

Inharmonious Region Localization via Recurrent Self-Reasoning

Penghao Wu
wupenghao@sjtu.edu.cn

Li Niu*
ustcnewly@sjtu.edu.cn

Jing Liang
leungjing@sjtu.edu.cn

Liqing Zhang
zhang-lq@cs.sjtu.edu.cn

MoE Key Lab of Artificial Intelligence
Shanghai Jiao Tong University
Shanghai, China

Abstract

Synthetic images created by image editing operations are prevalent, but the color or illumination inconsistency between the manipulated region and background may make it unrealistic. Thus, it is important yet challenging to localize the inharmonious region to improve the quality of synthetic image. Inspired by the classic clustering algorithm, we aim to group pixels into two clusters: inharmonious cluster and background cluster by inserting a novel Recurrent Self-Reasoning (RSR) module into the bottleneck of UNet structure. The mask output from RSR module is provided for the decoder as attention guidance. Finally, we adaptively combine the masks from RSR and the decoder to form our final mask. Experimental results on the image harmonization dataset demonstrate that our method achieves competitive performance both quantitatively and qualitatively.

1 Introduction

Thanks to the rapid development of digital photography and editing software, synthetic images created by image editing operations (*e.g.*, crop and paste, appearance adjustment) are prevalent in our daily lives. However, one crucial issue of some synthetic images is that the color and illumination characteristics of the manipulated regions are inconsistent with other regions, which could severely degrade the quality of synthetic images.

Following the definition in [26], in a synthetic image, the region incompatible with the background in terms of color or illumination is named the *inharmonious region*. Examples of synthetic images with inharmonious regions are shown in Figure 1. To avoid ambiguity, following [26], we assume that the area of the inharmonious region is smaller than 50% of the whole image. Otherwise, the background is viewed as the inharmonious region. The *inharmonious region localization* task aims to localize the inharmonious region. After the inharmonious region is localized, we can manually adjust the inharmonious region or utilize off-the-shelf image harmonization techniques to make the synthetic image more harmonious.

*Corresponding author.




Figure 1: Examples of inharmonious images and the corresponding masks.

Thus, inharmonious region localization task is necessary for image harmonization when the foreground mask is not available [10]. To the best of our knowledge, the only existing work on inharmonious region localization is DIRM [26]. However, it did not exploit the uniqueness of this task, *i.e.*, the discrepancy between the inharmonious region and the background region.

Considering the uniqueness of the inharmonious region localization task, we treat it as a clustering problem and draw inspiration from typical clustering algorithm K-means [20]. Specifically, we aim to assign all the pixels into two clusters: inharmonious cluster and background cluster. The idea of K-means is iteratively performing the following two steps. 1) *Assignment step*: for each pixel, calculate the similarity between it and all the centroids followed by assigning it to its nearest centroid. 2) *Update step*: re-calculate the centroid of each cluster based on its associated samples. Inspired by K-means [20], we design a Recurrent Self-Reasoning (RSR) module and insert it into the bottleneck of UNet [38] structure.

Now we briefly introduce our RSR module. We extract conventional feature map and style feature map from the encoder of UNet, in which the style feature map is expected to contain color and illumination information while there is no restriction for the conventional feature map. Given the style feature map and an initial inharmonious mask, we calculate the background style feature by averaging pixel-wise features within the background region. Then, we calculate a similarity map between the background style feature and the whole style feature map. This is similar to the *assignment step* in K-means because the similarity map roughly indicates which pixels belong to the background cluster. However, directly using this similarity map for assignment is inaccurate, because localizing the inharmonious region requires rich context information and the style feature may not be adequately informative. Therefore, we employ a convolutional Gate Recurrent Unit (GRU) cell [11, 12], which takes in the similarity map, current inharmonious mask, and conventional feature map to conduct self-reasoning about the location of inharmonious region and outputs an updated inharmonious mask. With a new inharmonious mask, we can update the background style feature, which is similar to re-calculating the centroids in the *update step* of K-means. We repeat the above procedure iteratively and the quality of inharmonious mask is gradually improved.

After introducing our core Recurrent Self-Reasoning (RSR) module, we elaborate on our Recurrent Self-Reasoning Network (RSRNet). As mentioned above, we insert our RSR module into the bottleneck of UNet structure. After the recurrent steps in RSR module, we feed the estimated mask from RSR to the decoder as attention guidance and the decoder outputs a refined mask. We observe that the mask from RSR can provide a compact shape of inharmonious region, while the mask from decoder can provide sharper details and edges. Since these two masks are complementary to each other, we combine them adaptively to

form our final mask. The main contributions of our work can be summarized as follows:

- We treat inharmonious region localization as a clustering problem and draw inspiration from typical clustering algorithm, which provides a new perspective for this task.
- We propose a novel Recurrent Self-Reasoning (RSR) module to gradually improve the quality of inharmonious mask. We also propose to adaptively combine the mask from RSR module and the mask from decoder, leading to further improvement.
- Extensive experiments on iHarmony4 dataset show that our RSRNet achieves the best performance both quantitatively and qualitatively.

2 Related Work

2.1 Image Manipulation Localization

Existing image manipulation localization methods can be categorized according to the type of manipulation (*e.g.*, copy-move, removal, enhancement, and splicing), which share some similarities with inharmonious region localization task. Traditional methods are mainly aiming at localizing a specific type of manipulation. These methods are generally based on the detection of specific clues or traces in the manipulated images including noise patterns [30, 31], JPEG compression differences [25, 28], and colour filter array interpolation patterns [13, 15]. Recently, there are many deep learning-based methods [9, 23, 24, 39, 45, 48, 51], which have shown impressive performance on image manipulation detection and localization. However, these methods do not specifically consider the inconsistency due to color or illumination characteristics.

2.2 Recurrent Neural Network

Recurrent Neural Network (RNN) is a type of neural networks that recursively process sequential data. The Long Short-Term Memory (LSTM) [22] and Gated Recurrent Unit (GRU) [10] are two commonly used recurrent neural networks. Although these recurrent models are originally designed to solve machine translation and sequence modeling tasks, they are then widely employed to solve computer vision tasks, such as object recognition [27], action detection [40], motion prediction [32], optical flow estimation [42], and image manipulation localization [8]. The convolutional based GRU was proposed in [10]. Different from the above works, we are the first to apply RNN for the inharmonious region localization task.

2.3 Image Harmonization and Inharmonious Region Localization

Given a synthetic image which has incompatible foreground and background due to different color and illumination, image harmonization aims to adjust the appearance of the foreground to make it compatible with the background. Recently, abundant deep learning based image harmonization methods [1, 4, 9, 11, 17, 18, 19, 29] have been proposed. To name a few, Tsai *et al.* [43] first proposed an end-to-end convolutional neural network to extract context and semantic information to generate the harmonized image. Cun and Pan [11] proposed a spatial separated attention module S²AM to separately learn the features in foreground and

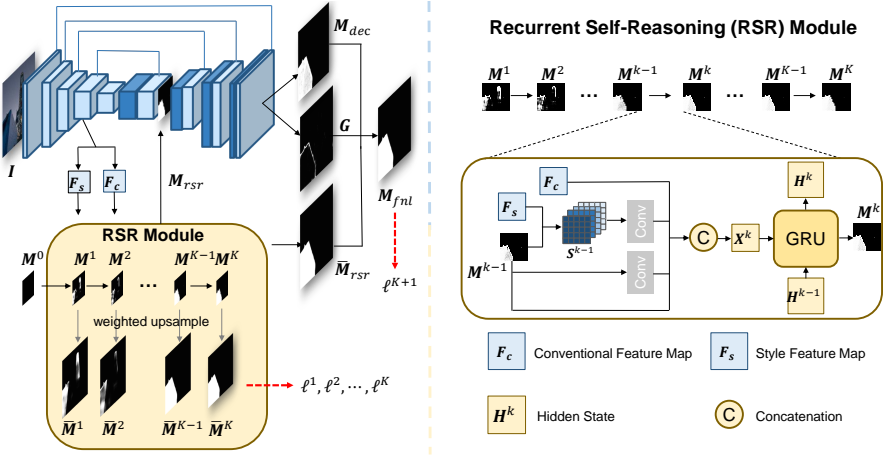


Figure 2: The left part shows the overall pipeline of our RSRNet which consists of a UNet structure and a RSR module inserted at the bottleneck. The right part shows the detailed iterative process in the RSR module.

background. Cong *et al.* [8, 9] introduced the concept of domain to distinguish incompatible regions and translate the foreground to the same domain as the background. Combining semantic features into the harmonization network [41] has shown great performance. Transformer based structure is utilized in [46] to tackle the harmonization task. Most image harmonization methods require the inharmonious region mask as input, which are usually hard to acquire in real-life scenario. Although S^2AM used the attention block to predict the inharmonious region mask, the quality of its predicted mask is far from satisfactory.

DIRL [24] is the first method focusing on inharmonious region localization, which effectively fused multi-scale features and suppressed redundant information to better localize the inharmonious region. However, it is a rather general model without exploiting the uniqueness of this task, that is, the discrepancy between the harmonious and inharmonious region.

3 Methodology

3.1 Overview

Given an input synthetic image $I \in \mathbb{R}^{H \times W \times 3}$, our goal is to estimate its inharmonious region mask $M \in \mathbb{R}^{H \times W \times 1}$. Our whole model shown in Figure 2 contains two parts: a UNet structure network and a Recurrent Self-Reasoning (RSR) module. The encoder contains five res-blocks. The first four res-blocks are adopted from ResNet34 [41] without pooling layer. After the fourth res-block, we get a feature map of shape $\frac{H}{8} \times \frac{W}{8} \times 512$. This feature map is then fed into two separate feature heads which both contain a 3×3 convolution followed by a Relu activation and a 1×1 convolution. These two feature heads extract a style feature map F_s and a conventional feature map F_c of shape $\frac{H}{8} \times \frac{W}{8} \times 256$, respectively. F_s , F_c , and the initial inharmonious mask M^0 are delivered to Recurrent Self-Reasoning (RSR) module to refine the inharmonious mask iteratively. The mask M_{rsr} from RSR is then provided for the decoder as attention guidance to output a refined mask M_{dec} . The final output M_{fnl} will

be the weighted combination of \mathbf{M}_{dec} and upsampled $\bar{\mathbf{M}}_{rsr}$. Next, we will introduce our Recurrent Self-Reasoning (RSR) module in Section 3.2 and the weighted combination of two masks in Section 3.3.

3.2 Recurrent Self-Reasoning Module (RSR)

In this module, we will iteratively update the inharmonious region mask via self-reasoning. At first, the inharmonious mask is initialized as all zeros, *i.e.*, $\mathbf{M}^0 = \mathbf{0}$. We denote the estimated mask after the k -th iteration as \mathbf{M}^k .

In the k -th iteration, we update the background style feature \mathbf{f}_{bg}^{k-1} , which is the average of style features within the background region based on the style feature map \mathbf{F}_s and current inharmonious mask from last iteration \mathbf{M}^{k-1} . This step is similar to the *update step* in K-Means algorithm where we update the centroid (the centroid is the background style feature representing the background harmonious centroid in our case) based on the current assignment of all samples (the assignment is the current estimated mask in our case). A Multi-scale Similarity Map (MSM) \mathbf{S}^{k-1} is constructed to measure how similar each pixel in \mathbf{F}_s is to the background style feature \mathbf{f}_{bg}^{k-1} . Then, MSM \mathbf{S}^{k-1} , conventional feature \mathbf{F}_c , and inharmonious mask \mathbf{M}^{k-1} will be passed through a GRU cell to produce an updated inharmonious mask \mathbf{M}^k . This corresponds to the *assignment step* in K-Means where each sample is assigned to its corresponding centroid based their similarity. The detailed process of an iteration in RSR module is shown in the right part of Figure 2.

3.2.1 Multi-scale Similarity Map

In the k -th iteration, given the current inharmonious mask \mathbf{M}^{k-1} , we calculate the background style feature by averaging the style features within the background region. Since the values of inharmonious mask are between 0 and 1 (1 for inharmonious and 0 for background), we set a threshold $\varepsilon = 0.5$ to determine pixels belonging to the background region. Then, the background style feature can be calculated as follows,

$$\mathbf{f}_{bg}^{k-1} = \frac{1}{\sum_i \delta(\mathbf{M}^{k-1}(i) < \varepsilon)} \sum_i \delta(\mathbf{M}^{k-1}(i) < \varepsilon) \mathbf{F}_s(i), \quad (1)$$

where $\mathbf{M}^{k-1}(i)$ (*resp.*, $\mathbf{F}_s(i)$) is the value of i -th pixel in \mathbf{M}^{k-1} (*resp.*, \mathbf{F}_s). $\delta(s)$ is an indicator function, *i.e.*, $\delta(s) = 1$ if s is true and 0 otherwise. $\sum_i \delta(\mathbf{M}^{k-1}(i) < \varepsilon)$ denotes the total number of background pixels.

When calculating the similarity between \mathbf{f}_{bg}^{k-1} and each pixel-level style feature, we consider the average feature in a local neighbourhood of each pixel, due to the potential noise of pixel-level features. Therefore, we use different neighborhood sizes to take multi-scale information into account. For scale l , we define $\mathcal{N}^l(i)$ as a square neighborhood centered at the i -th pixel with side length $2l + 1$. Then, we can calculate the cosine similarity between \mathbf{f}_{bg}^{k-1} and the local average feature at the i -th pixel:

$$\mathbf{S}^{k-1,l}(i) = \frac{\mathbf{f}_{bg}^{k-1} \cdot \sum_{j \in \mathcal{N}^l(i)} \mathbf{F}_s(j)}{\|\mathbf{f}_{bg}^{k-1}\|_2 \|\sum_{j \in \mathcal{N}^l(i)} \mathbf{F}_s(j)\|_2}, \quad (2)$$

in which $\mathbf{S}^{k-1,l}(i)$ is the value of the i -th pixel in the similarity matrix $\mathbf{S}^{k-1,l}$ for scale l in the k -th iteration, and \cdot means the dot product between two vectors. In practice, we take

a set of scales $l = 0, 1, 2, 3$, corresponding to the side length of neighborhood $2l + 1 = 1, 3, 5, 7$, respectively. As a result, we obtain four similarity matrices, which are concatenated to form our Multi-scale Similarity Map (MSM) \mathbf{S}^{k-1} .

Note that the style feature map \mathbf{F}_s will remain unchanged in the iterative process, but \mathbf{S}^{k-1} will be changed because \mathbf{f}_{bg}^{k-1} will be updated based on the current inharmonious mask \mathbf{M}^k in each iteration. Recall that the initial \mathbf{M}^0 is an all-zero mask and thus \mathbf{f}_{bg}^0 is the global average feature over the entire image. Since the area of the inharmonious region is smaller than background region as defined in Section 1, the initial \mathbf{f}_{bg}^0 is closer to the average feature of background region than that of inharmonious region. In the iterative process, \mathbf{f}_{bg}^{k-1} will gradually approach the true average feature of the background region, offering guidance for inharmonious region localization in the subsequent iterations.

3.2.2 GRU Based Recurrent Unit

In the k -th iteration, we have the current inharmonious mask \mathbf{M}^{k-1} , conventional feature map \mathbf{F}_c , MSM \mathbf{S}^{k-1} , and the hidden state from last iteration \mathbf{H}^{k-1} . \mathbf{M}^{k-1} and \mathbf{S}^{k-1} will go through two 3×3 convolutions, respectively. The outputs will be concatenated with \mathbf{F}_c and \mathbf{M}^{k-1} to form the input \mathbf{X}^k . After that, \mathbf{X}^k and \mathbf{H}^{k-1} are fed into a convolution GRU [14]. In each iteration, we learn the residual of the estimated mask $\Delta\mathbf{M}^k$, so $\mathbf{M}^k = \mathbf{M}^{k-1} + \Delta\mathbf{M}^k$. The update process in convolution GRU is

$$\begin{aligned}\mathbf{Z}^k &= \sigma(\text{Conv}([\mathbf{H}^{k-1}, \mathbf{X}^k], \mathbf{W}_z)), \\ \mathbf{R}^k &= \sigma(\text{Conv}([\mathbf{H}^{k-1}, \mathbf{X}^k], \mathbf{W}_r)), \\ \tilde{\mathbf{H}}^k &= \tanh(\text{Conv}([\mathbf{R}^k \odot \mathbf{H}^{k-1}, \mathbf{X}^k], \mathbf{W}_h)), \\ \mathbf{H}^k &= (1 - \mathbf{Z}^k) \odot \mathbf{H}^{k-1} + \mathbf{Z}^k \odot \tilde{\mathbf{H}}^k,\end{aligned}$$

in which $\text{Conv}(\cdot, \mathbf{W})$ means 3×3 convolution with kernel parameters \mathbf{W} , $[\cdot, \cdot]$ means concatenation, σ is the sigmoid function, \odot means element-wise product. \mathbf{Z} , \mathbf{R} , $\tilde{\mathbf{H}}$, and \mathbf{H} are the latent variables defined in GRU [14]. $\Delta\mathbf{M}^k$ is obtained by applying a 3×3 convolution followed by a 1×1 convolution on \mathbf{H}^k .

The updated mask \mathbf{M}^k is of size $\frac{H}{8} \times \frac{W}{8}$. Similar to [14], we use two convolutional layers to predict an upsample weight map of shape $\frac{H}{8} \times \frac{W}{8} \times (8 \times 8 \times 9)$ to upsample \mathbf{M}^k to the full-resolution one $\tilde{\mathbf{M}}^k$ by taking the weighted combination over the 9 neighborhood pixels. We update \mathbf{M}^k at each iteration to $\tilde{\mathbf{M}}^k$, so that the mask supervision can be employed at full resolution (see Section 3.4).

3.3 Adaptive Combination of Two Masks

We denote the mask from the last iteration in our RSR module as \mathbf{M}_{rsr} and its upsampled version is $\tilde{\mathbf{M}}_{rsr}$. \mathbf{M}_{rsr} contains the general shape and location of the inharmonious region, but lacks accurate edges and details (see examples in Supplementary) because the process is performed at the resolution of $\frac{1}{8}$. We feed \mathbf{M}_{rsr} into the decoder by concatenation to provide attention guidance. The decoder can focus on the inharmonious region provided by RSR module and integrate multi-scale encoder features through skip connection, producing the refined inharmonious mask \mathbf{M}_{dec} with sharper details and edges.

In our experiments, we observe that although \mathbf{M}_{dec} is more accurate about edges and details, it may contain some holes and misdetected regions (see examples in Supplementary). Thus, \mathbf{M}_{dec} and $\bar{\mathbf{M}}_{rsr}$ are complementary to each other and should be utilized simultaneously. For two complementary outputs, it is more beneficial to adaptively combine them instead of simply calculating the average [8, 14]. Thus, we learn a combination mask \mathbf{G} with values between 0 and 1 to adaptively combine \mathbf{M}_{dec} and $\bar{\mathbf{M}}_{rsr}$. Then the final mask \mathbf{M}_{fnl} can be obtained by

$$\mathbf{M}_{fnl} = \mathbf{G} \odot \mathbf{M}_{dec} + (1 - \mathbf{G}) \odot \bar{\mathbf{M}}_{rsr}.$$

The combination mask is expected to utilize the accurate details and edges from \mathbf{M}_{dec} as well as the general shape and location from $\bar{\mathbf{M}}_{rsr}$ to form our final mask.

3.4 Loss Function

In our RSR module, we perform K iterations and obtain K upsampled masks $\bar{\mathbf{M}}^k$ for $k = 1, \dots, K$. We choose $K = 12$ by cross-validation. We denote the loss for the k -th mask as $\ell^{(k)}$. Besides, we denote the loss for the final mask \mathbf{M}_{fnl} as $\ell^{(K+1)}$. Given the ground-truth mask \mathbf{M} and an estimated mask $\bar{\mathbf{M}}^k$, the loss $\ell^{(k)}$ is comprised of three parts following [36]: $\ell^{(k)} = \ell_{bce}^{(k)} + \ell_{ssim}^{(k)} + \ell_{iou}^{(k)}$.

$\ell_{bce}^{(k)}$ is the binary cross-entropy loss, which is commonly used in binary classification task and segmentation task. $\ell_{ssim}^{(k)}$ is the structural similarity loss, which can better represent the structural information of the ground-truth mask. $\ell_{iou}^{(k)}$ is the Intersection over Union (IoU) loss. We refer readers to [36] for details about these three losses. Our total loss is a weighted sum of losses for all the estimated masks:

$$\mathcal{L} = \ell^{(K+1)} + \sum_{k=1}^K \lambda^{K-k} \ell^{(k)}, \quad (3)$$

where λ is a weight factor set as 0.8 following [14]. The loss weights for the masks $\bar{\mathbf{M}}^k$ are increasing exponentially since the mask is initially coarse and gradually getting more accurate, and we assign higher weights to the more accurate masks.

4 Experiments

4.1 Dataset and Evaluation Metrics

Following [26], we conduct our experiments on the the image harmonization dataset **iHarmony4** [8], which contains inharmonious images with corresponding masks and harmonious images. iHarmony4 [8] consists of four sub-datasets: HCOCO, HFlickr, HAdobe5K, and Hday2night. The iHarmony4 dataset is suitable for inharmonious region localization task since the inconsistency between foreground and background is mainly due to incompatible color or illumination [26]. For HCOCO and HFlickr sub-datasets, the inharmonious images are generated by applying color transfer methods [14, 35, 47, 46] to transfer the foreground color of real images. For HAdobe5K sub-dataset, real images are retouched by five professional photographers to obtain the corresponding inharmonious images. For Hday2night sub-dataset, the inharmonious image is made by overlaying the foreground of a real image with the corresponding region in another image, which has the same scene captured under a

different condition. To avoid the ambiguity of the definition of inharmonious region, we follow [26] to only keep the images with foreground area larger than 50%. We follow the same train-test split as [26] to get 64255 training images and 7237 test images. Following [26], we adopt the evaluation metrics including Average Precision (AP), F_1 score, and Intersection over Union (IoU).

4.2 Implementation Details

We adopt the first four blocks from ImageNet-pretrained ResNet34 [21] as the first four blocks in our encoder, and the last block in our encoder is similar to the first block. We also use 3×3 convolution to replace the 7×7 convolution and remove the pooling layer to keep the resolution and retain more details in the shallow layer. We use Pytorch [4] to implement our model. We use Adam optimizer with $\beta_1 = 0.9$, $\beta_2 = 0.999$, weight decay = $1e-4$, and initial learning rate = $1e-4$. We train our model with batch size 32 for 60 epochs on a Ubuntu 18.04.4 machine with 4 GeForce GTX TITAN X GPUs with the learning rate reduced by 0.5 after 30, 40, 50, 55 epochs respectively. The random seed set for numpy and Pytorch is 42.

4.3 Comparison with the State-of-the-art

To the best of our knowledge, DURL [26] is the only work focusing on the inharmonious region localization task. Therefore, following [26], we also compare our model with the state-of-the-art methods from other related fields. We choose three groups of baselines for comparison: 1) Popular segmentation networks: UNet [58], DeepLabv3 [6], HRNet-OCR [49], SegFormer [4]. 2) Image manipulation localization methods: MFCN [49], MantraNet [45], MAGritte [24], H-LSTM [8], SPAN [23]. 3) Salient object detection methods: F3Net [24], GATENet [50], MINet [53]

4.3.1 Quantitative Evaluation

We report AP, F_1 , and IoU of all methods for each sub-datasets in Table 1. Following [26], we use ResNet34 as backbone for ResNet based methods, HRNet30 for HRNet-OCR, and SegFormer-B3 for SegFormer. Based on Table 1, our method achieves the best performance on the whole dataset. Our method beats the best method MAGritte in the image manipulation localization field by a large margin. The general segmentation methods and MINet in the salient object detection field also show competitive performance, but they are still worse than our model. For the strongest baseline DURL, our model has 3.19% (*resp.*, 3.96%, 5.57%) improvement for AP, F_1 , and IoU.

In addition, we report model related statistics including amount of parameters, model inference speed, and GFlops of our model and other strong baselines in the Supplementary.

4.3.2 Qualitative Evaluation

To better verify the advantage of our model, we visualize the predicted masks from our model and top five baseline methods in Figure 3. It can be seen that in various challenging scenarios including occlusion, fine-grained structure, and small object, our RSRNet can better localize the inharmonious region with clear and sharp boundary.

Method	HCOCO			HAdobe5k			HFlickr			Hday2night			All		
	AP \uparrow	$F_1\uparrow$	IoU \uparrow	AP \uparrow	$F_1\uparrow$	IoU \uparrow	AP \uparrow	$F_1\uparrow$	IoU \uparrow	AP \uparrow	$F_1\uparrow$	IoU \uparrow	AP \uparrow	$F_1\uparrow$	IoU \uparrow
UNet	68.11	0.5869	56.57	89.26	0.8380	80.85	80.72	0.7683	74.58	35.74	0.2362	19.60	74.90	0.6717	64.74
DeepLabv3	69.09	0.6070	58.21	90.20	0.8591	81.56	80.01	0.7698	74.91	35.87	0.2550	21.38	75.69	0.6902	66.01
HRNet-OCR	68.89	0.5981	57.69	89.63	0.8387	80.98	79.62	0.7489	74.55	34.98	0.2477	21.34	75.33	0.6765	65.49
SegFormer	72.46	0.6578	58.78	89.43	0.8531	80.44	85.19	0.7986	75.02	45.16	0.3856	32.75	78.05	0.7249	66.55
MFCN	37.36	0.3030	25.18	62.75	0.5365	36.63	49.89	0.4209	28.34	19.71	0.1426	11.88	45.63	0.3794	28.54
MantraNet	56.55	0.4811	41.04	81.07	0.7510	68.50	67.52	0.6302	58.51	28.88	0.2019	16.71	64.22	0.5691	50.31
MAGritte	64.75	0.6058	51.77	85.50	0.8630	76.36	75.02	0.7725	70.25	31.20	0.2549	17.05	71.16	0.6907	60.14
H-LSTM	52.29	0.4336	37.81	77.62	0.7056	65.19	63.12	0.5918	54.93	26.59	0.1977	15.91	60.21	0.5239	47.07
SPAN	58.41	0.4906	45.07	82.57	0.7786	72.49	69.22	0.6510	62.20	29.58	0.2171	19.41	65.94	0.5850	54.27
F3Net	54.17	0.4703	40.03	74.31	0.6944	60.08	72.53	0.6582	59.31	30.08	0.2563	20.83	61.46	0.5506	47.48
GATENet	55.07	0.4568	38.89	75.19	0.6634	59.18	74.13	0.6256	57.51	30.98	0.2174	19.38	62.43	0.5296	46.33
MINet	71.74	0.6022	55.79	89.58	0.8379	77.23	83.86	0.7761	72.51	37.82	0.2710	19.38	77.51	0.6822	63.04
DIREL	74.25	0.6701	60.85	92.16	0.8801	84.02	84.21	0.7786	73.21	38.74	0.2396	20.11	80.02	0.7317	67.85
RSRNet	78.42	0.7131	65.85	93.10	0.8901	85.96	87.11	0.8048	76.84	47.34	0.3028	26.34	82.57	0.7607	71.63

Table 1: Performance comparison with different methods on the iHarmony4 dataset. \uparrow means the larger, the better. The best results are denoted in **bold**.




Figure 3: Qualitative comparison of our model with other state-of-the-art methods from related fields. GT is the ground-truth inharmonious mask.

4.4 Ablation Studies

Row 1 in Table 2 is the basic UNet including an encoder and decoder. The result is slightly different from UNet in Table 1 because the original UNet method only has ℓ_{bce} . In row 2, we insert our RSR module into the bottleneck and evaluate the mask \mathbf{M}_{dec} output from decoder. Comparing row 1 & 2, the quality of \mathbf{M}_{dec} is significantly improved after inserting RSR module. In row 3, we only iteratively update similarity matrix without using other information, which is closer to K-means clustering. Specifically, we directly use a single similarity matrix with scale $l = 0$ as the updated inharmonious mask and update it for 12 iterations. The results demonstrate that the similarity map is not informative enough and imposing supervision on it may degrade the performance. In row 4, we simply average $\bar{\mathbf{M}}_{rsr}$ and \mathbf{M}_{dec} as \mathbf{M}_{fnl} , which is better than row 2 but much worse than our full method in row 9. The results show that it is necessary to combine $\bar{\mathbf{M}}_{rsr}$ and \mathbf{M}_{dec} and sophisticated combination strategy performs more favorably.

#	UNet	RSR	Mask Combination	Evaluation		
				AP(%) \uparrow	$F_1 \uparrow$	IoU(%) \uparrow
1	+			78.35	0.7130	65.99
2	+	+		80.53	0.7426	69.72
3	+	only similarity map		77.24	0.6973	64.55
4	+	+	simple average	80.82	0.7431	69.69
5	+	+ (w/o GRU)	+	81.66	0.7564	70.52
6	+	+ (w/o MSM)	+	80.62	0.7451	69.81
7	+	+ (w/o F_c)	+	81.08	0.7440	69.93
8	+	+ (w/o weighted upsample)	+	81.24	0.7515	70.57
9	+	+	+	82.57	0.7607	71.63

Table 2: Ablation study on different components and combinations in our method. See Section 4.4 for the detailed explanation.

#	RSR	\bar{M}_{rsr}			\bar{M}_{dec}			\bar{M}_{fnl}		
		AP(%) \uparrow	$F_1 \uparrow$	IoU(%) \uparrow	AP(%) \uparrow	$F_1 \uparrow$	IoU(%) \uparrow	AP(%) \uparrow	$F_1 \uparrow$	IoU(%) \uparrow
1	w/o GRU	81.29	0.7495	70.21	77.88	0.7526	70.38	81.66	0.7564	70.52
2	w/o MSM	80.27	0.7398	68.92	75.57	0.7400	69.15	80.62	0.7451	69.81
3	w/o F_c	80.74	0.7392	69.11	76.79	0.7413	69.58	81.08	0.7440	69.93
4	w/o weighted upsample	80.57	0.7423	69.26	76.81	0.7486	70.20	81.24	0.7517	70.57
5	full module	82.19	0.7547	70.62	78.32	0.7591	71.42	82.57	0.7607	71.63

Table 3: The evaluation results of three masks \bar{M}_{rsr} , \bar{M}_{dec} , \bar{M}_{fnl} using ablated RSR module.

In row 5, we remove the GRU cell and only use two convolution layers to output the mask. In row 6, we remove MSM, and only feed F_c and \bar{M}^{k-1} into the GRU cell. In row 7, we replace F_c with F_s . In row 8, we use bilinear upsample instead of the weighted upsample for \bar{M}^k . The results in row 5-8 are all worse than our full method, which justifies the necessity and effectiveness of our RSR module.

For each experiment from row 5 to row 9 in Table 2, we further report the results of \bar{M}_{rsr} and \bar{M}_{dec} to be compared with \bar{M}_{fnl} in Table 3. We can see that \bar{M}_{rsr} usually has a higher AP while \bar{M}_{dec} has higher F_1 and IoU. The reason is that the AP is an average value computed at different thresholds while F_1 and IoU are both at only one threshold 0.5, and those uncertain or misdetected region in \bar{M}_{dec} will have worse performance when the threshold is high. Combining these two complementary masks leads to better performance of \bar{M}_{fnl} .

5 Conclusion

In this paper, we propose a Recurrent Self-Reasoning based Network (RSRNet) to achieve inharmonic region localization. Inspired by K-means algorithm, we design a recurrent module to iteratively reason about the inharmonic region. We also design an adaptive combination mask to selectively combine the mask from our RSR module and that from the decoder. Experiments on iHarmony4 demonstrate the superiority of our proposed model.

Acknowledgement

The work was supported by the Shanghai Municipal Science and Technology Major/Key Project, China (2021SHZDZX0102, 20511100300) and National Natural Science Foundation of China (Grant No. 61902247).

References

- [1] Nicolas Ballas, Li Yao, Chris Pal, and Aaron Courville. Delving deeper into convolutional networks for learning video representations. *arXiv preprint arXiv:1511.06432*, 2015.
- [2] Zhongyun Bao, Chengjiang Long, Gang Fu, Daquan Liu, Yuanzhen Li, Jiaming Wu, and Chunxia Xiao. Deep image-based illumination harmonization. In *CVPR*, 2022.
- [3] Jawadul H Bappy, Cody Simons, Lakshmanan Nataraj, BS Manjunath, and Amit K Roy-Chowdhury. Hybrid lstm and encoder–decoder architecture for detection of image forgeries. *IEEE Transactions on Image Processing*, 28(7):3286–3300, 2019.
- [4] Junyan Cao, Wenyan Cong, Li Niu, Jianfu Zhang, and Liqing Zhang. Deep image harmonization by bridging the reality gap. In *BMVC*, 2022.
- [5] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Re-thinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017.
- [6] Yanhua Cheng, Rui Cai, Zhiwei Li, Xin Zhao, and Kaiqi Huang. Locality-sensitive deconvolution networks with gated fusion for rgb-d indoor semantic segmentation. In *CVPR*, 2017.
- [7] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014.
- [8] Wenyan Cong, Jianfu Zhang, Li Niu, Liu Liu, Zhixin Ling, Weiyuan Li, and Liqing Zhang. Dovenet: Deep image harmonization via domain verification. In *CVPR*, 2020.
- [9] Wenyan Cong, Li Niu, Jianfu Zhang, Jing Liang, and Liqing Zhang. Bargainnet: Background-guided domain translation for image harmonization. In *ICME*, 2021.
- [10] Wenyan Cong, Xinhao Tao, Li Niu, Jing Liang, Xuesong Gao, Qihao Sun, and Liqing Zhang. High-resolution image harmonization via collaborative dual transformations. In *CVPR*, 2022.
- [11] Xiaodong Cun and Chi-Man Pun. Improving the harmony of the composite image by spatial-separated attention module. *IEEE Transactions on Image Processing*, 29: 4759–4771, 2020.
- [12] Tianyu Ding, Luming Liang, Zhihui Zhu, and Ilya Zharkov. Cdfi: Compression-driven network design for frame interpolation. In *CVPR*, 2021.
- [13] Ahmet Emir Dirik and Nasir Memon. Image tamper detection based on demosaicing artifacts. In *ICIP*, 2009.
- [14] Ulrich Fecker, Marcus Barkowsky, and André Kaup. Histogram-based prefiltering for luminance and chrominance compensation of multiview video. *IEEE Transactions on Circuits and Systems for Video Technology*, 18(9):1258–1267, 2008.

- [15] Pasquale Ferrara, Tiziano Bianchi, Alessia De Rosa, and Alessandro Piva. Image forgery localization via fine-grained analysis of cfa artifacts. *IEEE Transactions on Information Forensics and Security*, 7(5):1566–1577, 2012.
- [16] Zonghui Guo, Dongsheng Guo, Haiyong Zheng, Zhaorui Gu, Bing Zheng, and Junyu Dong. Image harmonization with transformer. In *ICCV*, 2021.
- [17] Zonghui Guo, Haiyong Zheng, Yufeng Jiang, Zhaorui Gu, and Bing Zheng. Intrinsic image harmonization. In *CVPR*, 2021.
- [18] Yucheng Hang, Bin Xia, Wenming Yang, and Qingmin Liao. Scs-co: Self-consistent style contrastive learning for image harmonization. In *CVPR*, 2022.
- [19] Guoqing Hao, Satoshi Iizuka, and Kazuhiro Fukui. Image harmonization with attention-based deep feature modulation. In *BMVC*, 2020.
- [20] John A Hartigan and Manchek A Wong. Algorithm as 136: A k-means clustering algorithm. *Journal of the royal statistical society. series c (applied statistics)*, 28(1): 100–108, 1979.
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [22] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [23] Xuefeng Hu, Zhihan Zhang, Zhenye Jiang, Syomantak Chaudhuri, Zhenheng Yang, and Ram Nevatia. Span: Spatial pyramid attention network for image manipulation localization. In *ECCV*, 2020.
- [24] Vladimir V Kniaz, Vladimir Knyaz, and Fabio Remondino. The point where reality meets fantasy: Mixed adversarial generators for image splice detection. 2019.
- [25] Weihai Li, Yuan Yuan, and Nenghai Yu. Passive detection of doctored jpeg image via block artifact grid extraction. *Signal Processing*, 89(9):1821–1829, 2009.
- [26] Jing Liang, Li Niu, and Liqing Zhang. Inharmonious region localization. In *ICME*, 2021.
- [27] Ming Liang and Xiaolin Hu. Recurrent convolutional neural network for object recognition. In *CVPR*, 2015.
- [28] Zhouchen Lin, Junfeng He, Xiaou Tang, and Chi-Keung Tang. Fast, automatic and fine-grained tampered jpeg image detection via dct coefficient analysis. *Pattern Recognition*, 42(11):2492–2501, 2009.
- [29] Jun Ling, Han Xue, Li Song, Rong Xie, and Xiao Gu. Region-aware adaptive instance normalization for image harmonization. In *CVPR*, 2021.
- [30] Siwei Lyu, Xunyu Pan, and Xing Zhang. Exposing region splicing forgeries with blind local noise estimation. *International journal of computer vision*, 110(2):202–221, 2014.

- [31] Babak Mahdian and Stanislav Saic. Using noise inconsistencies for blind image forensics. *Image and Vision Computing*, 27(10):1497–1503, 2009.
- [32] Julieta Martinez, Michael J Black, and Javier Romero. On human motion prediction using recurrent neural networks. In *CVPR*, 2017.
- [33] Youwei Pang, Xiaoqi Zhao, Lihe Zhang, and Huchuan Lu. Multi-scale interactive network for salient object detection. In *CVPR*, 2020.
- [34] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017.
- [35] François Pitié, Anil C Kokaram, and Rozenn Dahyot. Automated colour grading using colour distribution transfer. *Computer Vision and Image Understanding*, 107(1-2):123–137, 2007.
- [36] Xuebin Qin, Zichen Zhang, Chenyang Huang, Chao Gao, Masood Dehghan, and Martin Jagersand. Basnet: Boundary-aware salient object detection. In *CVPR*, 2019.
- [37] Erik Reinhard, Michael Adhikhmin, Bruce Gooch, and Peter Shirley. Color transfer between images. *IEEE Computer graphics and applications*, 21(5):34–41, 2001.
- [38] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 2015.
- [39] Ronald Salloum, Yuzhuo Ren, and C-C Jay Kuo. Image splicing localization using a multi-task fully convolutional network (mfcn). *Journal of Visual Communication and Image Representation*, 51:201–209, 2018.
- [40] Bharat Singh, Tim K Marks, Michael Jones, Oncel Tuzel, and Ming Shao. A multi-stream bi-directional recurrent neural network for fine-grained action detection. In *CVPR*, 2016.
- [41] Konstantin Sofiiuk, Polina Popenova, and Anton Konushin. Foreground-aware semantic representations for image harmonization. In *WACV*, 2021.
- [42] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *ECCV*, 2020.
- [43] Yi-Hsuan Tsai, Xiaohui Shen, Zhe Lin, Kalyan Sunkavalli, Xin Lu, and Ming-Hsuan Yang. Deep image harmonization. In *CVPR*, 2017.
- [44] Jun Wei, Shuhui Wang, and Qingming Huang. F³net: Fusion, feedback and focus for salient object detection. In *AAAI*, 2020.
- [45] Yue Wu, Wael AbdAlmageed, and Premkumar Natarajan. Mantra-net: Manipulation tracing network for detection and localization of image forgeries with anomalous features. In *CVPR*, 2019.
- [46] Xuezhong Xiao and Lizhuang Ma. Color transfer in correlated color space. In *VRCAI*, pages 305–309, 2006.

-
- [47] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. *NeurIPS*, 2021.
 - [48] Chao Yang, Huizhou Li, Fangting Lin, Bin Jiang, and Hao Zhao. Constrained r-cnn: A general image manipulation detection model. In *ICME*, 2020.
 - [49] Yuhui Yuan, Xilin Chen, and Jingdong Wang. Object-contextual representations for semantic segmentation. 2020.
 - [50] Xiaoqi Zhao, Youwei Pang, Lihe Zhang, Huchuan Lu, and Lei Zhang. Suppress and balance: A simple gated network for salient object detection. In *ECCV*, 2020.
 - [51] Peng Zhou, Bor-Chun Chen, Xintong Han, Mahyar Najibi, Abhinav Shrivastava, Ser-Nam Lim, and Larry Davis. Generate, segment, and refine: Towards generic manipulation segmentation. In *AAAI*, 2020.