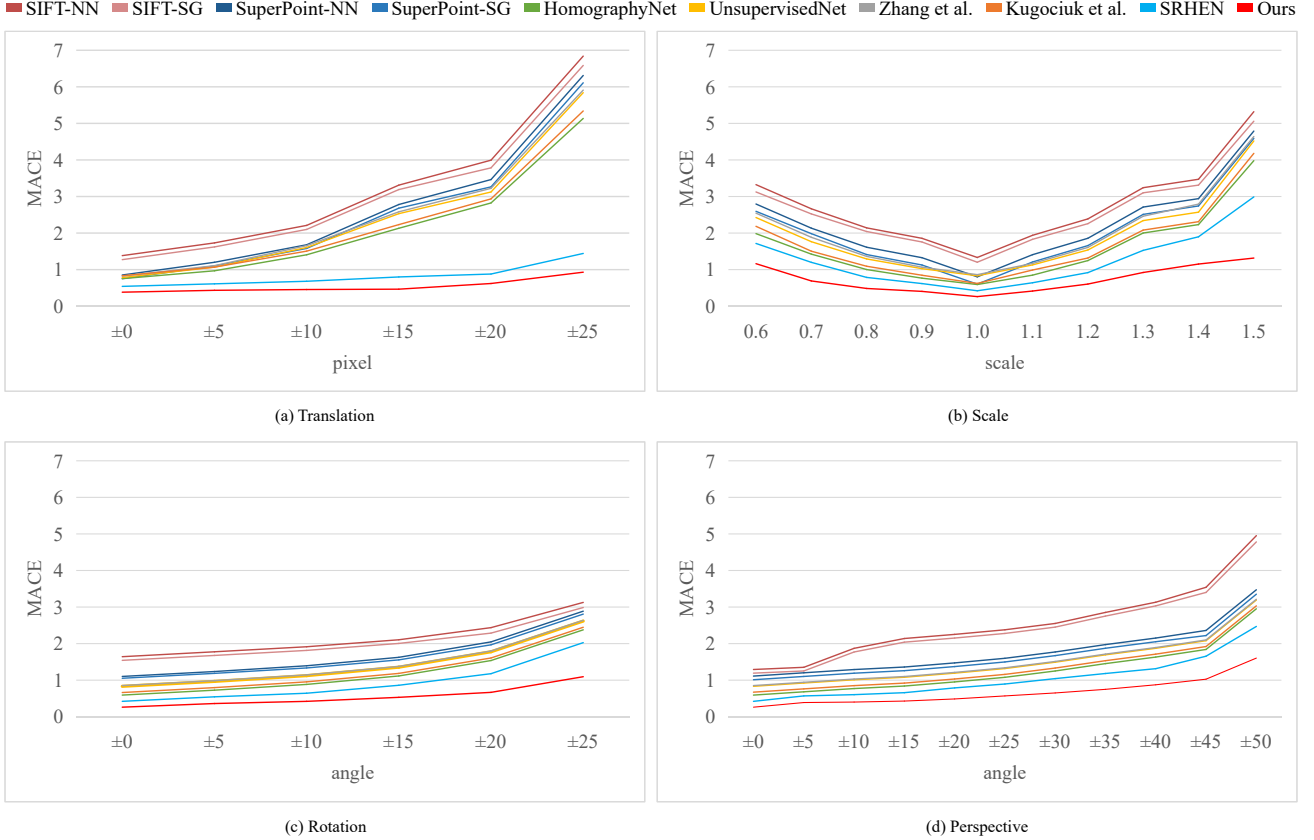

Fig. 7. Results of different methods under four types of geometric transformations: (a) Translation, (b) Scale, (c) Rotation, and (d) Perspective.

parameters λ_1 and λ_2 in Eq. 8 are empirically set to 0.5 and 0.25, respectively.

C. Results on S-COCO Dataset

We first conduct experiments on the S-COCO dataset. Fig. 6 reports the MACE of different methods on this dataset. In general, deep learning methods outperform keypoint based methods on this dataset, mainly due to the powerful feature representation capability of deep learning features. Compared with those deep models without the feature matching module, SRHEN with the feature matching module achieves a better result. This demonstrates the necessity of the feature matching module not only in the traditional homography estimation pipeline but also in deep homography estimation models. Our model contains not only the feature matching module but also an outlier removal module. It leverages the strength of the traditional homography estimation pipeline and the strength

of deep learning features. It obtains a much better result than SRHEN on the S-COCO dataset, which indicates that the outlier removal module is as important as the feature matching module for deep homography estimation models.

To analyze the robustness of different models to different geometric transformations, we generated four subsets according to the data generation rule in [20]. Each subset corresponds to a specific type of transformation, namely translation, scale, rotation and perspective transformaitons. We plot the results in Fig. 7. It manifests that all methods perform well when the magnitude of the geometric transformation is small. However, the performance degrades as the magnitude of the geometric transformation increases. Specifically, deep learning models that do not employ the feature matching module perform poorly for large geometric transformations, especially for large translations. Overall, our model is more robust to these four types of geometric transformations than other methods.

(a) Results on the I-COCO dataset.
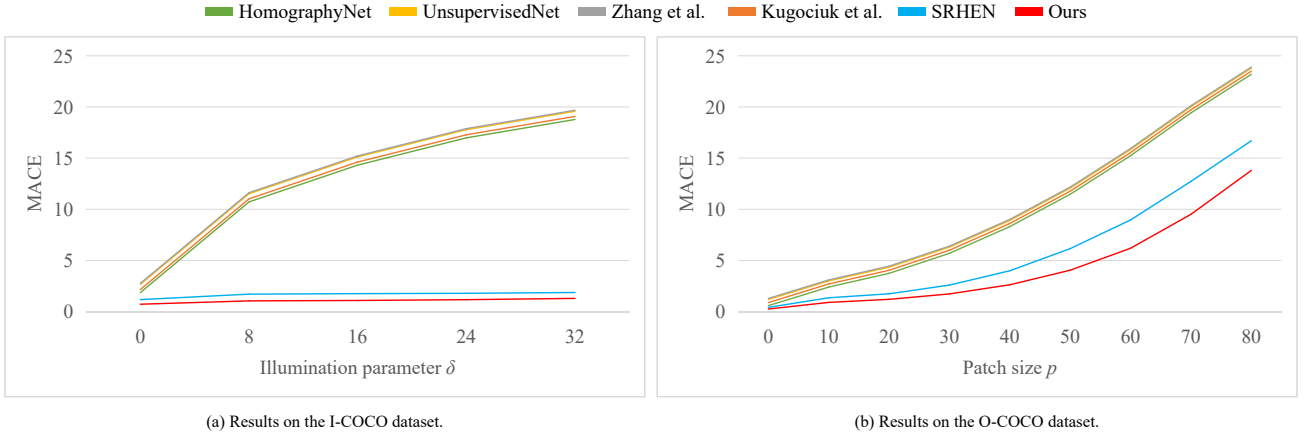
(b) Results on the O-COCO dataset.

Fig. 8. Results of the first experimental configuration: all models are trained on the S-COCO dataset and tested on the (a) I-COCO dataset and (b) O-COCO dataset.

(a) Results on the I-COCO dataset.
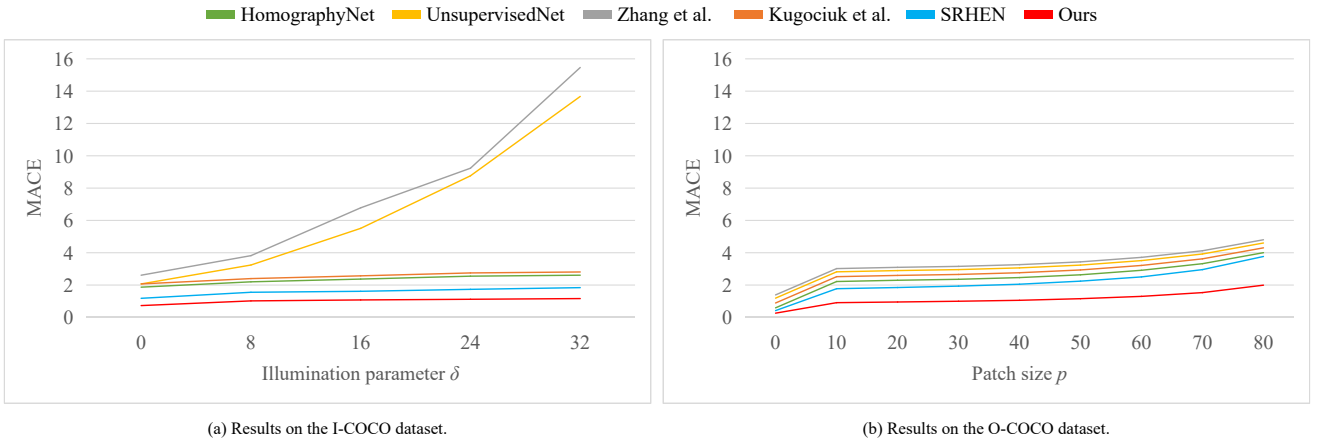
(b) Results on the O-COCO dataset.

Fig. 9. Results of the second experimental configuration: all models are (a) trained and tested on the I-COCO dataset and (b) trained and tested on the O-COCO dataset.

D. Robustness to Occlusion and Illumination

To investigate the robustness or the generalization ability of deep learning models to illumination and occlusion, we conduct experiments on the I-COCO and O-COCO datasets. We employ two experimental configurations. In the first experiment, we train all deep learning models on the S-COCO dataset and test them on the I-COCO/O-COCO dataset. In the second experiment, all models are trained and tested on the I-COCO/O-COCO dataset. The experimental results are shown in Fig. 8 and Fig. 9.

Interestingly, we observe that the deep learning models (*i.e.* HomographyNet, UnsupervisedNet, Zhang's *et al.* and Kugociuk's *et al.* models) without the feature matching module obtain significantly different results under the two experimental configurations. This demonstrates that learning the mapping function directly from image features to the homography matrix may lead to the generalization problem. In contrast, SRHEN and our model achieve more stable results under these two experimental configurations. Compared to SRHEN, our model is more robust to illumination and occlusion, as the outlier removal module is able to recover the "correct" clean cost volume in most cases (see Fig. 10).

E. Results on HPatches Dataset

In the previous sections, we have tested different methods on three synthetic datasets. In order to investigate the generalization ability of these methods to real images, we conduct experiments on the HPatches dataset. Since there are no large-scale real dataset for training the deep learning models, we train all the models on the S-COCO dataset and test them on the HPatches dataset. The experimental results are reported in Fig. 11.

As shown in Fig. 11, our model achieves a distinct advantages on the HPatches dataset among all deep learning models. However, it is inferior to SIFT and SuperPoint. A major reason may be that the distributions of training images and test images are quite different. The images in the HPatches dataset contains a lot of local illumination variations and non-planar regions. These challenging factors are difficult to simulate by synthetic means. We show some success and failure cases of our model in Fig. 12.

Currently, deep learning models do not achieve satisfactory results on the HPatches dataset. However, they have the advantage of being able to generalize to a variety of scenarios when training on a large number of images, as demonstrated by
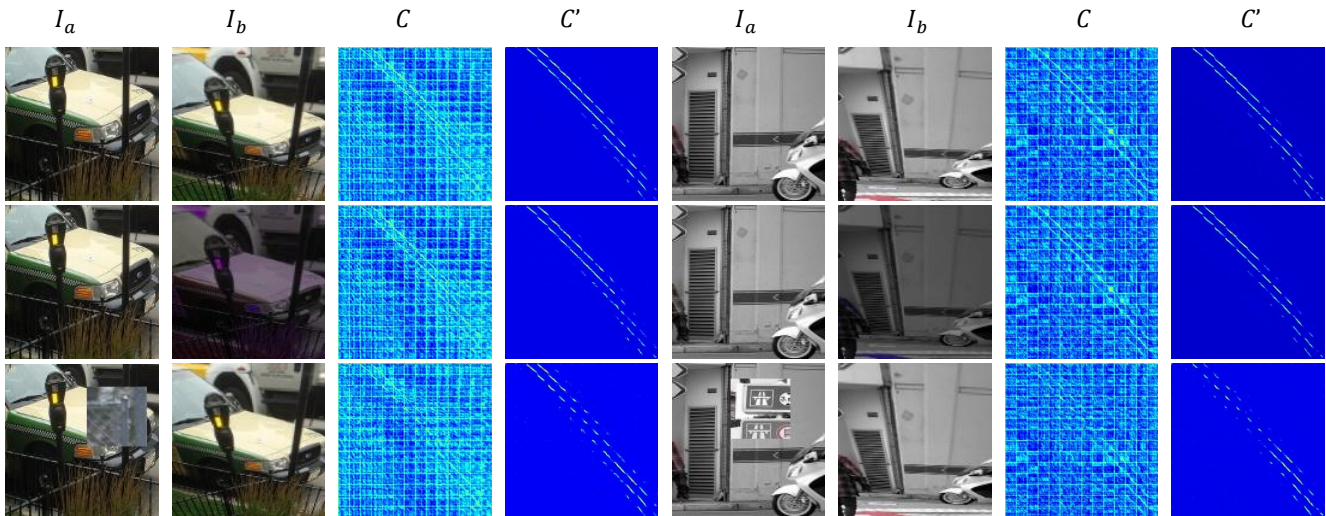
Fig. 10. All image pairs in this figure are related by the same homography matrix. While the cost volume C produced by the feature matching module may be different due to illumination changes (the second row) or occlusions (the third row), the outlier removal model is able to remove most of the outliers in C and produce the "correct" clean cost volume C'.
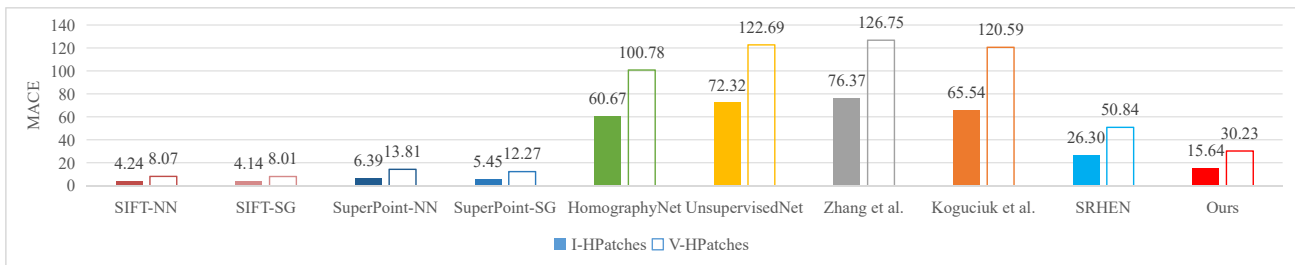

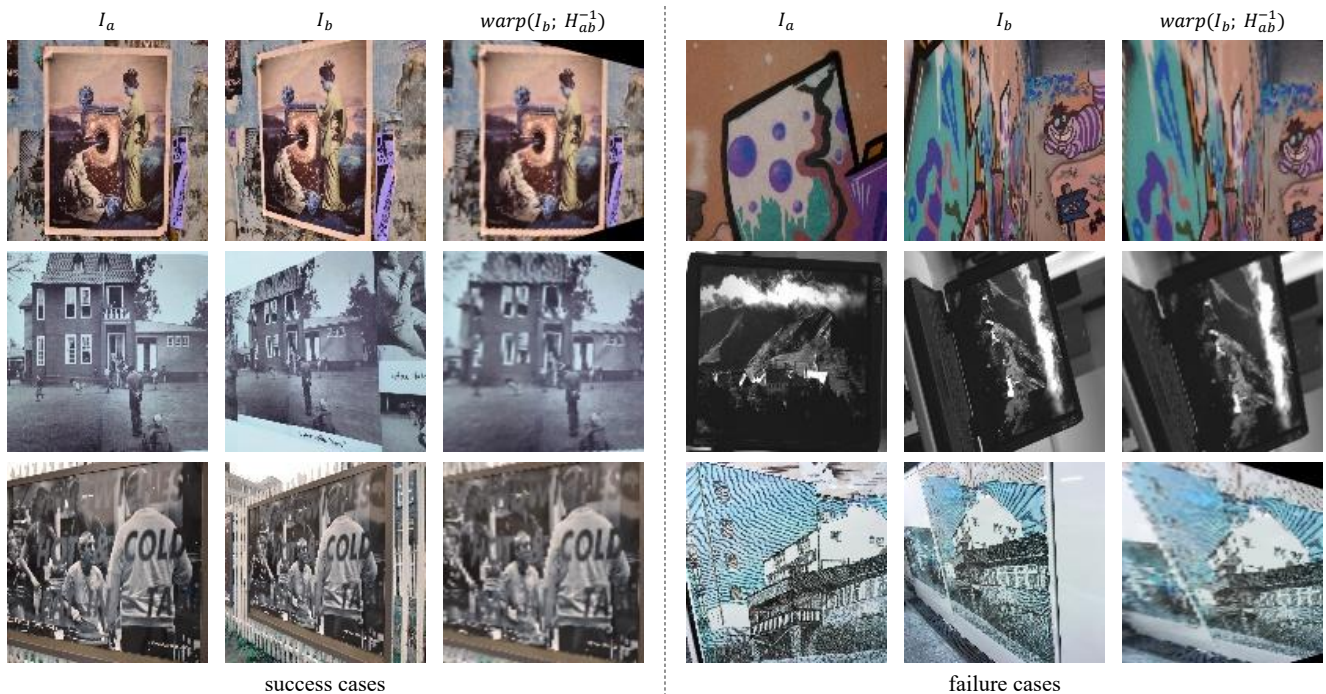Fig. 11. MACE for different methods on the HPatches dataset. Lower MACE means better performance.

Fig. 12. Some success and failure cases of our model on the HPatches dataset. We warp the image I_b with the inverse of the predicted homography matrix H_{ab}. The more accurate the predicted homography matrix, the more aligned I_a and the warped image will be.

	FH	FMH	FMRH-s	FMRH-ss
MACE	1.79	1.38	1.26	**0.73**

previous experiments on the synthetic datasets. We believe that the performance of our model, as well as other deep learning models, on the HPatches dataset will be further improved if large-scale real dataset are available for training in the future.

F. Ablation Study

In this section we conduct experiments to investigate the effectiveness of each proposed technique in our model. To this end, we perform ablation experiments on four variants of our model:

- **FH**, which only contains the feature extractor and the homography estimation. It learns the mapping function from image features to the homography matrix.
- **FMH**, which augments **FH** with a feature matching module. It learns the mapping function from the cost volume produced by the feature matching module to the homography matrix.
- **FMRH-s**, which further augments **FMH** with an outlier removal module. It mimics the traditional homography estimation pipeline in a single deep learning model. It is trained with the training loss in Eq. 9.
- **FMRH-ss**, has exactly the same network structure as **FMRH-s**, but is trained with the training loss in Eq. 8.

Table II reports comparison results on the S-COCO dataset between these four variants. **FH**, without using the feature matching module and the outlier removal module, is similar to HomographyNet with different feature learning structures. Thus, **FH** performs similarly to HomographyNet on the S-COCO dataset. The performance comparison between **FH** and **FMH** manifests the benefit of leveraging the feature matching module to bridge the gap between image features and the homography matrix. Although **FMRH-s** is further equiped with an outlier removal module based on **FMH**, the outlier removal module can not be well trained with the supervised loss in Eq. 9. As a consequence, **FMRH-s** is only slightly better than **FMH**. The performance is improved marginally from **FMRH-s** to **FMRH-ss**, which is benefited from the proposed self-supervised loss for training the outlier removal module.

Besides the network structure and the loss function, in practice, we find that the backbone of the feature extractor and the outlier removal module also has a significant impact on the performance of our model. Early deep learning models usually employ VGG as the backbone of the feature extractor. However, recent studies [24], [25] have shown that replacing VGG with ResNet-34 always results in better performance in the homography estimation task. Therefore, we choose ResNet-34 as the backbone of the feature extractor in our model. Using the ResNet-34, our model achieves a MACE of 0.73 on the S-COCO dataset. While using the VGG, the MACE is 2.36. For fair comparison, in previous experiments, we implemented all deep learning models to use the same ResNet-34 as the backbone of the feature extractor. As for the outlier removal module,we started with a classic deep denoising model DnCNN [39], but later found that this model performed slightly worse than the currently used model UNet.

V. CONCLUSION

In this work, we have presented a new deep learning model that mimics the traditional homography estimation pipeline. Specifically, the proposed model consists of four components: a feature extractor, a feature matching module, an outlier removal module, and a homography estimator. The feature matching module is implemented using the cost volume technique and has no trainable parameters. The outlier removal module is built based on the UNet structure to remove outliers in the cost volume. We propose a novel self-supervised loss to train the outlier removal module. The entire model can be trained in an end-to-end fashion with a combination of a supervised and the self-supervised loss. Extensive experiments on synthetic and real datasets show that the proposed model significantly outperforms existing deep learning models.

REFERENCES

[1] G. Klein and D. W. Murray, "Parallel tracking and mapping for small ar workspaces," in *IEEE International Symposium on Mixed and Augmented Reality*, 2007, pp. 1–10.
[2] F. Zhou, H. B. Duh, and M. Billinghurst, "Trends in augmented reality tracking, interaction and display: A review of ten years of ismar," in *IEEE International Symposium on Mixed and Augmented Reality*, 2008, pp. 193–202.
[3] F. Tang, Y. Wu, X. Hou, and H. Ling, "3d mapping and 6d pose computation for real time augmented reality on cylindrical objects," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 30, no. 9, pp. 2887–2899, 2019.
[4] C. Chang, C. Chou, and E. Y. Chang, "Clkn: Cascaded lucas-kanade networks for image alignment," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 3777–3785.
[5] B. Chung and C. Yim, "Bi-sequential video error concealment method using adaptive homography-based registration," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 30, no. 6, pp. 1535–1549, 2019.
[6] L. Nie, C. Lin, K. Liao, S. Liu, and Y. Zhao, "Depth-aware multi-grid deep homography estimation with contextual correlation," *IEEE Transactions on Circuits and Systems for Video Technology*, 2021.
[7] W. Xue, W. Xie, Y. Zhang, and S. Chen, "Stable linear structures and seam measurements for parallax image stitching," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 32, no. 1, pp. 253–261, 2021.
[8] R. Murartal, J. M. M. Montiel, and J. D. Tardos, "Orb-slam: A versatile and accurate monocular slam system," *IEEE Transactions on Robotics*, vol. 31, no. 5, pp. 1147–1163, 2015.
[9] R. Murartal and J. D. Tardos, "Orb-slam2: An open-source slam system for monocular, stereo, and rgb-d cametas," *IEEE Transactions on Robotics*, vol. 33, no. 5, pp. 1255–1262, 2017.
[10] X. Shao, L. Zhang, T. Zhang, Y. Shen, and Y. Zhou, "Mofisslam: A multi-object semantic slam system with front-view, inertial and surround-view sensors for indoor parking," *IEEE Transactions on Circuits and Systems for Video Technology*, 2021.
[11] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *International Journal of Computer Vision*, vol. 60, no. 2, pp. 91–110, 2004.
[12] H. Bay, T. Tuytelaars, and L. V. Gool, "Surf: Speeded up robust features," in *European Conference on Computer Vision*, 2006, pp. 404–417.
[13] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski, "Orb: An efficient alternative to sift or surf," in *IEEE International Conference on Computer Vision*, 2011, pp. 1–9.

[14] J. Ye, S. Zhang, T. Huang, and Y. Rui, "Cdbin: Compact discriminative binary descriptor learned with efficient neural network," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 30, no. 3, pp. 862–874, 2019.

[15] D. Detone, T. Malisiewicz, and A. Rabinovich, "Superpoint: Self-supervised interest point detection and description," in *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2018, pp. 224–236.

[16] E. Brachmann, A. Krull, S. Nowozin, J. Shotton, F. Michel, S. Gumhold, and C. Rother, "Dsac-differentiable ransac for camera localization," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 6684–6692.

[17] M. A. Fischler and R. C. Bolles, "Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography," *Communications of the ACM*, vol. 24, no. 6, pp. 381–395, 1981.

[18] X. Zhao, J. Liu, X. Wu, W. Chen, F. Guo, and Z. Li, "Probabilistic spatial distribution prior based attentional keypoints matching network," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 32, no. 3, pp. 1313–1327, 2021.

[19] P.-E. Sarlin, D. DeTone, T. Malisiewicz, and A. Rabinovich, "Superglue: Learning feature matching with graph neural networks," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2020, pp. 4938–4947.

[20] D. DeTone, T. Malisiewicz, and A. Rabinovich, "Deep image homography estimation," *arXiv*, 2016.

[21] Y. Li, W. Pei, and Z. He, "Srhen: Stepwise-refining homography estimation network via parsing geometric correspondences in deep latent space," in *ACM International Conference on Multimedia*, 2020, pp. 3063–3071.

[22] D. Sun, X. Yang, M.-Y. Liu, and J. Kautz, "Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 8934–8943.

[23] T. Nguyen, S. W. Chen, S. S. Shivakumar, C. J. Taylor, and V. Kumar, "Unsupervised deep homography: A fast and robust homography estimation model," in *IEEE International Conference on Robotics and Automation*, 2018, pp. 2346–2353.

[24] J. Zhang, C. Wang, S. Liu, L. Jia, N. Ye, J. Wang, J. Zhou, and J. Sun, "Content-aware unsupervised deep homography estimation," in *European Conference on Computer Vision*, 2020, pp. 653–669.

[25] D. Koguciuk, E. Arani, and B. Zonooz, "Perceptual loss for robust unsupervised homography estimation," in *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2021, pp. 4274–4283.

[26] S. Baker, A. Datta, and T. Kanade, "Parameterizing homographies," Carnegie Mellon University, Tech. Rep. CMU-RI-TR-06-11, March 2006.

[27] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft coco: Common objects in context," in *European Conference on Computer Vision*, 2014, pp. 740–755.

[28] N. Japkowicz, F. E. Nowruzi, and R. Laganiere, "Homography estimation from image pairs with hierarchical convolutional networks," in *IEEE International Conference on Computer Vision Workshops*, 2017, pp. 913–920.

[29] H. Le, F. Liu, S. Zhang, and A. Agarwala, "Deep homography estimation for dynamic scenes," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2020, pp. 7652–7661.

[30] R. Shao, G. Wu, Y. Zhou, Y. Fu, L. Fang, and Y. Liu, "Localtrans: A multiscale local transformer network for cross-resolution homography estimation," *arXiv*, 2021.

[31] J. Justin, A. Alexandre, and F.-F. Li, "Perceptual losses for real-time style transfer and super-resolution," in *European Conference on Computer Vision*, 2016, pp. 694–711.

[32] N. Ye, C. Wang, H. Fan, and S. Liu, "Motion basis learning for unsupervised deep homography estimation with subspace projection," *arXiv*, 2021.

[33] R. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2000.

[34] K. Mikolajczyk, T. Tuytelaars, C. Schmid, A. Zisserman, J. Matas, F. Schaffalitzky, T. Kadir, and L. V. Gool, "A comparison of affine region detectors," *International Journal of Computer Vision*, vol. 65, no. 1, pp. 43–72, 2005.

[35] V. Balntas, K. Lenc, A. Vedaldi, and K. Mikolajczyk, "Hpatches: A benchmark and evaluation of handcrafted and learned local descriptors," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 5173–5182.

[36] J. Ma, X. Jiang, A. Fan, J. Jiang, and J. Yan, "Image matching from handcrafted to deep features: A survey," *International Journal of Computer Vision*, vol. 129, no. 1, pp. 23–79, 2021.

[37] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg, "Ssd: Single shot multibox detector," in *European Conference on Computer Vision*, 2016, pp. 21–37.

[38] D. Barath, J. Matas, and J. Noskova, "Magsac: marginalizing sample consensus," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 10 197–10 205.

[39] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang, "Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising," *IEEE Transactions on Image Processing*, vol. 26, no. 7, pp. 3142–3155, 2017.