

Dense Cross-Scale Image Alignment with Fully Spatial Correlation and Just Noticeable Difference Guidance

Jinkun You¹, Jiaxue Li², Jie Zhang¹, Yicong Zhou^{1*}

¹Department of Computer and Information Science, University of Macau

²School of Artificial Intelligence, China University of Geosciences (Beijing)

yujinkun09@gmail.com, lijiaxue7@gmail.com, jiezh1997@gmail.com, yicongzhou@um.edu.mo

Abstract

Existing unsupervised image alignment methods exhibit limited accuracy and high computational complexity. To address these challenges, we propose a dense cross-scale image alignment model. It takes into account the correlations between cross-scale features to decrease the alignment difficulty. Our model supports flexible trade-offs between accuracy and efficiency by adjusting the number of scales utilized. Additionally, we introduce a fully spatial correlation module to further improve accuracy while maintaining low computational costs. We incorporate the just noticeable difference to encourage our model to focus on image regions more sensitive to distortions, eliminating noticeable alignment errors. Extensive quantitative and qualitative experiments demonstrate that our method surpasses state-of-the-art approaches.

Introduction

Image alignment has received increasing attention in recent years (Jiang et al. 2022b; Kim et al. 2024; Li and Zhou 2025). This task involves spatially registering images of the same scene captured from different viewpoints, aiming to maximize the consistency of overlapping regions. It forms the foundation for various applications, including deep camera calibration (Liao et al. 2023), surround-view imaging systems (Kumar et al. 2023), and multimodal image processing (Zhao et al. 2023; You, Zhang, and Zhou 2025). This growing importance has driven efforts to optimize alignment accuracy and efficiency for practical deployment.

Traditional alignment methods rely on hand-crafted features to identify matched points between input images (Li and Zhou 2022; Li, Tofghi, and Monga 2019; Lin et al. 2016). They establish transformations based on the point correspondences. Early approaches employ a single homography to achieve satisfactory results when the images share a common plane or differ only in camera rotation (Li and Zhou 2024). However, their performance degrades under parallax conditions. To overcome this limitation, non-linear local deformations were introduced to offer greater flexibility by predicting a mesh instead of a single homography. The thin-plate spline transformation has been adopted to enhance the alignment performance. Higher accuracy leads to

greater content similarity in the overlapping regions of the aligned images. Nevertheless, focusing solely on accuracy can lead to projective distortion in non-overlapping areas. It causes straight lines to bend and objects to be over-stretched. Recent advances have incorporated geometric constraints to better preserve object shapes (Jia et al. 2021).

While traditional methods yield promising results, they struggle with low-texture and large-parallax cases (Nie et al. 2023). Deep networks have been utilized to address these limitations since they have powerful feature extraction and representation abilities (Deng, Luo, and Chen 2024; You and Zhou 2024; Ou et al. 2024). They extract feature maps to generate a correlation tensor that encodes the matching relationships between images. This tensor is used to predict position offsets between image pixels to do image warping effectively (DeTone, Malisiewicz, and Rabinovich 2016). Unsupervised methods are preferred recently since they do not need the ground truth offsets to train deep networks and are more suitable for real-world image pairs (Nie et al. 2021; Jia et al. 2023). Early unsupervised methods estimate a single homography by maximizing content similarity in the overlapping regions of aligned images (Nguyen et al. 2018). To better handle parallax problems, the non-linear local alignment strategy is integrated in a coarse-to-fine manner (Nie et al. 2022b). Global offsets are first estimated from global features to maintain global consistency. Next, local offsets are computed from fine features for precise alignment. The global and local offsets are integrated to obtain the mesh used for image warping. Shape preservation is enforced by constraining the magnitudes and directions of the mesh.

Existing unsupervised methods have achieved impressive alignment accuracy, but they still face three key limitations. First, they do not explicitly account for scale variations. The apparent size of the same object can differ across images due to parallax, complicating the alignment process. Current approaches typically predict the transformation parameters using features of identical spatial dimensions. These intra-scale features cannot effectively capture scale discrepancies. Second, widely used correlation calculation methods are either ineffective or inefficient. The correlation layer is efficient but destroys spatial information (Rocco, Arandjelovic, and Sivic 2017). The contextual correlation layer and cost volume capture long-range correlation information to achieve better performance at the cost of high FLOPs and

*Corresponding author.

runtime, respectively (Nie et al. 2022b). Their advantages diminish when other modules also capture context information. Third, prevailing methods treat all image pixels equally during optimization. Not all pixels exhibit the same tolerance to distortion (Jiang et al. 2022a). Alignment errors are more noticeable to the human visual system when they occur in pixels that are more sensitive to distortion. As a result, the content similarity is more easily reduced for the overlapping areas of the aligned images.

To tackle these challenges, we introduce an unsupervised image alignment method with the following contributions:

- We design a dense cross-scale image alignment model that leverages cross-scale features to enhance alignment accuracy. It allows for flexible trade-offs between effectiveness and efficiency based on user requirements.
- We introduce a fully spatial correlation module that flexibly utilizes the spatial information of the input features to further increase alignment accuracy. It can also maintain low computational costs.
- The just noticeable distortion estimation is used to guide image alignment. It helps an alignment model prioritize pixels with lower distortion tolerance, thus further increasing alignment accuracy.
- Extensive ablation studies and comparisons are conducted to validate the effectiveness and efficiency of the proposed method.

Proposed Method

Overall Structure

The non-linear local deformation is employed to warp the target image I_{tar} for the alignment with the reference image I_{ref} . It is implemented via a mesh $M^f \in \mathbb{R}^{U \times V \times 2}$, which has a spatial size of $U \times V$ and contains the horizontal and vertical positions. This alignment can be formulated as the optimization problem below:

$$M_{opt}^f = \min_{M^f} P(I_{ref}, \text{Warp}(I_{tar}, M^f)), \quad (1)$$

where $P(\cdot, \cdot)$ evaluates the alignment error and $\text{Warp}(\cdot, \cdot)$ performs differentiable image warping.

The dense cross-scale alignment network is introduced to address the optimization problem more effectively. Instead of regressing the mesh directly, it predicts position offsets. As illustrated in Figure 1(a), the network employs a coarse-to-fine strategy. Specifically, it extracts fine and coarse feature maps from both reference and target images to regress global offsets $O_g \in \mathbb{R}^{4 \times 2}$ and local offsets $O_l \in \mathbb{R}^{U \times V \times 2}$, respectively. The global offsets O_g are used to apply a homography transformation on a regular mesh $M \in \mathbb{R}^{U \times V \times 2}$ to ensure global consistency. Subsequently, O_l is added to the transformation result to generate M^f for fine-grained warping. Our method leverages cross-scale information from fine feature maps to enhance local deformation. A fully spatial correlation module is designed to assist in predicting both global and local offsets. Additionally, just noticeable difference (JND) guidance is incorporated to reduce the alignment errors in visually sensitive regions.

Dense Cross-Scale Regression

Large parallax results in scale variations for the same object across the reference and target images. Existing methods rely solely on intra-scale features and fail to accurately establish their relationships. This limits alignment performance and motivates the development of the dense cross-scale regression module to alleviate alignment difficulty. The cross-scale information is incorporated into the local offset O_l , which is expressed as

$$O_l = O_l^{intra} + O_l^{cross}, \quad (2)$$

where O_l^{intra} and $O_l^{cross} \in \mathbb{R}^{U \times V \times 2}$ are the offsets calculated from the intra-scale and cross-scale features, respectively. Notably, intra-scale information is integrated at both coarse and fine stages. The proposed module predicts O_l^{cross} in a residual manner by leveraging diverse cross-scale information captured from the fine feature maps.

Figure 1(b) illustrates the structure of the dense cross-scale regression module. Let F_{ref}^l and $F_{tar}^l \in \mathbb{R}^{c \times h \times w}$ denote the fine features extracted from the reference and target images, respectively. The homography transformation is applied to F_{tar}^l using the global offset O_g to obtain the pre-aligned version F_{tar}^w . The cross-scale regression module takes F_{ref}^l and F_{tar}^w as inputs and first downsamples them. This downsampling operation is implemented using max pooling for efficiency while preserving important information. This process is repeated N times, reducing the spatial size by half with each iteration. The results are expressed as:

$$\begin{aligned} F_i^l &= \text{MaxPool}_i(\dots \text{MaxPool}_1(F_{ref}^l) \dots), \\ F_i^w &= \text{MaxPool}_i(\dots \text{MaxPool}_1(F_{tar}^w) \dots), \end{aligned} \quad (3)$$

where $\text{MaxPool}_i(\cdot)$ represents the i -th max pooling operation and i belongs to the positive integer set $\{1, 2, \dots, N\}$; F_i^l and F_i^w are of size $\mathbb{R}^{c \times \frac{h}{2^i} \times \frac{w}{2^i}}$. For simplicity, we set F_0^l and F_0^w to F_{ref}^l and F_{tar}^w , respectively. The cross-scale regression module establishes dense cross-scale relationships between the above features of the reference and target images. In this way, it fully utilizes the cross-scale information to enhance alignment accuracy. Specifically, a set of offsets is calculated based on F_m^l and F_n^w for $\forall m, n \in \{0, 1, \dots, N\}$ and $m \neq n$. In this way, we can obtain $N^2 + N$ outputs. These offsets are aggregated to obtain O_l^{cross} . This process can be expressed as

$$O_l^{cross} = \sum_{m, n \in [N] \text{ and } m \neq n} \text{Net}(F_m^l, F_n^w), \quad (4)$$

where $\text{Net}(\cdot, \cdot)$ regresses local offsets from the input features and is implemented by the proposed fully spatial correlation module; $[N]$ denotes the set of natural numbers $\{0, 1, \dots, N\}$. Only intra-scale features are used for alignment when N is 0.

Fully Spatial Correlation

Let $F_{ref} \in \mathbb{R}^{c \times h_1 \times w_1}$ and $F_{tar} \in \mathbb{R}^{c \times h_2 \times w_2}$ represent the feature maps extracted from the reference and target images, respectively. A correlation tensor is computed from

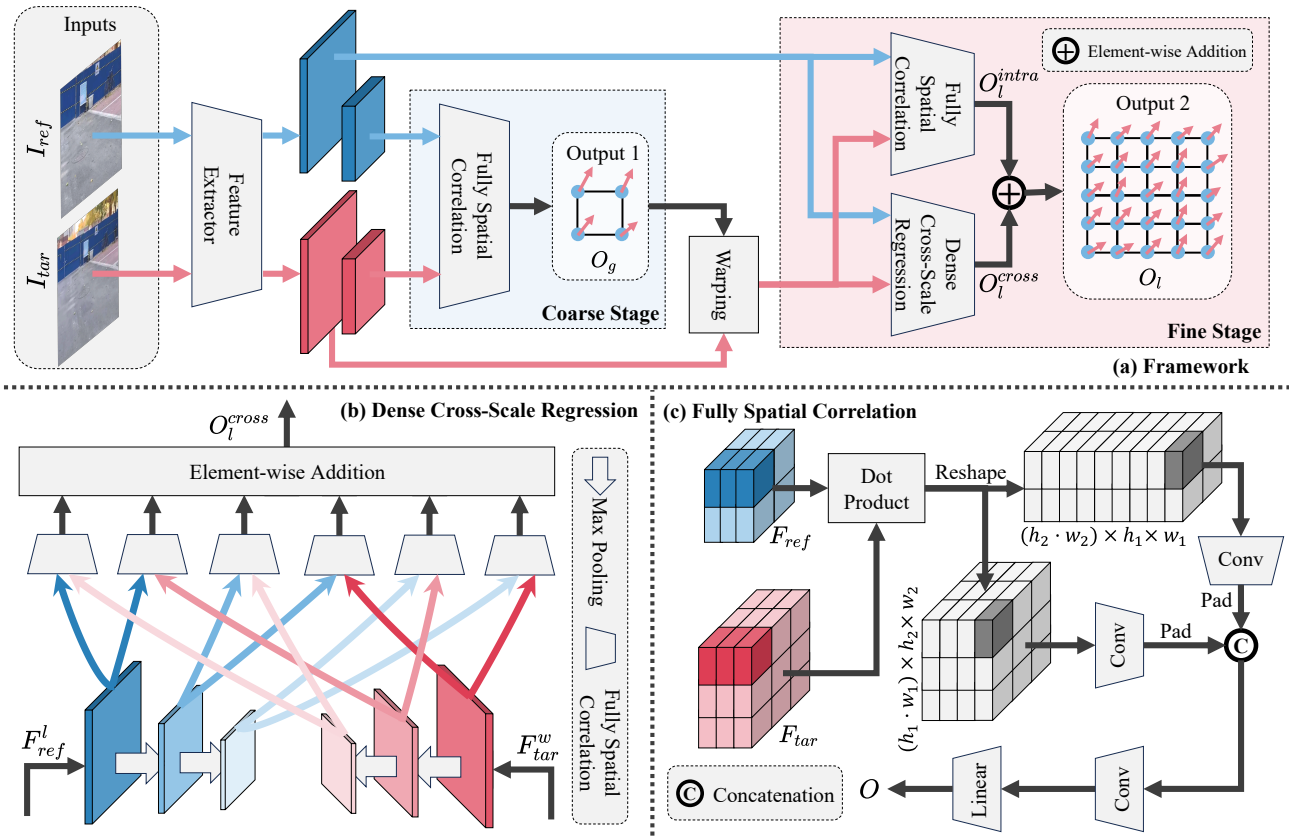


Figure 1: Illustration of the proposed alignment network. (a) The overall framework. (b) The proposed cross-scale regression. (c) The proposed fully spatial correlation. Our method uses a coarse-to-fine strategy to estimate the offsets of a mesh. The dense cross-scale regression module integrates the cross-scale information into the local offsets. The fully spatial correlation module utilizes the spatial information of both input features. The dot product is performed between the dark blue and dark red feature vectors to obtain the element in dark gray.

F_{ref} and F_{tar} to regress offsets. The contextual correlation layer (CCL) constructs this tensor by decomposing F_{tar} into $h_2 \cdot w_2$ overlapping kernels and convolving them with F_{ref} . The cost volume (CV) performs a Hadamard product on F_{tar} and F_{ref} using a sliding window approach. Both methods capture contextual information. However, the dense cross-scale regression module also establishes long-range relationships. It diminishes the advantages of the CCL and CV. Moreover, the large value of $h_2 \cdot w_2$ results in numerous convolution operations, and the Hadamard product is computed repeatedly. They lead to high FLOPs and long runtime, respectively. Our dense cross-scale regression module requires $N^2 + N$ such operations. The computational overhead is significant when the CCL and CV are employed.

To further improve effectiveness while maintaining low computational costs, we propose the fully spatial correlation module. Figure 1(c) illustrates its architecture. Specifically, the module computes a 4D correlation tensor $T \in \mathbb{R}^{h_2 \times w_2 \times h_1 \times w_1}$ to encode pairwise correlations through dot products between the feature vectors in F_{ref} and F_{tar} . This tensor T must be processed through several convolutional and linear layers to regress offsets. To facilitate this, the correlation layer reshapes the 4D tensor T into a 3D tensor

$T_1 \in \mathbb{R}^{(h_2 \cdot w_2) \times h_1 \times w_1}$. Nevertheless, this reshaping operation destroys the spatial information of F_{tar} . Convolutional layers can only utilize the spatial information of F_{ref} , limiting alignment accuracy. To fully leverage the spatial information from both feature maps, our module reshapes T to construct another tensor $T_2 \in \mathbb{R}^{(h_1 \cdot w_1) \times h_2 \times w_2}$. Both T_1 and T_2 are processed independently by convolutional layers for feature compression and refinement. Given the large dimensions of $h_1 \cdot w_1$ and $h_2 \cdot w_2$, we reduce the channel numbers to a more manageable size to enhance efficiency. The resulting outputs are then padded to match spatial dimensions for concatenation along the channel axis. The concatenation result is given by

$$T' = \text{Cat}(\text{Pad}(\text{Conv}(T_1)), \text{Pad}(\text{Conv}(T_2))), \quad (5)$$

where $\text{Cat}(\cdot, \cdot)$, $\text{Pad}(\cdot)$, and $\text{Conv}(\cdot)$ represent the concatenation operation, padding operation, and convolution layers, respectively. Finally, T' undergoes further processing through a series of convolutional layers and linear layers to regress the offsets:

$$O = \text{Linear}(\dots(\text{Linear}((\text{Conv}(\dots\text{Conv}(T')\dots))\dots)), \quad (6)$$

where $\text{Linear}(\cdot)$ denotes the linear layer. In this way, the

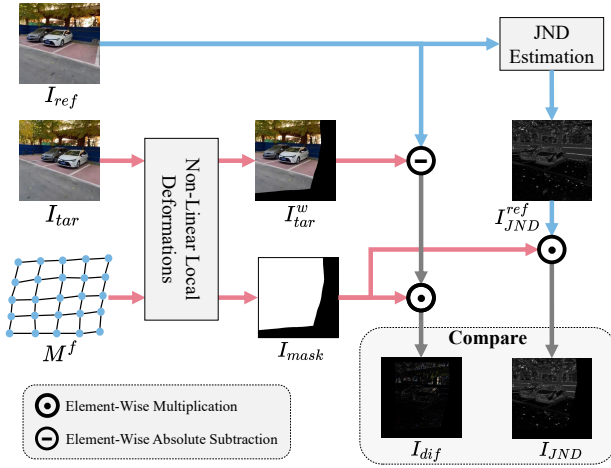


Figure 2: Illustration of the proposed JND guidance. The JND map is estimated for the reference image. The difference is calculated for the overlapping area to compare with the JND map to update network parameters.

proposed module can better utilize the spatial information from F_{ref} and F_{tar} with low computational costs.

Just Noticeable Distortion Guidance

As shown in Eq. (1), an alignment model aims to find an optimal mesh that minimizes alignment errors. With the optimal mesh M_{opt}^f and the one M^f predicted by a model, we obtain the corresponding warped target images as follows:

$$I_{opt}^w = \text{Warp}(I_{tar}, M_{opt}^f), \quad (7)$$

$$I_{tar}^w = \text{Warp}(I_{tar}, M^f) = I_{opt}^w + I_n, \quad (8)$$

where I_n denotes the noise introduced by M^f . Consequently, the optimization problem in Eq. (1) can be reframed as minimizing I_n . Alignment accuracy is quantified by the similarity between the overlapping regions of the reference and warped target images. The noise I_n degrades I_{opt}^w to reduce this similarity, leading to alignment errors. However, the impact of I_n on similarity varies across different instances, even at the same noise level, due to the characteristics of the human visual system (Jiang et al. 2022a). Distortions are more readily perceived in certain regions of an image, resulting in greater degradation of similarity in those areas. Inspired by this insight, we leverage the JND to further enhance alignment accuracy. The JND serves as an indicator of distortion tolerance for each pixel. A pixel change is imperceptible when it falls below the threshold value specified by the JND. By utilizing these threshold values, we encourage our method to focus on image regions that are more sensitive to distortions. Figure 2 illustrates the proposed JND guidance strategy. It aims to ensure that distortions in the overlapping areas between the reference and warped target images remain below the JND thresholds.

We adopt the methodology from (Wu et al. 2017) to estimate the JND map I_{JND}^{ref} for the reference image. Detailed

estimation procedures are provided in the subsection ‘‘JND Estimation’’ of the supplementary material. Next, we calculate the difference between the reference image I_{ref} and the warped target image I_{tar}^w within the overlapping areas. A pixel’s distortion is considered for parameter updates only if it exceeds the corresponding threshold in I_{JND}^{ref} . This consideration can be expressed as:

$$I(x) = \begin{cases} 0 & \text{if } I_{dif}(x) \leq I_{JND}(x) \\ I_{dif}(x) - I_{JND}(x) & \text{if } I_{dif}(x) > I_{JND}(x) \end{cases}, \quad (9)$$

$$I_{dif} = |I_{ref} - I_{tar}^w| \odot I_{mask}, \quad (10)$$

$$I_{JND} = I_{JND}^{ref} \odot I_{mask}, \quad (11)$$

where x denotes a pixel position; \odot and $|\cdot|$ represent the Hadamard product and absolute value operation, respectively; I_{mask} is a binary mask indicating overlapping and non-overlapping areas with values of one and zero, respectively. When updating the parameters of the alignment model, our goal is to minimize each element in I as much as possible. We incorporate I into the loss function to guide the parameter update. The loss function of the JND is defined as

$$L_{JND} = \text{Mean}(\text{ReLU}(I_{dif} - I_{JND})), \quad (12)$$

where $\text{Mean}(\cdot)$ and $\text{ReLU}(\cdot)$ represent the average operation and rectified linear unit, respectively.

Loss Function

The loss function is designed to optimize three critical objectives in image alignment tasks. First, it enforces precise spatial correspondence by maximizing image similarity between overlapping regions of the reference and warped target images. Second, it maintains the visual plausibility of the warped target image through naturalness constraints that preserve realistic structures. Third, it incorporates perceptual regularization to ensure that the introduced distortions remain below the threshold of human visual perception as possible. These requirements are mathematically formalized through three complementary loss components as follows:

$$L = L_{content} + \alpha L_{shape} + \beta L_{JND}, \quad (13)$$

where $L_{content}$ and L_{shape} are the content and shape preservation losses employed in (Nie et al. 2023) and (Nie et al. 2022a); L_{JND} is the distortion loss given in Eq. (12). The details of $L_{content}$ and L_{shape} are shown in the subsection ‘‘Loss Function’’ of the supplementary material.

Experiments

Experimental Settings

Datasets. Following existing unsupervised methods, the UDIS-D dataset (Nie et al. 2021) is used to train and assess our image alignment model. It is the only large-scale public dataset that contains real-world image pairs with parallax. No ground truth is provided. There are 10,440 image pairs for training and 1,106 pairs for testing, with each image sized 512×512 . The small traditional dataset (Liao and Li 2020) is also used to perform cross-dataset validation.

Method	PSNR \uparrow				SSIM \uparrow				FLOPs \downarrow	Time \downarrow
	Easy	Moderate	Hard	Average	Easy	Moderate	Hard	Average		
$N = 0$	30.46	26.04	21.74	25.63	0.9363	0.8805	0.7453	0.8426	443.1 G	0.0174
$N = 1$	30.88	26.58	22.26	26.13	0.9405	0.8915	0.7653	0.8552	527.1 G	0.0246
$N = 2$	30.95	26.62	22.29	26.18	0.9408	0.8918	0.7657	0.8556	562.9 G	0.0305
$N = 2$ w/ CCL	30.72	26.48	22.16	26.01	0.9381	0.8882	0.7588	0.8509	1396.8 G	0.0315
$N = 2$ w/ CL	30.80	26.46	22.10	26.01	0.9394	0.8887	0.7591	0.8516	556.3 G	0.0285
$N = 2$ w/ CV	—	—	—	—	—	—	—	—	431.9 G	0.5580
$N = 2$ w/o JND	30.82	26.50	22.12	26.03	0.9403	0.8895	0.7599	0.8524	562.9 G	0.0305

Table 1: Ablation studies on each proposed module. The fully spatial correlation and JND guidance are employed by default. N denotes the number of scales used to predict local offsets. CCL, CL, and CV represent the contextual correlation layer, correlation layer, and cost volume, respectively. From “Easy” to “Hard”, the parallax becomes larger.

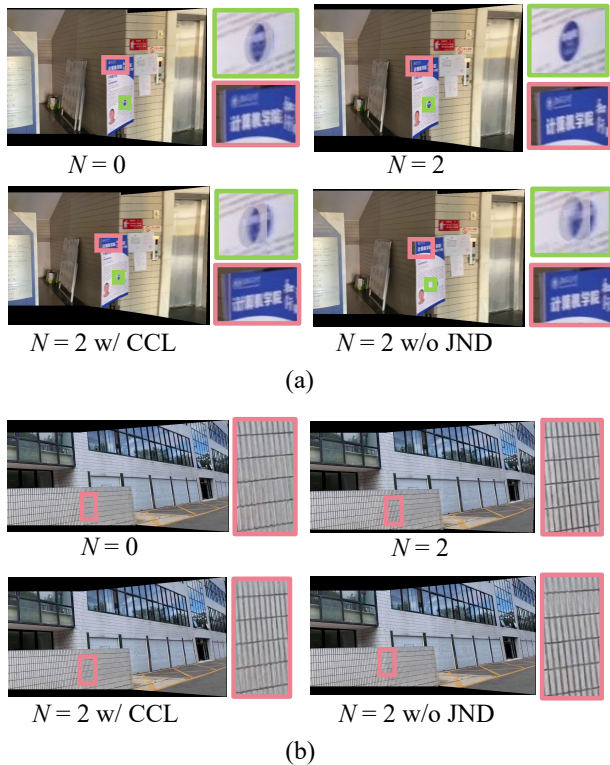


Figure 3: Ablation studies on our method. The red and green boxes zoom in on the regions with alignment errors.

Implementations. Our alignment network is implemented in PyTorch and trained for 150 epochs on a single NVIDIA RTX A100 GPU. We use the Adam optimizer for parameter updates, with coefficients set to 0.9 and 0.999. The ResNet (He et al. 2016) is employed as the feature extractor. The learning rate, batch size, mesh size, α , β , and N are set to 10^{-4} , 4, 13×13 , 10, 1, and 2, respectively.

Metrics. Due to the lack of ground truth, existing unsupervised methods measure the accuracy of image alignment by the similarity between overlapping regions of the reference and warped target images. The PSNR and SSIM are two commonly used metrics for this evaluation.

Ablation Studies

We conduct ablation studies to validate the effectiveness of each proposed module. Table 1 shows the alignment accuracy and computational costs.

Cross-Scale Regression. Our method with $N = 0$ only uses the intra-scale information for alignment, while those with $N > 0$ incorporate the cross-scale information. Parallax causes the scale variation problem. As shown in Table 1, alignment accuracy improves across different parallax scenarios as N increases. These results underscore the importance of the cross-scale information in addressing the variation problem. The computational costs also increase with increasing N . The additional cross-scale information reduces efficiency. Thus, our method allows for various trade-offs between accuracy and computational costs by adjusting N . Notably, the accuracy gains become marginal when increasing N from 1 to 2. This leads us to set N to 2.

Fully Spatial Correlation. According to Table 1, replacing the fully spatial correlation module with the CCL and CL results in decreased alignment accuracy. One advantage of the CCL is its ability to capture contextual information through convolution operations. The dense cross-scale module can utilize cross-scale features to establish long-range relationships. This enables the CL to achieve performance comparable to the CCL. The fully spatial correlation module can better facilitate offset regression through its flexible spatial information integration ability. Thus, it outperforms the CCL and CL. Furthermore, our method with $N = 1$ has higher alignment accuracy than the CCL and CL with $N = 2$. In this case, the CCL and CL incur longer inference times and higher FLOPs. Especially, the CCL has approximately 2.5 times the FLOPs of our method. While parallel computation techniques can help reduce inference time, they do not alleviate the increased energy demands associated with higher FLOPs. Besides, the varying spatial sizes of cross-scale features render the CV unsuitable for our approach, as it necessitates input features with consistent spatial dimensions. Although the CV has lower FLOPs, it exhibits nearly 20 times the runtime compared to our method. Overall, our approach achieves superior alignment accuracy while maintaining low computational costs.

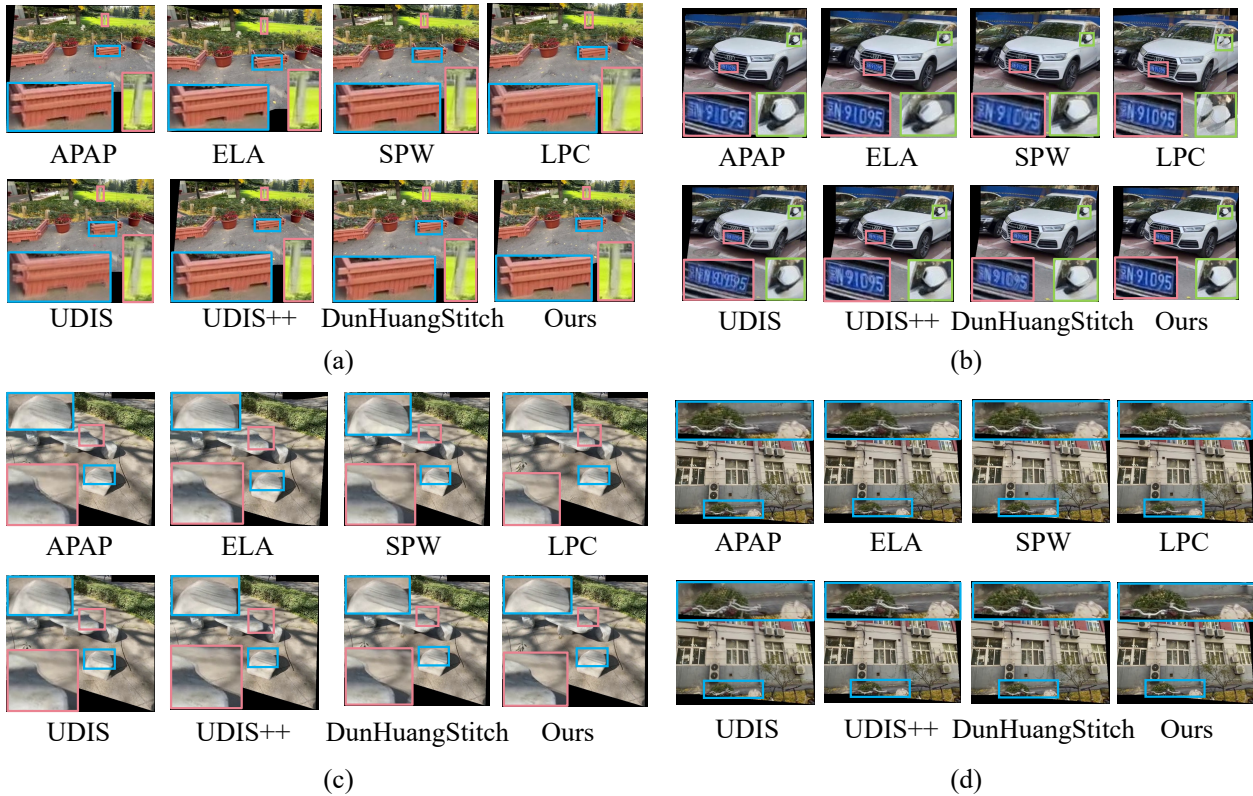


Figure 4: Visual comparisons of different alignment methods on the UDIS-D dataset (Nie et al. 2021). The colorful boxes zoom in on the area with alignment errors.

JND Guidance. The alignment accuracy is further enhanced under the guidance of the JND. It makes the alignment errors less perceptible in line with the human visual system. Besides, the JND guidance does not increase inference time and FLOPs, as it is used only for model training.

Visualization. Figure 3 presents the visualization results. To highlight alignment errors, we apply the average fusion to the reference and warped target images. In Figure 3(a), our method with $N = 2$ exhibits superior alignment performance compared to the other three methods. These three methods introduce distortions in the text, board contour, and blue circle. In Figure 3(b), the method with $N = 2$ does not produce obvious artifacts. In contrast, the alignment errors cause blurring: at the top for the method with the CCL, at the bottom for the method without JND guidance, and throughout the whole area for the method with $N = 0$.

Comparisons with Baselines

Our method is compared with state-of-the-art unsupervised image alignment methods and homography estimation methods. The baselines include APAP (Zaragoza et al. 2013), ELA (Li et al. 2018), SPW (Liao and Li 2020), LPC (Jia et al. 2021), UDIS (Nie et al. 2021), DAMG (Nie et al. 2022b), UDIS++ (Nie et al. 2023), DunHuangStitch (Mei et al. 2024), IHN (Cao et al. 2022), RHWF (Cao et al. 2023) and MCNet (Zhu et al. 2024).




Figure 5: Validation results on the cross-dataset image “19” (Liao and Li 2020). The red box zooms in on the area with alignment errors.

Quantitative Results. Table 2 reports the alignment accuracy of different methods. Compared to the baselines, our method improves the PSNR and SSIM values by at least 0.74/0.007, 0.79/0.016, and 0.28/0.025 for easy, moderate, and hard parallax scenarios, respectively. The average improvements in PSNR and SSIM are 0.75 and 0.018, respectively. Parallax causes the scale variation problem for the same object across the image pair. Our method outperforms all the baselines in various parallax scenarios. It can better

Method	PSNR \uparrow				SSIM \uparrow			
	Easy	Moderate	Hard	Average	Easy	Moderate	Hard	Average
$I_{3\times 3}$	15.87	12.76	10.68	12.86	0.530	0.286	0.146	0.303
APAP	27.99	24.29	20.17	23.74	0.899	0.832	0.683	0.792
ELA	29.74	25.36	19.47	24.30	0.923	0.862	0.704	0.816
SPW	28.15	23.39	19.86	23.32	0.907	0.799	0.631	0.764
LPC	27.26	22.87	19.38	22.78	0.883	0.776	0.623	0.747
IHN	22.16	18.17	14.53	17.90	0.747	0.558	0.391	0.547
RHWF	20.10	16.42	13.51	16.35	0.688	0.481	0.331	0.483
MCNet	22.52	18.68	15.45	18.53	0.768	0.591	0.422	0.576
UDIS	25.16	20.96	18.36	21.17	0.834	0.669	0.495	0.648
DAMG	29.52	25.24	21.20	24.89	0.923	0.859	0.708	0.817
UDIS++	<u>30.21</u>	<u>25.83</u>	<u>21.60</u>	<u>25.43</u>	<u>0.934</u>	<u>0.876</u>	<u>0.739</u>	<u>0.838</u>
DunHuangStitch	29.47	25.60	<u>22.01</u>	25.31	0.921	0.859	0.741	0.830
Ours	30.95	26.62	22.29	26.18	0.941	0.892	0.766	0.856

Table 2: Alignment accuracy comparisons of different methods on the UDIS-D dataset (Nie et al. 2021). The best and second-best results are written in bold and underlined, respectively. N denotes the number of scales used by our method. $I_{3\times 3}$ indicates that no alignment is performed. From “Easy” to “Hard”, the parallax becomes larger.

address the scale variation problem.

Qualitative Results. Figure 4 displays the visualization results of different methods on the UDIS-D dataset. The average fusion is applied to highlight alignment errors. In Figure 4(a), the planter outlined in blue shows misalignment across all competing methods. ELA, UDIS++, and DunHuangStitching exhibit less blurring at the top but more at the bottom. Other baselines demonstrate artifacts across the whole planter. The competing methods also cause noticeable distortions to the pole outlined in red. In Figure 4(b), significant misalignment is observed in the license plate, except for ELA, LPC, and our method. However, ELA and LPC display artifacts in the side mirror. It is also observed from Figures 4(c) and 4(d) that misalignment occurs in the areas of the stone table, the stone chair, and the handlebar crossbar for the competing methods. The scale of the stone chair and desk in Figure 4(c) differs across the reference and target images. Our method achieves more plausible visual results and thus can better address the scale variation problem.

Cross-Dataset Validation. We also evaluate the generalization of the proposed method on cross-dataset images (Liao and Li 2020). Following (Nie et al. 2023), the iterative optimization strategy is adopted. Figures 5 and 6 present the visualization results. In Figure 5, noticeable artifacts appear at the top for ELA, SPW, and LPC, as well as at the bottom for UDIS++ and DunHuangStitch. Figure 6 reveals evident misalignment in areas with bicycles for SPW, LPC, and DunHuangStitch. All competing methods exhibit significant artifacts on the left side of the window fence. We also perform user studies to compare our method with UDIS++ and DunHuangStitch. Details are provided in the subsection “Cross-Dataset Validation” of the supplementary material. Feedback indicates that our method produces more visually plausible results than the two competing approaches.

Runtime. We show the runtime for images of varying resolutions in the subsection “Runtime Comparison” of the supplementary material. Our method is competitive.




Figure 6: Validation results on the cross-dataset image “40” (Liao and Li 2020). The red and blue boxes zoom in on the area with alignment errors.

Conclusion

In this paper, we introduced a dense cross-scale image alignment method. It captured the cross-scale features to increase image alignment. The number of scales can be adjusted to balance accuracy and computational costs. A fully spatial module was designed to further increase accuracy while maintaining low computational costs. We also exploited the JND to encourage the alignment model to focus on regions sensitive to distortions, thus improving performance. Ablation studies validated the effectiveness of the proposed modules. Extensive experiments were conducted to demonstrate the superiority of the proposed method over the state-of-the-art approaches in both effectiveness and efficiency.

Acknowledgements

This work was funded by the Science and Technology Development Fund, Macau SAR (File no. 0049/2022/A1, 0050/2024/AGJ), by the University of Macau Development Foundation (File no. MYRG-GRG2024-00181-FST-UMDF).

References

- Cao, S.-Y.; Hu, J.; Sheng, Z.; and Shen, H.-L. 2022. Iterative deep homography estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 1879–1888.
- Cao, S.-Y.; Zhang, R.; Luo, L.; Yu, B.; Sheng, Z.; Li, J.; and Shen, H.-L. 2023. Recurrent homography estimation using homography-guided image warping and focus transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9833–9842.
- Deng, Z.; Luo, L.; and Chen, H. 2024. Enable the right to be forgotten with federated client unlearning in medical imaging. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 240–250. Springer.
- DeTone, D.; Malisiewicz, T.; and Rabinovich, A. 2016. Deep image homography estimation. *arXiv preprint arXiv:1606.03798*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Jia, Q.; Feng, X.; Liu, Y.; Fan, X.; and Latecki, L. J. 2023. Learning pixel-wise alignment for unsupervised image stitching. In *Proceedings of the 31st ACM International Conference on Multimedia*, 1392–1400.
- Jia, Q.; Li, Z.; Fan, X.; Zhao, H.; Teng, S.; Ye, X.; and Latecki, L. J. 2021. Leveraging Line-point Consistency to Preserve Structures for Wide Parallax Image Stitching. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 12181–12190.
- Jiang, Q.; Liu, Z.; Wang, S.; Shao, F.; and Lin, W. 2022a. Toward Top-Down Just Noticeable Difference Estimation of Natural Images. *IEEE Transactions on Image Processing*, 31: 3697–3712.
- Jiang, Z.; Zhang, Z.; Fan, X.; and Liu, R. 2022b. Towards all weather and unobstructed multi-spectral image stitching: Algorithm and benchmark. In *Proceedings of the 30th ACM international conference on multimedia*, 3783–3791.
- Kim, S.; Hong, S.; Park, K.; Cho, H.; and Sohn, K.-A. 2024. OmniStitch: Depth-Aware Stitching Framework for Omnidirectional Vision with Multiple Cameras. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 10210–10219.
- Kumar, V. R.; Eising, C.; Witt, C.; and Yogamani, S. K. 2023. Surround-view fisheye camera perception for automated driving: Overview, survey & challenges. *IEEE Transactions on Intelligent Transportation Systems*, 24(4): 3638–3659.
- Li, J.; Wang, Z.; Lai, S.; Zhai, Y.; and Zhang, M. 2018. Parallax-Tolerant Image Stitching Based on Robust Elastic Warping. *IEEE Transactions on Multimedia*, 20(7): 1672–1687.
- Li, J.; and Zhou, Y. 2022. Automatic color image stitching using quaternion rank-1 alignment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19720–19729.
- Li, J.; and Zhou, Y. 2024. Automatic Quaternion-Domain Color Image Stitching. *IEEE Transactions on Image Processing*, 33: 1299–1312.
- Li, J.; and Zhou, Y. 2025. Seam-Adaptive Structure-Preserving Image Stitching for Drone Images. *IEEE Transactions on Geoscience and Remote Sensing*, 63: 1–12.
- Li, Y.; Tofighi, M.; and Monga, V. 2019. Robust Alignment for Panoramic Stitching Via an Exact Rank Constraint. *IEEE Transactions on Image Processing*, 28(10): 4730–4745.
- Liao, K.; Nie, L.; Huang, S.; Lin, C.; Zhang, J.; Zhao, Y.; Gabbouj, M.; and Tao, D. 2023. Deep learning for camera calibration and beyond: A survey. *arXiv preprint arXiv:2303.10559*.
- Liao, T.; and Li, N. 2020. Single-Perspective Warps in Natural Image Stitching. *IEEE Transactions on Image Processing*, 29: 724–735.
- Lin, K.; Jiang, N.; Cheong, L.-F.; Do, M.; and Lu, J. 2016. Seagull: Seam-guided local alignment for parallax-tolerant image stitching. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part III 14*, 370–385. Springer.
- Mei, Y.; Yang, L.; Wang, M.; Yu, T.; and Wu, K. 2024. DunHuangStitch: Unsupervised Deep Image Stitching of Dunhuang Murals. *IEEE Transactions on Visualization and Computer Graphics*, 1–17.
- Nguyen, T.; Chen, S. W.; Shivakumar, S. S.; Taylor, C. J.; and Kumar, V. 2018. Unsupervised deep homography: A fast and robust homography estimation model. *IEEE Robotics and Automation Letters*, 3(3): 2346–2353.
- Nie, L.; Lin, C.; Liao, K.; Liu, S.; and Zhao, Y. 2021. Unsupervised Deep Image Stitching: Reconstructing Stitched Features to Images. *IEEE Transactions on Image Processing*, 30: 6184–6197.
- Nie, L.; Lin, C.; Liao, K.; Liu, S.; and Zhao, Y. 2022a. Deep rectangling for image stitching: A learning baseline. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5740–5748.
- Nie, L.; Lin, C.; Liao, K.; Liu, S.; and Zhao, Y. 2022b. Depth-Aware Multi-Grid Deep Homography Estimation With Contextual Correlation. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(7): 4460–4472.
- Nie, L.; Lin, C.; Liao, K.; Liu, S.; and Zhao, Y. 2023. Parallax-Tolerant Unsupervised Deep Image Stitching. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, 7365–7374.
- Ou, F.-Z.; Li, C.; Wang, S.; and Kwong, S. 2024. CLIB-FIQA: Face Image Quality Assessment with Confidence Calibration. In *Proceedings of the IEEE/CVF Conference on*

Computer Vision and Pattern Recognition (CVPR), 1694–1704.

Rocco, I.; Arandjelovic, R.; and Sivic, J. 2017. Convolutional neural network architecture for geometric matching. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 6148–6157.

Wu, J.; Li, L.; Dong, W.; Shi, G.; Lin, W.; and Kuo, C.-C. J. 2017. Enhanced Just Noticeable Difference Model for Images With Pattern Complexity. *IEEE Transactions on Image Processing*, 26(6): 2682–2693.

You, J.; Zhang, J.; and Zhou, Y. 2025. Hffnet: Hierarchical Feature Fusion Network for Thermal Image Super-Resolution. *Available at SSRN 5372220*.

You, J.; and Zhou, Y. 2024. Two-Stage Watermark Removal Framework for Spread Spectrum Watermarking. *IEEE Transactions on Multimedia*, 26: 7687–7699.

Zaragoza, J.; Chin, T.-J.; Brown, M. S.; and Suter, D. 2013. As-Projective-As-Possible Image Stitching with Moving DLT. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Zhao, Z.; Bai, H.; Zhang, J.; Zhang, Y.; Xu, S.; Lin, Z.; Timofte, R.; and Van Gool, L. 2023. Cddfuse: Correlation-driven dual-branch feature decomposition for multi-modality image fusion. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5906–5916.

Zhu, H.; Cao, S.-Y.; Hu, J.; Zuo, S.; Yu, B.; Ying, J.; Li, J.; and Shen, H.-L. 2024. Mcnet: Rethinking the core ingredients for accurate and efficient homography estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 25932–25941.