

RecStitchNet: Learning to Stitch Images with Rectangular Boundaries

Yun Zhang

Communication University of Zhejiang

zhangyun@cuz.edu.cn

Lang Nie

Beijing Jiaotong University

nielang@bjtu.edu.cn

Yu-Kun Lai

Cardiff University

Yukun.Lai@cs.cardiff.ac.uk

Fang-Lue Zhang

Victoria University of Wellington

fanglue.zhang@ecs.vuw.ac.nz

Lin Xu

University of South Australia

xuyly032@mymail.unisa.edu.au

Abstract

Irregular boundaries in image stitching naturally occur due to freely moving cameras. To deal with this problem, existing methods focus on optimizing mesh warping to make boundaries regular using the traditional explicit solution. However, previous methods always depend on hand-crafted features (e.g. keypoints and line segments). Thus, failures often happen in overlapping regions without distinctive features. In this paper, we address this problem by proposing *RecStitchNet*, a reasonable and effective network for image stitching with rectangular boundaries. Considering that both stitching and imposing rectangularity are non-trivial tasks in the learning-based framework, we propose a three-step progressive learning based strategy, which not only simplifies this task, but gradually achieves a good balance between stitching and imposing rectangularity. In the first step, we perform initial stitching by a pre-trained state-of-the-art image stitching model, to produce initially warped stitching results without considering the boundary constraint. Then, we use a regression network with a comprehensive objective regarding mesh, perception, and shape to further encourage the stitched meshes to have rectangular boundaries with high content fidelity. Finally, we propose an unsupervised instance-wise optimization strategy to refine the stitched meshes iteratively, which can effectively improve the stitching results in terms of feature alignment, as well as boundary and structure preservation. Due to the lack of stitching datasets and the difficulty of label generation, we propose to generate a stitching dataset with rectangular stitched images as pseudo-ground-truth labels, and the performance upper bound induced from the it can be broken by our

unsupervised refinement. Qualitative and quantitative results and evaluations demonstrate the advantages of our method over the state-of-the-art.

Keywords: image stitching, boundaries, convolutional neural network

1. Introduction

In recent decades, image stitching has been an active topic in computer graphics and vision. The goal of image stitching is to construct a wide field-of-view (FOV) scene from several overlapping images each having a limited FOV. This has a wide range of applications in virtual reality, autonomous driving, video surveillance, etc. Traditional image stitching methods mainly focus on accurate feature matching, natural warping, shape- and straight line-preservation [3, 5, 17]. Despite their great success, most of these methods rely on the performance of hand-crafted feature matching in overlapping regions, and thus have limited generalizability. These methods often struggle to stitch images with unclear textures, or taken in low light, or having low resolution. Additionally, preserving geometric structure and visual features necessitates complex optimization and intensive computation, further heightening the difficulty of image stitching.

To overcome the challenges posed by feature matching and structure preservation, learning-based methods have been extensively studied in recent years; they stitch images by adaptively learning high-level semantic features from big data. These methods can be roughly divided into three types: supervised [22, 38, 11], weakly-supervised [30], and unsupervised [23, 26] methods. They are able to robustly and efficiently stitch images, demonstrating high performance in terms of large parallax tolerance and geometry preservation. However, most of them do not take bound-

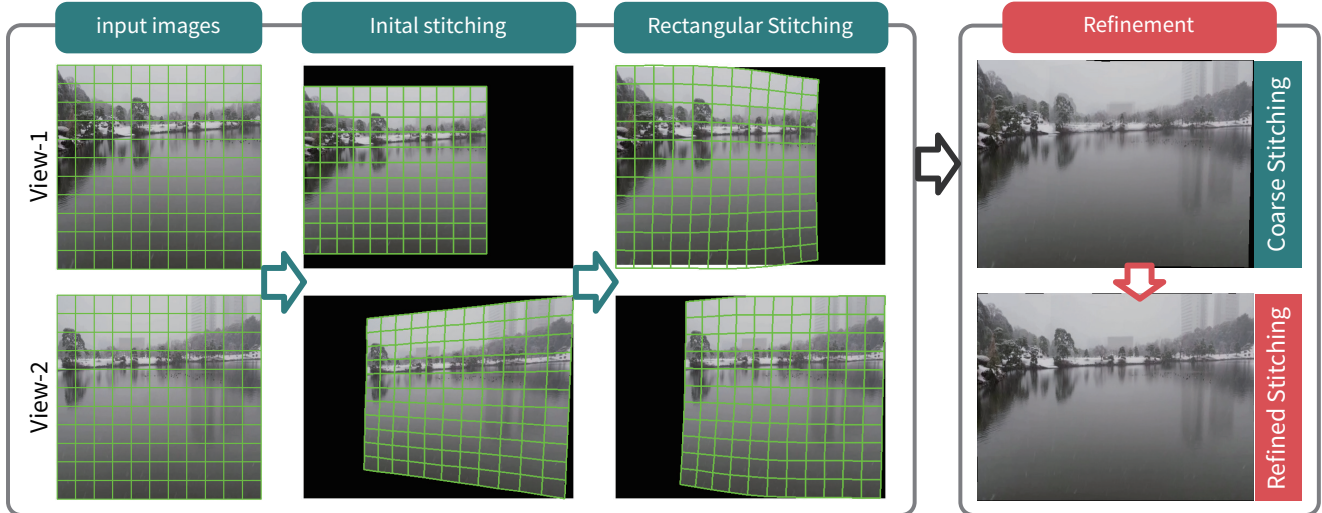


Figure 1. Pipeline of our method. We take two normal FOV images as input, and then stitch them using a pre-trained model. Taking the initial stitching result with irregular boundary as input, we use *RecStitchNet* to produce coarse stitching result with a rectangular boundary. We finally refine the coarse stitching result using an instance-wise unsupervised learning method.

ary regularity into consideration. Recently, following previous work on image stitching and imposing rectangularity of results based on conventional optimization frameworks [8, 37], Nie *et al.* [24] proposed the first deep learning solution for imposing image rectangularity, which was further extended to image rotation correction in [25]. They took well-stitched images as input and learned to rectify the irregular boundaries while preserving the high-level semantic features. However, their method does not consider optimizing stitching simultaneously with creation of rectangular image boundaries. This oversight could potentially amplify artifacts in the stitched input after applying warping-based rectification.

In this paper, we introduce *RecStitchNet*, a supervised learning network designed for stitching images while ensuring rectangular stitching boundaries. To enable an effective learning process, we have designed a three-step progressive stitching approach. Firstly, we conduct an initial stitching process using a state-of-the-art deep stitching technique, to give warped meshes for the image pair. Secondly, we use a regression network of our own design with a comprehensive objective regarding mesh, perception, and shape to encourage the stitched meshes to have rectangular boundaries with high content fidelity. The output of the network is the predicted mesh motions relative to the initially warped meshes. In this paper, the term ‘mesh motion’ refers to offsets of all vertex positions of the mesh on each image. Finally, to ensure the robustness of our method across various scenarios, we employ an unsupervised instance-wise network to improve the stitching result. This refinement process is guided by an objective function comprising a rectangular boundary term, a feature-matching term and a shape preser-

vation term, which collectively contribute to the production of high-quality stitching results.

Unlike typical stitching methods that often result in irregular boundaries, our objective is to achieve stitching results with rectangular boundaries. Generally, both stitching and imposing rectangularity are challenging tasks, necessitating a supervised network for effective learning. However, obtaining ground truth stitching results is difficult due to the absence of a publicly recognized standard. So that rectangular stitching results can be more standardized and universally recognized, we propose to generate pseudo-ground-truth using a state-of-the-art stitching technique with rectangular boundaries [37].

Fig. 1 shows the pipeline of our method. Given two normal FOV images as input, the proposed solution progressively stitches them from an initial stitched image with irregular boundaries to a coarse stitched image with rectangular boundaries to the final stitching result with further refinement of boundaries and alignment. Extensive experiments and evaluations in this paper show that our approach can effectively stitch images and obtain satisfactory results, with rectangular boundaries.

Compared to a previous stitching method imposing rectangularity [37], our method is more robust and efficient due to its effective high-level feature extraction and matching.

To sum up, our main contributions are as follows:

- We propose a novel deep stitching network called *RecStitchNet*, which does not rely on the fragile and expensive feature matching found in traditional methods, so is much more robust and efficient compared to these methods. As our extensive experimental results demonstrate later, our method achieves state-of-the-

art performance, both qualitatively and quantitatively, outperforming traditional methods and deep learning baselines.

- We introduce an unsupervised instance-wise learning strategy to iteratively optimize stitching results, to ensure high-quality stitching in a wide range of scenarios.
- Given the absence of an existing dataset for supervised learning, we have created a new dataset, which includes pseudo-ground-truth mesh warping results, strictly selected and re-rendered from traditional stitching results with rectangular boundaries.

2. Related Work

2.1. Traditional Image stitching

Image stitching refers to aligning multiple images with mutual overlaps and producing a new image with a larger FOV. The key problem of image stitching is to keep accurate feature alignment, with unnoticeable distortion. Earlier works based on a single homography [1] and dual-homographies are limited to parallax and perspective variations.

To compensate for the shortcomings of a globally projective model, a number of spatially varying warping models, which can better address local alignment, have been proposed, such as smoothly varying affine stitching [19], as-perspective-as-possible (APAP) stitching [34], piecewise planar region matching [20], and seam-guided warping [18, 35, 6].

To produce more natural stitching with less perspective distortion, several warping schemes have been proposed, characterized by shape-preserving half-projection (SPHP) [2], adaptive as-natural-as-possible (AANAP) [17], global similarity prior [3], quasi-homography [14], single perspective [16], and geometric structure preserving [4].

Recently, Jia *et al.* [10] considered global collinear structures, effectively preserving global and local structures while reducing distortions. Zhang *et al.* [36] proposed manifold preserving stitching: using on-manifold operations helps to reduce ghosting and distortion artifacts. To improve stitching results, seam-cutting methods have been applied to removing artifacts in overlapping regions [13, 6].

The most relevant work to our paper comes from Zhang *et al.* [37], in which boundary regularity constraints are incorporated into the stitching framework, helping to solve the irregular boundary problem in image stitching. Although successful in many examples including some challenging cases, the method in [37] may fail in situations such as those with unclear textures, low lighting and low resolution. In addition, the two-step energy optimization process is also time-consuming.

2.2. Deep Image stitching

Unlike the above methods, deep stitching learns to stitch images by extracting high-level features from large datasets, which avoids the difficulties in feature matching, global and local structure preservation, etc. We may roughly divide recent research into three main types:

2.2.1 Supervised learning

Nie *et al.* [22] and Zhao *et al.* [38] proposed view-free image stitching based on global homography learning, which improves upon the previous learning based stitching method [12] which is limited to relatively fixed views. To tolerate parallax in stitching, they generate a synthetic dataset from an existing real image dataset. Instead of homography based learning, Kweon *et al.* [11] recently proposed a novel deep stitching framework using a pixel-wise warp field, which can handle the large-parallax problem well.

2.2.2 Weakly Supervised Learning

To overcome the difficulties in dataset and ground truth generation, Song *et al.* [30] proposed a weakly-supervised learning method to train the stitching model without using real ground truth images. They have further extended their method to stitching multiple images and creating 360-degree panoramas.

2.2.3 Unsupervised Learning

Considering the difficulties in data label generation, some works focus on unsupervised learning methods, which train stitching models without labels. Nie *et al.* [23] proposed an unsupervised image stitching method, which consists of unsupervised coarse image stitching and image reconstruction. Very recently, Nie *et al.* [26] further proposed a parallax-tolerant unsupervised image stitching method which is characterized by combining homography and thin-plate splines (TPS) into a unified framework.

2.3. Enforcing image rectangularity

Enforcing image rectangularity aims to regulate the irregular boundaries caused by image stitching, rotation, etc. The pioneering work in imposing image rectangularity is [7], in which a content-aware warping method was proposed using a two-stage warping based optimization on meshes. Wu *et al.* [33] further extended imposing rectangularity to videos, incorporating temporal coherence into the warping-based optimization. Nie *et al.* [24] proposed a one-stage learning baseline of deep rectangularity imposition for image stitching. Compared to the two-stage method in [7], the method in [24] is more efficient and robust, and

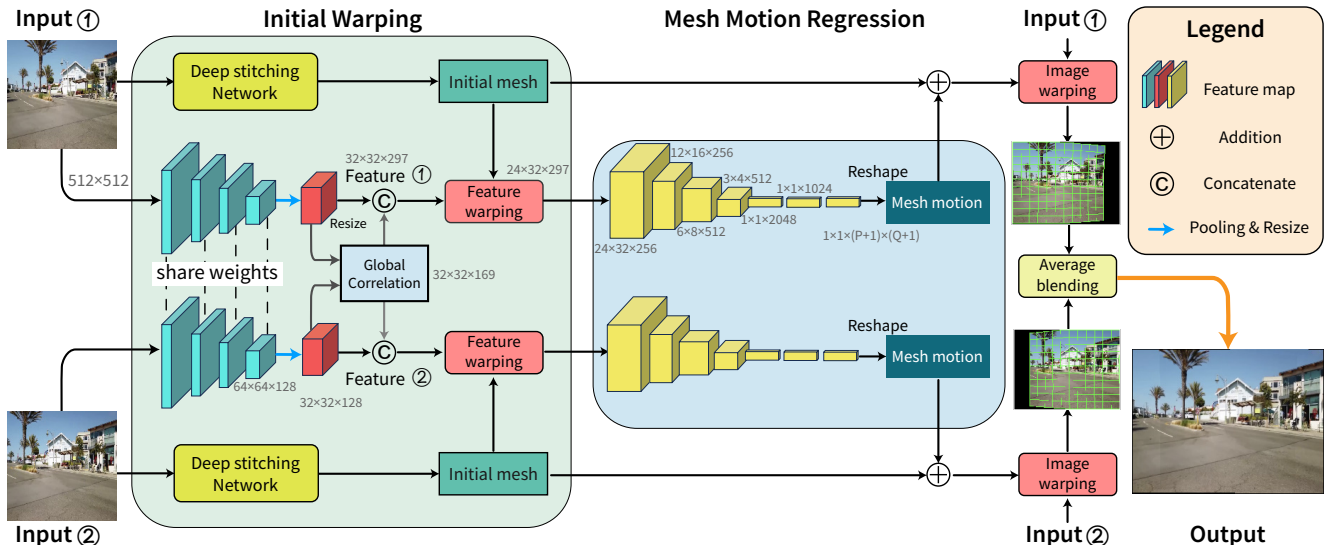


Figure 2. Overview of our network for deep stitching with rectangular boundaries. Our supervised learning network consists of initial warping and mesh motion regression stages. The first stage warps the high-level features extracted from the input images; the warping is guided by the meshes generated by the deep stitching model by Nie *et al.* 2023. Using the warped features as input, we further obtain the mesh motions (vertex offsets to apply to the initial meshes) through mesh motion regression. Final stitching results are obtained by averaging the images warped by the meshes produced by combining the initial meshes and mesh motions.

can well preserve non-linear structures thanks to high-level feature extraction in the learning framework. Liao *et al.* [?] proposed a rectangularity imposing rectification network, which applies the TPS module to perform non-linear and non-rigid transformations for imposing rectangularity on wide-angle rectified images. Very recently, Zhou *et al.* [39] combined stitching and imposing rectangularity into a unified end-to-end framework using a synthetic dataset. Although effective in producing stitching results with rectangular boundaries, it still suffers from content loss and ghosting effects in the overlapping regions.

3. Method

3.1. Overview

Like recent work on stitching and imposing rectangularity [3, 16, 37, 24, 26], we also stitch images by content-aware mesh warping. Mesh warping is widely used in image manipulation due to its simplicity and efficiency. Traditional methods [3, 16, 37] are based on energy optimization with constraints on all grid vertices of the mesh. Let $V = \{V^i, i = 1, \dots, N\}$ be the sets of all vertices of the input images, where N is the number of images. We aim to obtain warped mesh vertices $\hat{V} = \{\hat{V}^i\}$ by minimizing the energy function $E(\hat{V})$, which includes several content-aware constraints, such as feature alignment, shape preservation, straight line preservation etc. Such methods usually focus on designing energy terms that are effective in stitching and easy to optimize. Unlike traditional methods, deep learning based methods [24, 26] focus on dataset prepara-

tion, network construction, and mesh regression. For effective mesh regression, we have to focus on designing the objective function to effectively guide the training process, and the total loss, to ensure satisfactory convergence. To calculate the feature loss after mesh warping, an effective and efficient warping operation is required, which must also be differentiable for effective gradient propagation.

Inspired by the methods above, we propose a novel method to achieve stitching and imposing boundary rectangularity simultaneously in a learning-based framework. Fig. 2 shows an overview of our method for deep stitching with rectangular boundaries. We take two images of the same size with partial overlap as input; the output is a rectangular stitching result with no loss of content. We first perform initial stitching, which aims to warp the high-level features extracted from input images. The warping is guided by the initial meshes generated by a state-of-the-art deep stitching model [26]. We further learn to regulate the boundary of the stitching result by designing a regression network, which generates suitable mesh motion to be applied to the initial meshes. Finally, with the combination of initial meshes and these mesh motions, the final stitching result can easily be produced by warping and average blending.

3.2. Initial Warping

In this stage, the input images are initially stitched using a leading deep stitching model [26]. As illustrated in Fig. 3, given two input images $\{I^i, i = 1, 2\}$ to be stitched,




Figure 3. Mesh manipulations in initial stitching and mesh motion regression.

a uniform quad mesh v^i is placed on each image I^i . The initial warping produces warped meshes $\{\xi^i\}$ after the deep stitching process. To facilitate the deep stitching task, the first image is consistently kept unchanged, while the other image is warped into alignment with it. Next, high-level semantic features are extracted from each input image (without warping) through a series of convolution and pooling blocks (blue solid blocks in Fig. 2); each block comprises two convolution layers. After the first, second, and third blocks, a max-pooling layer is applied. We set the number of channels to 64 and 128 for the convolution layers in the first two and the last two blocks, respectively. Subsequently, following the last blocks, an adaptive pooling layer is employed to standardize the resolution of the features.

To establish the relationship between the two images in the overlapping regions, we concatenate the global correlation [22] with the features extracted for each image. Given the extracted features (F^1, F^2) of the input images, their global correlation refers to their feature-wise similarities, defined by:

$$\text{Cor}(x_1, x_2) = \frac{\langle F^1(x_1), F^2(x_2) \rangle}{|F^1(x_1)| |F^2(x_2)|}, \quad (1)$$

where x_1, x_2 represent the locations of the feature vector in each feature map. We limit the range of feature similarity comparison for fast calculation of global correlation. These features are then warped using the meshes $\{\xi^i\}$ obtained through the initial stitching process. These warped features $\{\kappa^i\}$ of the input images, representing the features of the initial stitching results, serve as input to the mesh motion regression process.

3.3. Mesh Motion Regression

In this stage, our goal is to obtain the mesh motion, as offsets from the initially warped mesh vertices; this helps regulate the shape of the final stitching boundaries. As indicated in the middle section of Fig. 2, the input to this stage consists of the extracted high-level features that have been warped by the meshes from the initial stitching. We provide a simple yet effective fully convolutional network to predict both the vertical and horizontal motion, denoted as

$\{\chi^i\}$, for all vertices relative to those in the initially warped meshes $\{\xi^i\}$. The output of the regression is of dimension $(P + 1) \times (Q + 1)$, where P and Q denote the resolution of the meshes. Subsequently, we determine the predicted meshes $\{\psi^i\}$ by combining the mesh motion $\{\chi^i\}$ with the initial warped meshes $\{\xi^i\}$, as shown in Fig. 3. With the incorporation of mesh motion, the outer boundary of the combined meshes more closely approximates a rectangle.

3.4. Loss Functions

3.4.1 Loss

Our regression network learns the motion of the mesh vertices that can ensure both feature alignment and boundary regularity. We use three loss terms and define the total loss as follows:

$$L_{\text{train}} = \varphi_m l_m + \varphi_p l_p + \varphi_s l_s, \quad (2)$$

where l_m, l_p, l_s refer to the mesh, perception, and shape preserving loss terms, and $\varphi_m, \varphi_p, \varphi_s$ are corresponding weights.

For our supervised training framework, we have prepared a large dataset (refer to Section 3.6), which contains input image pairs, pseudo-ground-truth mesh labels representing warped mesh vertices, and pseudo-ground-truth stitching result labels. The pseudo-ground-truth data is produced by a leading traditional stitching method which produces rectangular boundaries [37].

3.4.2 Mesh Loss

Given the predicted meshes $\{\psi^i\}$, we simply constrain them to be close to the ground truth labels of meshes Ψ^i , using:

$$l_m = \sum_{i=1}^2 \sum_{j=1}^{(P+1) \times (Q+1)} \|\psi_j^i - \Psi_j^i\|_1, \quad (3)$$

where ψ_j^i and Ψ_j^i refer to the vertex positions of the predicted mesh and the pseudo-ground-truth mesh.

3.4.3 Perceptual Loss

We further constrain the result to be visually appealing, and to preserve the structure in the input image, such as linear or salient structures. We define the corresponding loss as follows:

$$l_p = \sum_{i=1}^2 \|\Gamma(\Omega_{\text{tps}}(\psi^i, I^i)) - \Gamma(\Omega_{\text{tps}}(\Psi^i, I^i))\|_1, \quad (4)$$

where $\Omega_{\text{tps}}(\cdot)$ refers to the TPS transformation [9], which is used to warp the input images $\{I^i\}$ guided by the warped mesh, and $\Gamma(\cdot)$ refers to the VGG-19 [29] feature extractor.

3.4.4 Shape Preserving Loss

Following [24], we also preserve the shape of the mesh using intra-grid and inter-grid shape similarity constraints, using:

$$l_s = l_s^{\text{intra}} + l_s^{\text{inter}}. \quad (5)$$

The intra-grid constraint is employed to enforce both the scale and direction of the grid edges, and is defined as follows:

$$l_s^{\text{intra}} = \sum_{i=1}^2 \frac{\sum_{\vec{e}_j \in \vec{h}_i} \vartheta(\Delta_x(\vec{e}_j) + \sigma W/Q)}{(P+1)Q} + \sum_{i=1}^2 \frac{\sum_{\vec{e}_k \in \vec{v}_i} \vartheta(\Delta_y(\vec{e}_k) + \sigma H/P)}{P(Q+1)}, \quad (6)$$

where \vec{e}_j and \vec{e}_k refer to all horizontal and vertical edges of a mesh respectively, $\Delta_x(\vec{e}_j)$ and $\Delta_y(\vec{e}_k)$ refer to projections of the edge vectors onto x and y directions. ϑ is the ReLU function, which is used to cause the direction of the horizontal and vertical edges to be right and bottom, and enforce their scale to be more than $\sigma W/Q$ and $\sigma H/P$; we set $\sigma = 0.8$ in this paper.

The inter-grid constraint aims to cause pairs of successive horizontal and vertical grid edges $\{\vec{e}_{t1}, \vec{e}_{t2}\}$ to undergo linear changes (i.e. encouraging their angle to be close to zero). It is defined as follows:

$$l_s^{\text{inter}} = \sum_{i=1}^2 \frac{1}{|\Lambda^i|} \sum_{\{\vec{e}_{t1}, \vec{e}_{t2}\} \in \Lambda^i} (1 - \cos(\vec{e}_{t1}, \vec{e}_{t2})), \quad (7)$$

where $\cos(\vec{e}_{t1}, \vec{e}_{t2})$ calculates the cosine of the angle between \vec{e}_{t1} and \vec{e}_{t2} , Λ^i refers to the set of all successive grid edges in the mesh of i^{th} image, and $|\Lambda^i|$ is the total number of successive grid edges.

3.5. Unsupervised Instance-wise Stitching Refinement

3.5.1 Approach

In the mesh motion regression step, our loss functions are designed to cause the predicted stitching result to be close to

both the pseudo-ground-truth mesh and the stitched image while preserving mesh shape. However, our experiments showed that some predicted results may not exhibit perfect boundary regularity and feature matching (see Fig. 4). Simply incorporating feature matching and rectangular boundary constraints into the network training process does not yield satisfactory results. This is because the refinement objective (unsupervised learning) is slightly contradictory to the original optimization goal of *RecStitchNet* (supervised learning using pseudo mesh labels), preventing the network parameters from being optimized. To enhance stitching performance and enable the transfer of the pretrained model to other datasets, we propose an instance-wise unsupervised learning method constrained by feature matching, rectangular boundary and shape preservation constraints, which are designed to further optimize the mesh grid, so as to refine the imperfect rectangular boundaries and the ghosting in stitched images.

As Fig. 4 shows, while the predicted stitching result appears quite satisfactory, it still exhibits irregular boundaries and misalignment in the overlapping regions. To further refine the stitching result, we introduce an instance-wise unsupervised learning scheme to iteratively optimize the stitching (see Alg. 1). The input consists of the predicted meshes $\{\psi^i\}$ generated by our pre-trained regression network, along with the corresponding input image pair $\{I^i\}$. The output comprises the optimized meshes $\{\Theta^i\}$, for $i = 1, 2$, and the stitching result Φ . To iteratively optimize the stitching boundary, we first need to obtain the boundary vertices of the stitching result. Drawing inspiration from [37], we treat the outer boundaries of the meshes $\{\psi^i\}$ as polygons $\{\hat{P}^i\}$. Subsequently, the outer boundary P of the two meshes is calculated using a polygon Boolean union operation [21], as follows:

$$\hat{P} = \hat{P}^1 \cup \hat{P}^2. \quad (8)$$

With the outer boundary vertices of the stitching results, we are able to construct an effective constraint for the rectangular boundary. In each iteration, we first predict the mesh motions $\{\chi_i\}$ relative to the current meshes $\{\psi^i\}$ using an unsupervised learning network with the same architecture as *RecStitchNet*. $\{\psi^i\}$ is then updated for the next iteration. We then compare the current loss value with the value from the previous iteration. The iterations terminate when the difference in loss is sufficiently small. Finally, we warp the input images using the final optimized meshes $\{\Theta^i\}$, and obtain the final stitching result Φ through average blending of the warped images.

In the refinement step, we use a different set of loss functions for the instance-wise unsupervised learning, defined as below.

Algorithm 1: Refine the stitching results.

Input: Predicted meshes $\{\psi^i\}$ produced by our pre-trained regression network, and input image pair $\{I^i\}$, $i = 1, 2$;
Output: Optimized meshes $\{\Theta^i\}$, $i = 1, 2$ and stitching results Φ ;
Let \hat{P}^i be the boundary vertices of ψ^i ;
Let \hat{P} be the outer boundary vertices of the two meshes $\{\psi^i\}$, $i = 1, 2$, calculated by Equ. 8;
foreach $j \in [1, 200]$ **do**
 foreach $i \in [1, 2]$ **do**
 $\chi_i = \text{RecStitchNet}(\psi^i, I^i)$;
 $\psi^i = \psi^i + \chi_i$;
 $\Theta^i = \psi^i$;
 end
 if $j == 1$ **then**
 $\text{Loss}_{\text{pre}} = \text{Loss}(\text{RecStitchNet})$;
 end
 else
 $\text{Loss}_{\text{now}} = \text{Loss}(\text{RecStitchNet})$;
 if $|\text{Loss}_{\text{now}} - \text{Loss}_{\text{pre}}| < e^{-5}$ **then**
 break;
 end
 $\text{Loss}_{\text{pre}} = \text{Loss}_{\text{now}}$
 end
end
foreach $i \in [1, 2]$ **do**
 $R^i = \Omega_{\text{tps}}(\Theta^i, I^i)$
end
 $\Phi = \text{AverageBlend}(R^1, R^2)$;

3.5.2 Feature Matching Loss

The feature matching constraint is designed to ensure that the features of the two images in the overlapping regions are well-aligned. It is defined as the difference between the warped image features in the overlapping regions, as follows:

$$l_f = \left\| \sum_{i=1}^2 (\Gamma(\Omega_{\text{tps}}(\psi^i, I^i) * M * (-1)^{i-1})) \right\|_1, \quad (9)$$

where $\Omega_{\text{tps}}(\cdot)$ refers to the TPS transformation, and $\Gamma(\cdot)$ refers to the VGG-19 [29] feature extractor. M is the intersection of the warped masks guided by the predicted meshes $\{\psi^i\}$.

3.5.3 Rectangular Boundary Loss

In [24], the rectangular boundary loss is simply defined as the difference between the $\{0, 1\}$ mask of the result and the all-one mask. However, we found in experiments

that incorporating this form of loss has almost no effect on shaping the rectangular boundary probably due to the difficulty of gradient propagation. To effectively optimize the stitched boundary, we first extract the outer boundary \hat{P} of the warped and overlaid meshes $\{\psi^i\}$. We assign several attributes to each vertex ν_k in \hat{P} , including their constraint directions $\rho(\nu_k) \in \{[1, 0], [0, 1]\}$ (in x and y directions), and their target values $\tau(\nu_k)$ (the values in the top/bottom/left/right directions). Finally, the loss is defined as the sum of the differences between all vertices and their target locations, as follows:

$$l_b = \left\| \sum_{\nu_k \in \hat{P}} (\nu_k \cdot \rho(\nu_k)) - \tau(\nu_k) \right\|_1. \quad (10)$$

In this stage, the total loss function is a linear combination of feature matching, rectangular boundary, and shape preserving constraints (see details in Sec. 3.4.4), as follows:

$$L_{\text{refine}} = \varphi_f l_f + \varphi_b l_b + \varphi_s l_s, \quad (11)$$

where φ_f , φ_b , φ_s are corresponding weights to control their relative importance.

3.6. Data Preparation

Recently, there have been very few datasets available for image stitching, and defining their labels (i.e., ground-truth results) is quite challenging. To the best of our knowledge, no dataset suitable for our method exists yet. Unlike traditional stitching methods, which often yield results with irregular boundaries, our objective is to achieve stitching results with rectangular boundaries. This makes it considerably easier to define labels for the stitching process.

To train a deep learning network for image stitching with rectangular boundaries, we have established a new dataset, which comprises input images, mesh labels, and image labels (refer to Fig. 5). Data preparation was carried out as follows:

- *Stitching*: Input image pairs were sourced from the *training* dataset of UDIS-D proposed in [23]. We performed stitching using the traditional warping-based method described in [37]. This method is capable of generating stitching results with rectangular boundaries, along with corresponding meshes for each input image. Given our focus on stitching with rectangular boundaries, we prefer to omit data where the stitching result contains an excessive amount of missing content. Furthermore, our stitching method does not have to account for piecewise rectangular boundaries as in [37].
- *Normalization*: For effective training, mesh labels should be constrained within a certain range. However, the scale of stitching results tends to vary greatly.

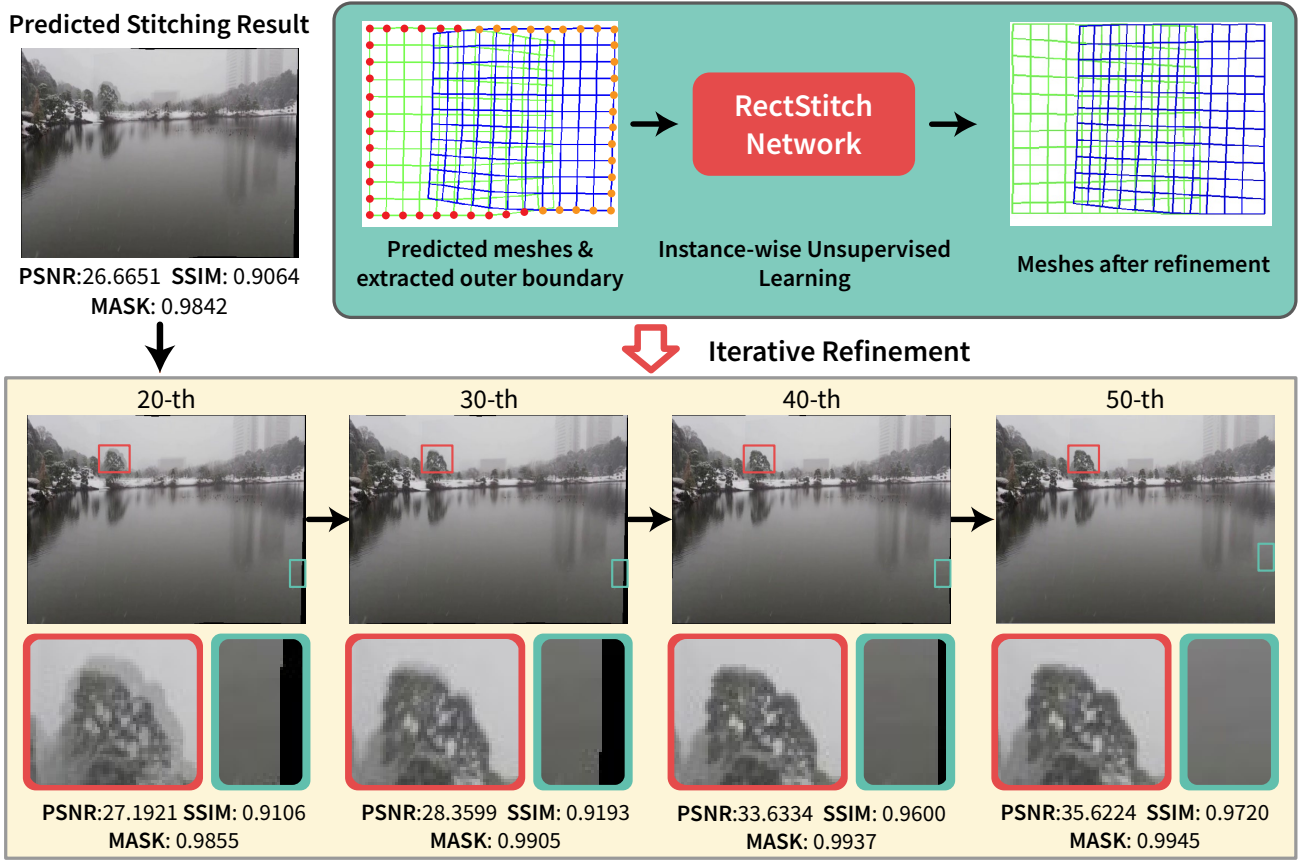


Figure 4. Refinement of stitching results. The input to this step comprises the predicted meshes and the corresponding input image pair. We first extract the outer boundary of the predicted meshes using a polygon Boolean operation, then predict the refined meshes using the instance-wise unsupervised learning framework in an iterative manner. Finally, we obtain an optimized stitching result by warping and blending.




Figure 5. Dataset preparation. Give a pair of input images, we first stitch them using the method from Zhang *et al.* 2021, and then output the resulting mesh labels and the corresponding warped image labels.

Consequently, we set the resolution of the stitching result to be $W_s \times H_s$; for each vertex of the mesh with coordinates (x, y) , we converted them to $(xW_s/w_t, yH_s/h_t)$, where w_t and h_t represented the outer boundary size of the stitching result.

- *Rendering*: We further rendered the stitching results by warping the input images guided by the normalized meshes. To achieve smoother stitching results, our rendering was performed using the TPS transformation [9], which provides more natural transitions and

smoother interpolation than mesh-based warping. At this stage, the resolution of each rendered image was also set to $W_s \times H_s$.

Actually, the stitching labels produced by [37] cannot be considered as ideal labels due to limitations in their approach. This may impact the performance of the training model. In this paper, we utilize these labels, considered as pseudo-ground-truth, for supervised training. To break the bottleneck of the pseudo-ground-truth, we further refined the stitching results using our unsupervised instance-wise

learning. Experimental results and evaluations in Section 4 demonstrate that our refined results improve upon the training labels produced by [37].

4. Experiments

4.1. Implementation details

In the data preparation and training stages, we set the mesh resolution of each image to 11×11 , and resolution of each input image was normalized to 512×512 . In the feature extraction and regression stages, we set kernel-size=3, stride=2 for all convolution blocks and kernel-size=2, stride=2 for all max pooling layers; we set the search range to 6 to efficiently calculate the correlation of two feature maps (32×32); the size of the correlation is ($4 \times 32 \times 32 \times 169$). In the training stage, we used a linear combination of conv4_2, conv3_2 and conv2_2 layers of the VGG-19 features as the high-level feature of an image. The weights of loss terms were set to $\varphi_m = 1$, $\varphi_p = 0.000006$, $\varphi_s = 0.8$. As for many CNN-based networks [24], we used the Adam optimizer with a learning rate initialized to 10^{-4} for 10^5 iterations, and a decay rate of 0.9. We set batch-size=4 and used ReLU as the activation function. In the stitching refinement stage, we set $\varphi_f = 0.0001$, $\varphi_b = 1$, $\varphi_s = 1$, and the decaying learning rate was initialized to 0.002 for fast refinement. All implementation was based on TensorFlow using a single GPU with an Nvidia RTX 4090. To better compare the performance of different stitching methods, we simply used average blending to composite the overlapping regions.

4.2. Evaluation

To assess the effectiveness of our method, we conducted both qualitative and quantitative evaluations. We compared our method to state-of-the-art methods that have public source code.

Fig. 6 displays several stitching results from examples sourced from the *testing* dataset of UDIS-D, which were unseen during training. Given a pair of input images, we shows results of performing stitching using the methods from [24, 37] and our method separately. For the method in [24], we initiate the process with an initial stitching using the deep stitching method from [26]. By subsequently utilizing the stitching result and corresponding mask, the final stitching result is obtained through the deep rectangularity imposition method from [24]. Zhang *et al.*'s method [37], carries out stitching through a global optimization process. To obtain our result, we first stitch images using the proposed *RecStitchNet*, and then refine the stitching to produce improved results. In comparison to [24, 37], our method excels at shaping the rectangular boundary and ensuring precise alignment in the overlapping regions. The marked red and green boxes (close-up views), along with the PSNR

(peak signal-to-noise ratio) and SSIM (structural similarity) metrics, highlight the advantages of our method in terms of shape preservation, boundary regularity, and feature alignment. To quantify the performance of boundary regularity, we define the Mask metric by calculating the proportion of white pixels inside the warped stitching mask, which serves as a demonstration of our proficiency in preserving rectangular boundaries.

To validate the effectiveness of our method, we conducted further tests on data unseen in UDIS-D, and some of which was previously utilized in certain traditional stitching methods [3, 37]. Fig. 7 showcases stitching results and makes comparisons to the methods presented in [24, 37]. The close-up views illustrate that our method excels in aligning salient structures, such as lines and characters, and better maintains rectangular boundaries. Fig. 8 provides more results and comparisons. The red rectangles highlight the shortcomings of [24, 37] in terms of structure preservation, feature alignment, and boundary regularity. Furthermore, we offer a quantitative evaluation in Table 1, which vividly compares the performance of different methods. Results obtained from [24] excel in feature alignment due to the complete separation of stitching and imposing rectangularity. However, they cannot guarantee rectangular boundaries and preservation of salient structures. In [37], where stitching and imposing rectangularity are accomplished through global optimization, the rectangular boundaries are well preserved, but artifacts tend to appear in terms of feature alignment.

We performed further extensive quantitative evaluations on the *testing* dataset of UDIS-D, as shown in Table 2. Pseudo-ground-truth in the column 2 refers to the metrics of the results from [37], which are used as training labels. The last two columns present the metrics of our stitching results before and after refinement. The metrics include PSNR, SSIM, and Mask, which are used to measure feature alignment in the overlapping regions as well as the boundary regularity. In experiments, we find that limitations of [37] may cause it to fail to generate stitching results when images exhibit characteristics such as low light, low texture, low contrast, low overlap etc. Out of the 1106 examples in the *testing* dataset of UDIS-D, 1068 examples were successfully stitched by [37], and the remaining 38 examples could not be stitched correctly. For a fair comparison, the quantitative evaluation was performed on the selected 1068 examples and the remaining 38 examples, separately. The results in Table 2 vividly show the advantages of our method.

Additionally, we select some stitching results produced from the remaining 38 examples of the *testing* dataset of UDIS-D, which exhibit characteristics such as low light, low texture, low contrast and low overlap. Both the visual results and metrics in Fig. 9 illustrate that our method is effective and robust in challenging scenarios. Both quan-

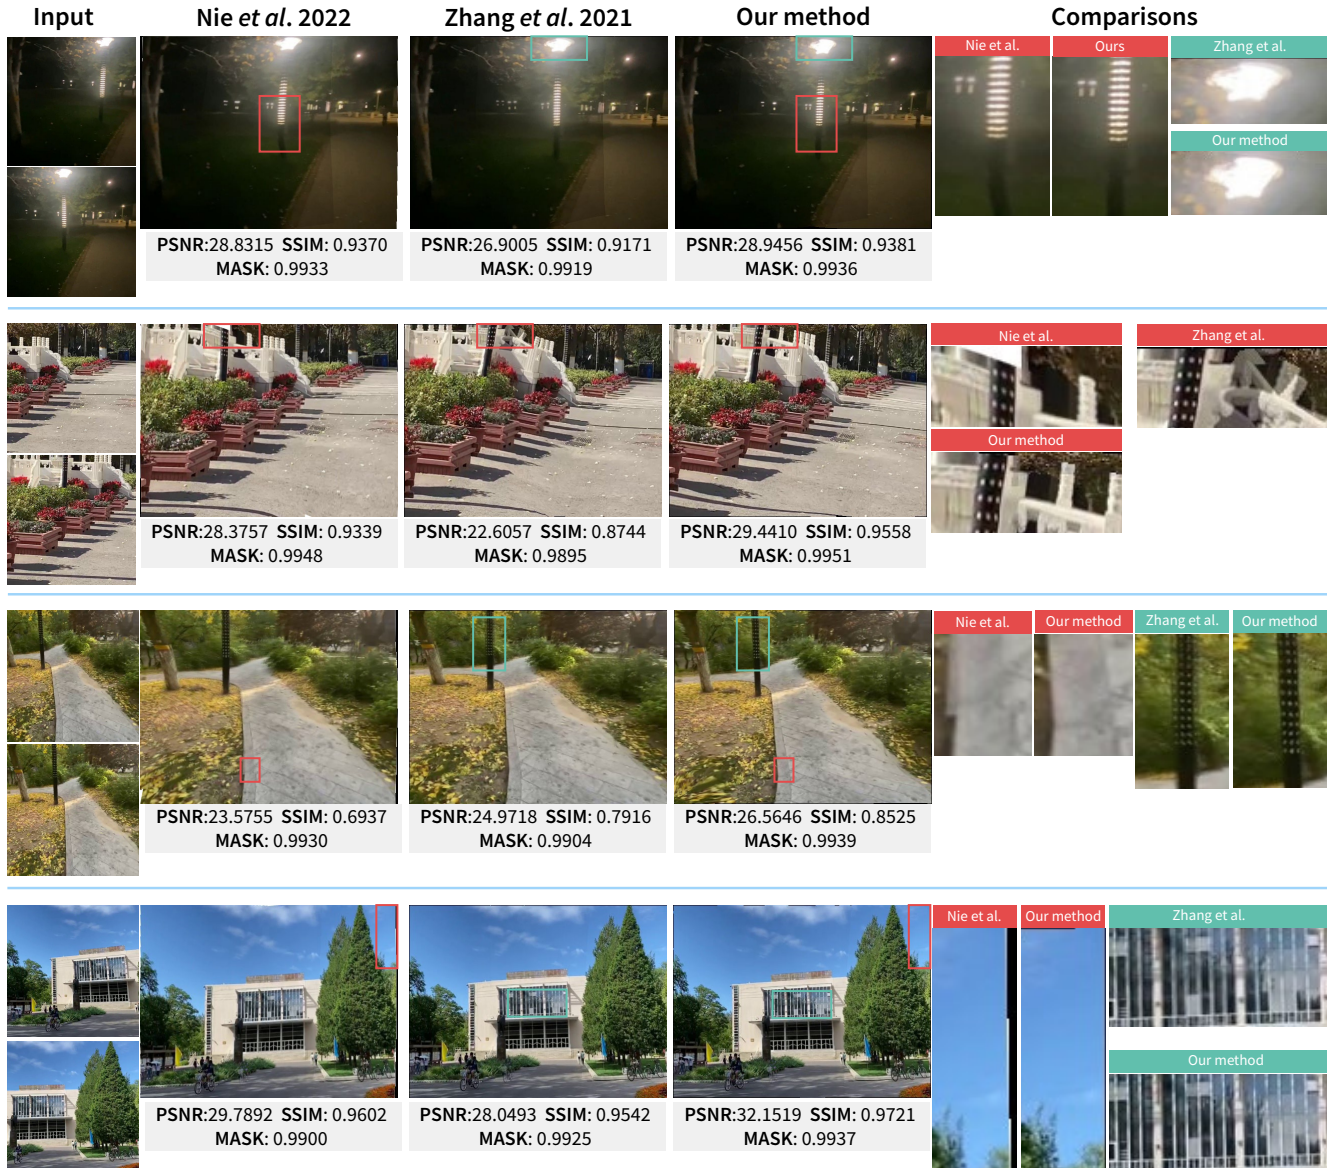


Figure 6. Stitching results and comparisons on the *testing* dataset of UDIS-D.

titative and qualitative results affirm the effectiveness of our method in terms of feature alignment, regular boundary preservation, and structure preservation.

4.3. Ablation Study

As in the case of the quantitative evaluation in Table 2, we also selected the *testing* dataset of UDIS-D for an ablation study.

We first provide visual results from the ablation study in Fig. 10. From the close-up views and quantitative metrics, it is easy to see that without the mesh and shape constraints, the results are completely unacceptable: there are significant artifacts in feature alignment and shape distortions. Without the perception and correlation constraints, the re-

sults are much better, but still have noticeable ghosting and irregular boundary artifacts.

We further conducted a quantitative evaluation for the ablation study test to assess the role of each constraint term and the global correlations, using as metrics PSNR, SSIM and Mask. In the ablation study, we observed that the ‘Mask’ metric may not accurately represent the regularity of boundaries, as the mesh vertices often exceed the target rectangular boundary, especially when there is no mesh label loss. Therefore, we additionally employed a Boundary metric, which measures the distance between the vertices on the outer boundary and their target positions (as detailed in Section 3.5.3). Table 3 shows that all the constraint terms and the global correlation play an important role in improv-

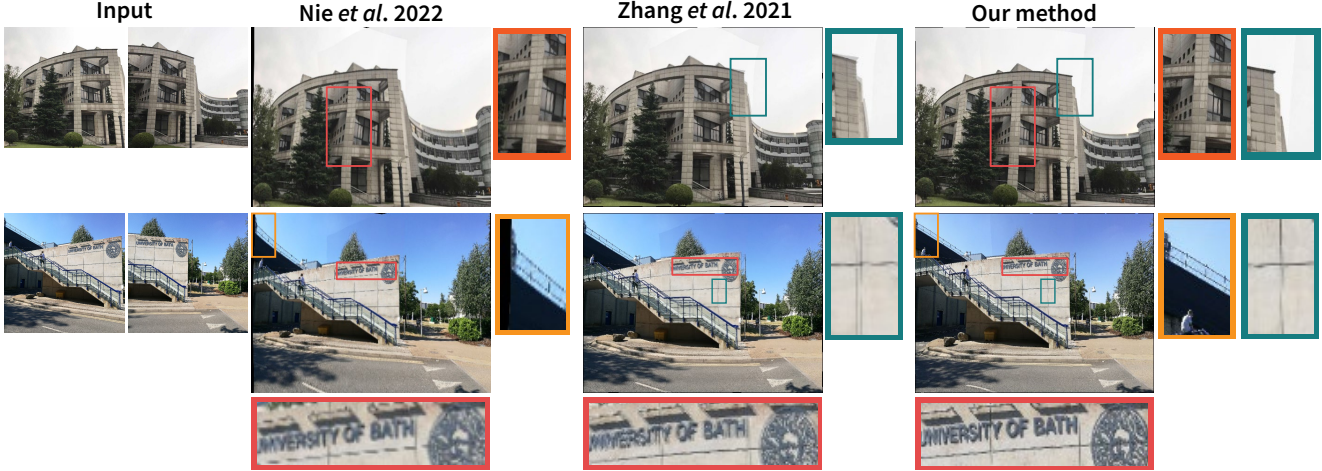


Figure 7. Stitching results and comparisons on data unseen in UDIS-D.

Table 1. Metrics for examples in Figs. 7 and 8.

Metrics	Fig. 7(1)	Fig. 7(2)	Fig. 8(1)	Fig. 8(2)	Fig. 8(3)	Fig. 8(4)	Fig. 8(5)	Fig. 8(6)
<i>Nie et al. 2022</i>								
PSNR	25.7282	26.8942	28.3332	27.8733	28.3954	30.1785	30.6975	29.9269
SSIM	0.9442	0.9390	0.9522	0.9430	0.9278	0.9490	0.9668	0.9499
Mask	0.9905	0.9873	0.9895	0.9873	0.9906	0.9925	0.9944	0.9927
<i>Zhang et al. 2021</i>								
PSNR	23.3887	23.1461	25.9588	24.9477	26.6624	27.2616	27.8547	28.2692
SSIM	0.8947	0.8597	0.9161	0.8844	0.9102	0.8901	0.9335	0.9171
Mask	0.9903	0.9908	0.9920	0.9924	0.9921	0.9920	0.9919	0.9920
<i>Ours</i>								
PSNR	25.4717	26.0382	29.0105	30.0713	28.8899	33.6945	32.9202	32.0031
SSIM	0.9618	0.9098	0.9619	0.9394	0.9454	0.9647	0.9660	0.9486
Mask	0.9951	0.9935	0.9936	0.9950	0.9940	0.9932	0.9947	0.9934

Table 2. Quantitative evaluation on the *testing* dataset of UDIS-D. The upper part and the lower part give quantitative evaluations on the selected 1068 examples and the remaining 38 examples of the *testing* dataset of UDIS-D.

Metrics	Pseudo-ground-truth	Nie <i>et al.</i> 2022	Ours	Ours+Refinement
<i>Selected 1068 testing examples of UDIS-D</i>				
PSNR	25.0656	25.9212	21.3544	27.7812
SSIM	0.8454	0.8581	0.7020	0.8958
Mask	0.9903	0.9913	0.9889	0.9941
Boundary	0.0002	0.00017	0.0016	0.00014
<i>Remaining 38 testing examples of UDIS-D</i>				
PSNR	-	24.7549	20.9781	26.7898
SSIM	-	0.8292	0.7281	0.8772
Mask	-	0.9909	0.9742	0.9920

ing the performance of stitching.

4.4. Speed

In terms of running times, our experiments show that the average time for image stitching for the *testing* dataset is 51

ms, which is significantly faster than the traditional method in [37]. Each iteration of stitching refinement requires 35 ms, with an average of 50 iterations. We provide a speed comparison on the *testing* dataset of UDIS-D in Table 4. For Nie *et al.*'s [24] method, the running time includes both




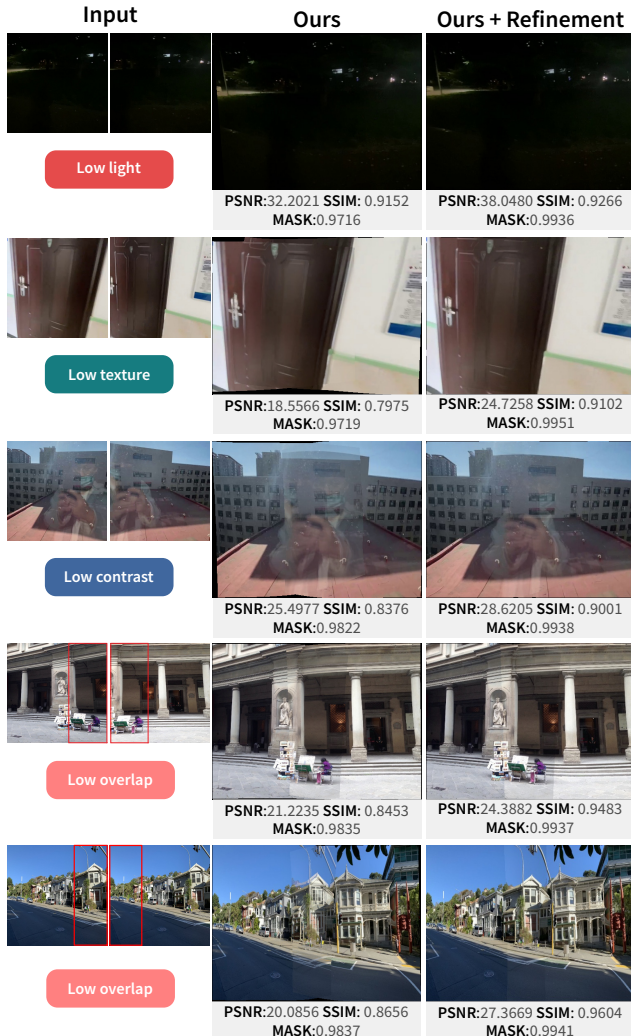
Figure 8. Further results and comparisons. The red rectangles point out artifacts in feature alignment, structure preservation, and rectangular boundary preservation.

the time spent in initial stitching using the learning-based

method [26] and their rectangularity imposition process. In

Table 3. Ablation study metrics.

Metrics	w/o shape	w/o perception	w/o mesh	w/o correlation	Our method
PSNR	18.1769	21.0704	17.8817	20.7431	21.3544
SSIM	0.6147	0.6690	0.5753	0.6758	0.7020
Mask	0.9617	0.9828	0.9937	0.9814	0.9889
Boundary	0.0382	0.0017	0.0756	0.002	0.0016

Figure 9. Challenging examples of the *testing* dataset of UDIS-D, including low light, low texture, low contrast and low overlap.

comparison, our learning-based method is shown to be more efficient.

Following refinement, our speed is comparable to that of the traditional method [37]. However, our results surpass it in terms of feature alignment and boundary regularity. Furthermore, our method demonstrates greater robustness in many challenging cases.

Table 4. Average running times.

Nie <i>et al.</i>	Zhang <i>et al.</i>	Ours	Ours+Refinement
0.256 s	1.413 s	0.211 s	1.921 s

4.5. Discussion

In this paper, we propose *RecStitchNet*, which combines imposing rectangularity and stitching in a unified learning based framework. It is natural to compare our method with the two-network cascade approach for stitching and imposing rectangularity. Actually, in this paper, the results of Nie *et al.* [24] are produced by the cascaded stitching and imposing rectangularity. Our method is superior to the cascaded one, and its advantages are as follows. Firstly, artifacts, such as feature misalignment, in the first stitching step cannot be fixed in the following rectangularity imposition step, and can be amplified by warping. In addition, the stitching performance of different methods may also affect the rectangularity imposition effects. Secondly, a cascaded solution cannot ensure globally optimal results in terms of shape preservation, rectangular boundary imposition and feature alignment, while our method takes two normal images as input, and learns to perform stitching and impose rectangularity in a unified framework. With a reasonable and effective network and the unsupervised refinement, our method can stably produce high-quality stitching results.

Having a supervised learning framework, we use pseudo-ground-truth as labels for training. The reason is that so far there is no recognized ground-truth for learning based stitching with rectangular boundary, and it is true that the pseudo-ground-truth is theoretically the upper limit to the boundary of the learned model in this step. However, we have to point out that it would be very difficult to train our *RecStitchNet* without labels, and after training using the pseudo-ground-truth, we can obtain acceptable stitching results at a very small cost, with only a few artifacts regarding feature alignment and rectangular boundary preservation, which also exist in the pseudo labels. To break the bottleneck of pseudo-labels and further improve stitching performance, we further refine the stitching results using an unsupervised learning method, which can produce high quality stitching results with better performance than the pseudo-ground-truth, as shown in our quantitative and qualitative evaluations.

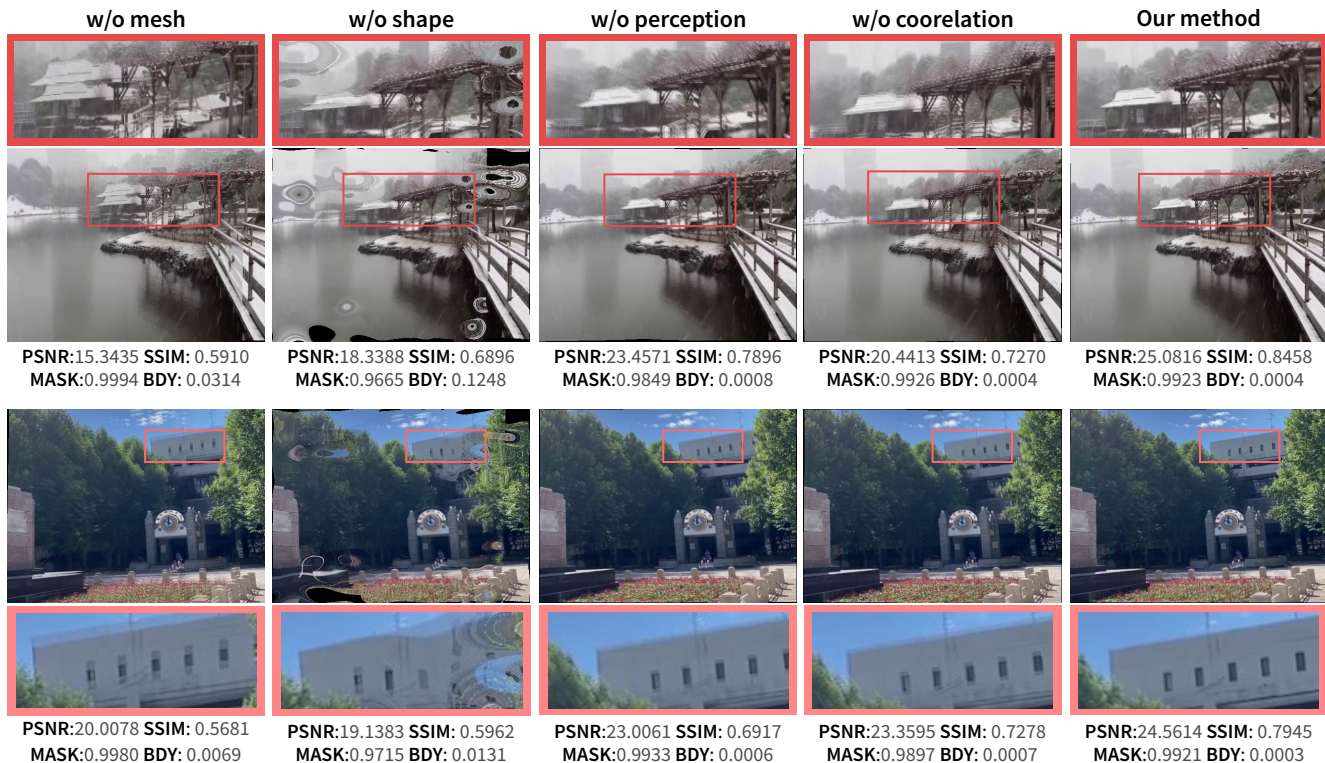


Figure 10. Output images in ablation study using the *testing* dataset of UDIS-D.



Figure 11. Our method may fail to preserve salient structures (e.g. straight lines) near stitching boundaries.

5. Conclusions

This paper has presented *RecStitchNet*, a novel learning-base framework for image stitching with regular boundaries. Compared to traditional stitching and recent learning-based methods, our method can effectively ensure feature alignment, boundary regularity, and salient structure preservation. Our stitching refinement stage enables our model to better adapt to various scenarios and datasets. Although simple yet effective, our method still has some limitations. See Fig. 11: our method may fail to preserve salient structures (e.g. straight lines) near stitching boundaries when there is large content loss. In addition, our method may fail to stitch correctly when there is very little overlap in the images, and this is also challenging for other methods.

In future, we hope to produce more image-stitching datasets with diverse scenarios and high-quality labels, and

further explore more effective networks and constraints for better stitching. In addition, we also would like to extend our learning based framework to video stitching [31, 27], in which stabilization [32] and feature tracking [28] across frames should be considered.

Acknowledgements

This research was supported by the Zhejiang Province Basic Public Welfare Research Program (No. LGG22F020009), Key Lab of Film and TV Media Technology of Zhejiang Province (No.2020E10015), Marsden Fund Council managed by the Royal Society of New Zealand (No. MFP-20-VUW-180). We also would like to express our special gratitude to Dr. Yaqi Wang for her great work in diagram enhancement.

References

- [1] M. Brown and D. G. Lowe. Automatic panoramic image stitching using invariant features. *Int. J. Comput. Vis.*, 74(1):59–73, 2007. 3
- [2] C. Chang, Y. Sato, and Y. Chuang. Shape-preserving half-projective warps for image stitching. In *2014 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2014, Columbus, OH, USA, June 23-28, 2014*, pages 3254–3261. IEEE Computer Society, 2014. 3
- [3] Y. Chen and Y. Chuang. Natural image stitching with the global similarity prior. In B. Leibe, J. Matas, N. Sebe, and

- M. Welling, editors, *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part V*, volume 9909 of *Lecture Notes in Computer Science*, pages 186–201. Springer, 2016. [1](#), [3](#), [4](#), [9](#)
- [4] P. Du, J. Ning, J. Cui, S. Huang, X. Wang, and J. Wang. Geometric structure preserving warp for natural image stitching. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 3678–3686. IEEE, 2022. [3](#)
- [5] J. Gao, S. J. Kim, and M. S. Brown. Constructing image panoramas using dual-homography warping. In *The 24th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2011, Colorado Springs, CO, USA, 20-25 June 2011*, pages 49–56. IEEE Computer Society, 2011. [1](#)
- [6] J. Gao, Y. Li, T. Chin, and M. S. Brown. Seam-driven image stitching. In M. A. Otaduy and O. Sorkine, editors, *34th Annual Conference of the European Association for Computer Graphics, Eurographics 2013 - Short Papers, Girona, Spain, May 6-10, 2013*, pages 45–48. Eurographics Association, 2013. [3](#)
- [7] K. He, H. Chang, and J. Sun. gling panoramic images via warping. *ACM Trans. Graph.*, 32(4):79:1–79:10, 2013. [3](#)
- [8] K. He, H. Chang, and J. Sun. Rectangling panoramic images via warping. *ACM Transactions on Graphics (TOG)*, 32(4):1–10, 2013. [2](#)
- [9] M. Jaderberg, K. Simonyan, A. Zisserman, and K. Kavukcuoglu. Spatial transformer networks. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 2017–2025, 2015. [6](#), [8](#)
- [10] Q. Jia, Z. Li, X. Fan, H. Zhao, S. Teng, X. Ye, and L. J. Latecki. Leveraging line-point consistence to preserve structures for wide parallax image stitching. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 12186–12195. Computer Vision Foundation / IEEE, 2021. [3](#)
- [11] H. Kweon, H. Kim, Y. Kang, Y. Yoon, W. Jeong, and K.-J. Yoon. Pixel-wise warping for deep image stitching. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(1):1196–1204, Jun. 2023. [1](#), [3](#)
- [12] W. Lai, O. Gallo, J. Gu, D. Sun, M. Yang, and J. Kautz. Video stitching for linear camera arrays. In *30th British Machine Vision Conference 2019, BMVC 2019, Cardiff, UK, September 9-12, 2019*, page 130. BMVA Press, 2019. [3](#)
- [13] N. Li, T. Liao, and C. Wang. Perception-based seam cutting for image stitching. *Signal Image Video Process.*, 12(5):967–974, 2018. [3](#)
- [14] N. Li, Y. Xu, and C. Wang. Quasi-homography warps in image stitching. *IEEE Trans. Multimed.*, 20(6):1365–1375, 2018. [3](#)
- [15] K. Liao, L. Nie, C. Lin, Z. Zheng, and Y. Zhao. Recre-net: Rectangling rectified wide-angle images by thin-plate spline model and dof-based curriculum learning. *CoRR*, abs/2301.01661, 2023.
- [16] T. Liao and N. Li. Single-perspective warps in natural image stitching. *IEEE Trans. Image Process.*, 29:724–735, 2020. [3](#), [4](#)
- [17] C. Lin, S. Pankanti, K. N. Ramamurthy, and A. Y. Aravkin. Adaptive as-natural-as-possible image stitching. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*, pages 1155–1163. IEEE Computer Society, 2015. [1](#), [3](#)
- [18] K. Lin, N. Jiang, L. Cheong, M. N. Do, and J. Lu. SEAGULL: seam-guided local alignment for parallax-tolerant image stitching. In B. Leibe, J. Matas, N. Sebe, and M. Welling, editors, *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part III*, volume 9907 of *Lecture Notes in Computer Science*, pages 370–385. Springer, 2016. [3](#)
- [19] W. Lin, S. Liu, Y. Matsushita, T. Ng, and L. F. Cheong. Smoothly varying affine stitching. In *The 24th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2011, Colorado Springs, CO, USA, 20-25 June 2011*, pages 345–352. IEEE Computer Society, 2011. [3](#)
- [20] Z. Lou and T. Gevers. Image alignment by piecewise planar region matching. *IEEE Trans. Multimed.*, 16(7):2052–2061, 2014. [3](#)
- [21] F. Martínez, A. J. Rueda, and F. R. Feito. A new algorithm for computing boolean operations on polygons. *Computers & Geosciences*, 35(6):1177–1185, 2009. [6](#)
- [22] L. Nie, C. Lin, K. Liao, M. Liu, and Y. Zhao. A view-free image stitching network based on global homography. *J. Vis. Commun. Image Represent.*, 73:102950, 2020. [1](#), [3](#), [5](#)
- [23] L. Nie, C. Lin, K. Liao, S. Liu, and Y. Zhao. Unsupervised deep image stitching: Reconstructing stitched features to images. *IEEE Trans. Image Process.*, 30:6184–6197, 2021. [1](#), [3](#), [7](#)
- [24] L. Nie, C. Lin, K. Liao, S. Liu, and Y. Zhao. Deep rectangling for image stitching: A learning baseline. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 5730–5738. IEEE, 2022. [2](#), [3](#), [4](#), [6](#), [7](#), [9](#), [11](#), [13](#)
- [25] L. Nie, C. Lin, K. Liao, S. Liu, and Y. Zhao. Deep rotation correction without angle prior. *IEEE Trans. Image Process.*, 32:2879–2888, 2023. [2](#)
- [26] L. Nie, C. Lin, K. Liao, S. Liu, and Y. Zhao. Parallax-tolerant unsupervised deep image stitching. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 7399–7408, October 2023. [1](#), [3](#), [4](#), [9](#), [12](#)
- [27] Y. Nie, T. Su, Z. Zhang, H. Sun, and G. Li. Dynamic video stitching via shakiness removing. *IEEE Trans. Image Process.*, 27(1):164–178, 2018. [14](#)
- [28] J. Rong, L. Zhang, H. Huang, and F. Zhang. Imu-assisted online video background identification. *IEEE Trans. Image Process.*, 31:4336–4351, 2022. [14](#)
- [29] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In Y. Bengio and Y. LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. [6](#), [7](#)

- [30] D. Song, G. Lee, H. Lee, G. Um, and D. Cho. Weakly-supervised stitching network for real-world panoramic image generation. In S. Avidan, G. J. Brostow, M. Cissé, G. M. Farinella, and T. Hassner, editors, *Computer Vision - ECCV 2022 - 17th European Conference, Tel Aviv, Israel, October 23-27, 2022, Proceedings, Part XVI*, volume 13676 of *Lecture Notes in Computer Science*, pages 54–71. Springer, 2022. [1](#), [3](#)
- [31] M. Wang, A. Shamir, G. Yang, J. Lin, S. Lu, and S. Hu. Biggerselfie: Selfie video expansion with hand-held camera. *IEEE Trans. Image Process.*, 27(12):5854–5865, 2018. [14](#)
- [32] M. Wang, G. Yang, J. Lin, S. Zhang, A. Shamir, S. Lu, and S. Hu. Deep online video stabilization with multi-grid warping transformation learning. *IEEE Trans. Image Process.*, 28(5):2283–2292, 2019. [14](#)
- [33] J. Wu, J. Shi, and L. Zhang. Rectangling irregular videos by optimal spatio-temporal warping. *Comput. Vis. Media*, 8(1):93–103, 2022. [3](#)
- [34] J. Zaragoza, T. Chin, Q. Tran, M. S. Brown, and D. Suter. As-projective-as-possible image stitching with moving DLT. *IEEE Trans. Pattern Anal. Mach. Intell.*, 36(7):1285–1298, 2014. [3](#)
- [35] F. Zhang and F. Liu. Parallax-tolerant image stitching. In *2014 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2014, Columbus, OH, USA, June 23-28, 2014*, pages 3262–3269. IEEE Computer Society, 2014. [3](#)
- [36] L. Zhang and H. Huang. Image stitching with manifold optimization. *IEEE Trans. Multim.*, 25:3469–3482, 2023. [3](#)
- [37] Y. Zhang, Y. Lai, and F. Zhang. Content-preserving image stitching with piecewise gular boundary constraints. *IEEE Trans. Vis. Comput. Graph.*, 27(7):3198–3212, 2021. [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#), [9](#), [11](#), [13](#)
- [38] Q. Zhao, Y. Ma, C. Zhu, C. Yao, B. Feng, and F. Dai. Image stitching via deep homography estimation. *Neurocomputing*, 450:219–229, 2021. [1](#), [3](#)
- [39] H. Zhou, Y. Zhu, X. Lv, Q. Liu, and S. Zhang. Rectangular-output image stitching. In *2023 IEEE International Conference on Image Processing (ICIP)*, pages 2800–2804, 2023. [4](#)