


Contrastive Semantic-Guided Image Smoothing Network

Jie Wang^{1,2}, Yongzhen Wang¹, Yidan Feng³, Lina Gong¹, Xuefeng Yan¹, Haoran Xie^{4†}, Fu Lee Wang⁵, Mingqiang Wei^{1,2†}

¹School of Computer Science and Technology, Nanjing University of Aeronautics and Astronautics, Nanjing, China

²Shenzhen Research Institute, Nanjing University of Aeronautics and Astronautics, Shenzhen, China

³Centre for Smart Health, School of Nursing, Hong Kong Polytechnic University, Hong Kong, China

⁴Department of Computing and Decision Sciences, Lingnan University, Hong Kong, China

⁵School of Science and Technology, Hong Kong Metropolitan University, Hong Kong, China

Abstract

Image smoothing is a fundamental low-level vision task that aims to preserve salient structures of an image while removing insignificant details. Deep learning has been explored in image smoothing to deal with the complex entanglement of semantic structures and trivial details. However, current methods neglect two important facts in smoothing: 1) naive pixel-level regression supervised by the limited number of high-quality smoothing ground-truth could lead to domain shift and cause generalization problems towards real-world images; 2) texture appearance is closely related to object semantics, so that image smoothing requires awareness of semantic difference to apply adaptive smoothing strengths. To address these issues, we propose a novel Contrastive Semantic-Guided Image Smoothing Network (CSGIS-Net) that combines both contrastive prior and semantic prior to facilitate robust image smoothing. The supervision signal is augmented by leveraging undesired smoothing effects as negative teachers, and by incorporating segmentation tasks to encourage semantic distinctiveness. To realize the proposed network, we also enrich the original VOC dataset with texture enhancement and smoothing labels, namely VOC-smooth, which first bridges image smoothing and semantic segmentation. Extensive experiments demonstrate that the proposed CSGIS-Net outperforms state-of-the-art algorithms by a large margin. Code and dataset are available at <https://github.com/wangjie6866/CSGIS-Net>.

CCS Concepts

• *Computing methodologies* → *Image processing*;

1. Introduction

Image smoothing is a fundamental technique widely used in various computer vision and graphics tasks, aiming to preserve salient structures while smoothing out insignificant details. Numerous algorithms have been proposed and applied to many downstream tasks such as texture removal [XYXJ12, SSW17], detail enhancement [FAR07], and clip-art restoration [WWH06, XLXJ11], which benefit from well-smoothed images. However, due to the complicated entanglement of structures and trivial details in natural images, image smoothing is still a challenging task that has attracted considerable attention in computer vision and computer graphics communities. Previous prior-based methods can be mainly divided into two classes: kernel-based and optimization-based methods. The kernel-based methods (e.g., Bilateral Filter [TM98] and Guided Filter [HST12]) utilize local image statistics to explore the effective kernels, which may introduce inconsistent smoothing effects in different local areas. The optimization-based methods iteratively optimize the smoothing result under the global regularization

terms, which trade smoothing quality with time and memory. These methods depend on manually designed image priors that describe low-level image features, which suffer from tedious parameter tuning and can be vulnerable to special input cases.

With the development of deep learning, data-driven approaches [XRY*15, ZLJ*19] have also been explored on image smoothing. The deep hierarchical network structure has a natural advantage in capturing semantic cues for extracting meaningful image structures. However, the supervision signal for the smoothing task is hard to define. [XLXJ11] only approximates the smoothing results from existing operators, and [LYB18] blends cartoon images with textures, which can cause over/under-smoothing on natural images due to the domain shift problem. These methods handle easy cases (salient structures and weak details) with better efficiency, but fail to tackle the coexistence of weak semantic structures and strong trivial details. Differently, [FDY*21, ZLJ*19] build high-quality ground-truth domains by delicately designed screening procedures, such as comprehensive user studies and patch-level quality measurements. Although achieved notable gain in visual performance, there are still two ignored problems: 1) since the desired smoothing ground-truths are hard to collect, current methods are biased to

† Co-corresponding authors: H. Xie and M. Wei.

the training domain when only supervised by the limited number of ground-truth data. 2) Current methods are dominated by pixel-wise constraints, which encourage smooth recovery of the overall structure, but neglect the fact that texture details are closely related to object semantics. For instance, the textures of grass and wall are totally different, and thus require different smoothing strengths.

To resolve the aforementioned problems, we propose to introduce contrastive learning and semantic guidance into the image smoothing problem. For the first problem, we augment the supervision signal by the undesired smoothing effects (i.e. over-smoothed images) as negative samples, where the predicted results are pushed away from those negative samples in a contrastive manner to encourage more robust representation learning for real-world image smoothing. For the second problem, we enhance the semantic awareness of the smoothing network by incorporating edge detection and semantic segmentation as auxiliary tasks. Edge detection emphasizes the global semantics and helps to localize the meaningful structure areas to be preserved. Semantic segmentation requires distinctiveness in different types of objects and consistency within the same type, and thus encourages the smoothing network to be semantic adaptive. Also, to support the supervision of the semantic branch, we build a dataset named VOC-smooth based on PASCAL VOC 2012 [EEVG*15] with both smoothing and segmentation labels, which makes the first trial to bridge image smoothing and semantic segmentation.

In summary, our main contributions are as follows:

- We propose a contrastive semantic-guided image smoothing network named CSGIS-Net that benefits from a contrastive learning paradigm and high-level semantic information, thereby enhancing the smoothing capacity of the model.
- We propose to build the VOC-smooth dataset, which contains 15,000 texture-enhanced images with both high-quality smoothing ground-truths and semantic segmentation labels.
- Extensive experiments demonstrate that the proposed CSGIS-Net performs favorably against the state-of-the-art image smoothing approaches.

2. Related work

2.1. Image Smoothing Methods

Kernel-based methods: Kernel-based approaches mainly apply the local image statistics to explore the effective kernels. For edge-preserving smoothing, Anisotropic Diffusion (AD) [PM90] defines a new scale-space and achieves a diffusion process iteratively. Bilateral Filter [TM98] has been widely used in various tasks, weighted by the Gaussian in both spatial and intensity distance. Gastal et al. propose the adaptive manifold filter (AMF) [GO12] to handle high-dimensional data as an extension to BF. They further introduce Domain Transform filtering (DT) [GO11] in real-time. Lee et al. [LJKL17] propose a gradient-guided algorithm and adaptively smooth image gradients via the interval gradient operator. However, the aforementioned local filters often result in gradient reversals and artifacts (such as halos) [FFLS08].

Optimization-based methods: To address the issue caused by the local filters, optimization-based approaches are designed, which

attempt to solve a global objective function consisting of a data term and a regularization term. TV regularization [ROF92] and its variant RTV (relative total variation) [XYXJ12] aim at preserving edges and the latter focuses on texture. The weighted least square smoothing (WLS) [FFLS08] is proposed for progressive coarsening of images and multi-scale detail extraction. [BZCC10] presents an optimization framework for image processing (e.g. smoothing), which mainly utilizes the gradient-domain analysis. Furthermore, Ham et al. solve a non-convex optimization problem, by combining a static guidance weight with Welsch's penalty [HW77], namely SD filter (SDF) [HCP17]. The latest work [LZL*21] introduces a truncated Huber penalty function and establishes a generalized framework, which is capable to achieve diverse smoothing natures under different parameter settings.

Data-driven methods: Recently, many deep learning-based methods have been proposed and achieved remarkable performance. Xu et al. [XRY*15] propose a deep edge-aware filter to approximate different operators based on a gradient domain training procedure. Shen et al. [SCTJ17] present a pyramid structure with an extended receptive field. In contrast, Zhu et al. [ZLJ*19] achieve an end-to-end architecture utilizing the ResNet [HZRS16] and VDCNN [CSBL16] as backbone. However, most of them fail to generate a novel smoothing effect since they are only supervised by the results of existing methods. Therefore, Fan et al. [FYW*18] present an unsupervised learning paradigm to learn an energy function. More recently, Fan et al. [FYH*17] and Feng et al. [FDY*21] combine the edge map, the former transfers edge prediction to a cascaded convolutional network, and the latter proposes a joint edge detection and image smoothing network.

2.2. Image Smoothing Benchmark

Several datasets such as BSDS300 [MFTM01], DIV2K [AT17], MIT5K [BPCD11] have been widely used in existing algorithms. Unfortunately, due to the lack of corresponding ground-truth images for smoothing, the evaluation is quite subjective. To address this issue, Zhu et al. [ZLJ*19] generate a set of ground-truths through traditional operators. Xu et al. [XLH*20] build a synthetic dataset (NKS) by blending cartoon images and common textures. However, they still fail to tackle the real-world data since the highly biased ground-truths. For evaluation, pixel-wised Peak-Signal-to-Noise-Ratio (PSNR) [WBSS04] and image-wised Structural Similarity index (SSIM) [WBSS04] are often employed to measure the difference between the reconstructed image and the ground-truth.

2.3. Contrastive Learning

Contrastive learning, fully exploring images' supervisory information, has been widely adopted in self-supervised learning [Hen20, TKI20]. For a given anchor, contrastive learning attempts to pull it together with positive pairs and push it far away from the negative ones. Previous works have realized such an idea in various computer vision tasks and achieved notable performance gains. The Noise Contrast Estimation (NCE) [GH10], InfoNCE [VdOLV18], and triplet loss [HBL17] are usually adopted as loss functions. Recently, [PEZZ20] demonstrates the effectiveness of contrastive learning for image reconstruction. However, few studies have fo-

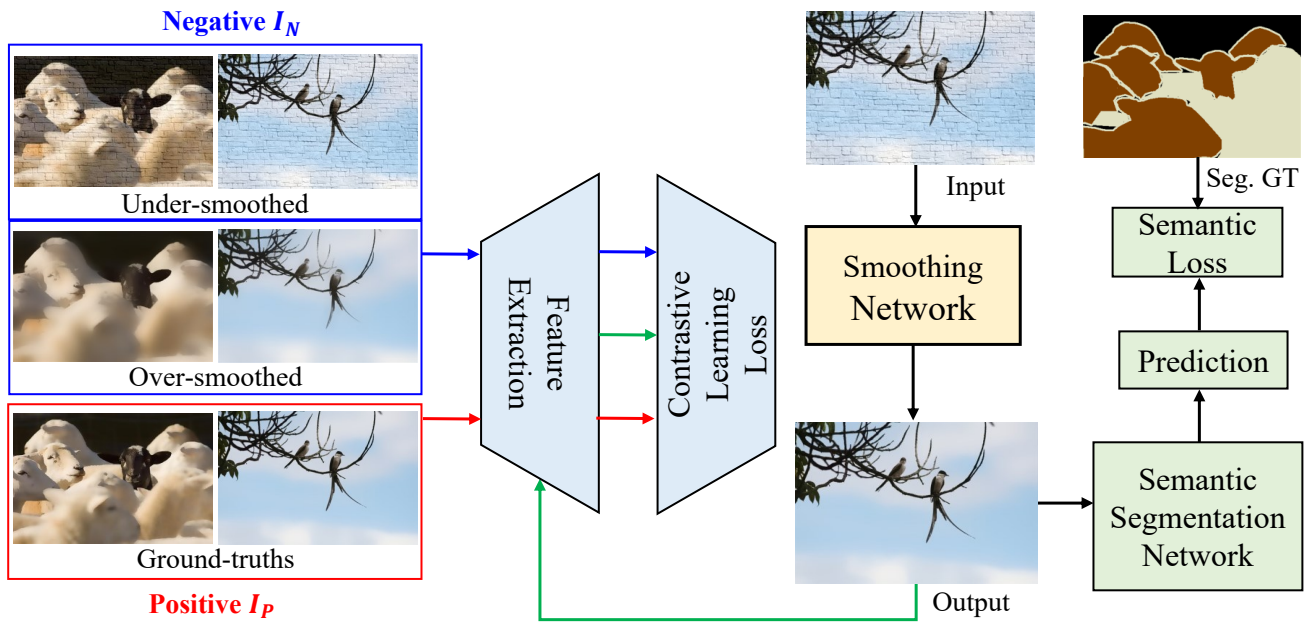


Figure 1: Pipeline of our CSGIS-Net. It includes a smoothing network, a semantic segmentation network and a contrastive module.

cused on image smoothing since it’s hard to define the positive/negative samples.

3. Proposed Method

In real-world smoothing scenes, strong details and weak structures of images are often intertwined, making the image smoothing task more challenging. Existing smoothing efforts are prone to cause over/under-smoothed problems since these algorithms are unable to understand the smoothing operation in a semantic manner. To address this issue, We regard the high-level semantic information as important prior knowledge to help the model better distinguish the texture and structure features. Meanwhile, the contrastive learning paradigm is adopted to augment the supervision signal by the unexplored information behind the negative samples (i.e. over/under-smoothed images), preventing the model’s output from falling into over/under-smoothed situations. To this end, a novel contrastive semantic-guided image smoothing network (CSGIS-Net) is proposed that leverages both contrastive learning paradigm and semantic knowledge for image smoothing.

We first present the overview of the proposed CSGIS-Net (see Fig. 1). Then the contrastive constraint and semantic structure preservation modules are described to exhibit how they work in our network. Finally, the loss function used in our training phase is introduced.

3.1. Overview

Fundamentally, the image smoothing task can be described as seeking a mapping function $\varphi(\cdot)$ from I to O , where $O = \varphi(I)$. It is an ill-posed problem since the smooth ground-truth is hard to define.

In our work, to facilitate robust image smoothing, we explore the data’s own supervision information from two different priors: one is from the contrastive learning paradigm, which introduces the negative samples I_N and positive ones I_P to alleviate the domain shift. The other is the semantic guidance information, i.e. the ground-truth and the prediction of semantic segmentation (S_{gt}, S_{pred}) and edge detection (E_{gt}, E_{pred}). It is beneficial to localize what to preserve and enhance the semantic awareness of the network. Hence, the mapping function can be expressed as:

$$O = \varphi(I, I_N, I_P, S_{gt}, S_{pred}, E_{gt}, E_{pred}) \quad (1)$$

We propose a Contrastive Semantic-guided Image Smoothing Network (CSGIS-Net) to preserve salient structures while removing insignificant details. Fig. 1 exhibits the training process of CSGIS-Net and demonstrates that the proposed network is able to remove the synthetic texture, even if the input images are meaningless in the real world. That is, CSGIS-Net is able to learn to distinguish the representations between unnecessary textures (e.g., bricks) and meaningful semantic structures. Hence, the model is well generalized to real-world scenes and deals with challenging textures. As demonstrated in Fig. 1, CSGIS-Net contains an image smoothing network, a contrastive learning strategy and a segmentation network. Specifically, we employ JESS [FDY*21] as our smoothing network and produce the result by bridging the edge detection. For the feature extractor, we adopt the VGG16 [SZ14] pre-trained on ImageNet [DDS*09]. For the semantic segmentation module, Deeplabv3+ [CZP*18] is employed.

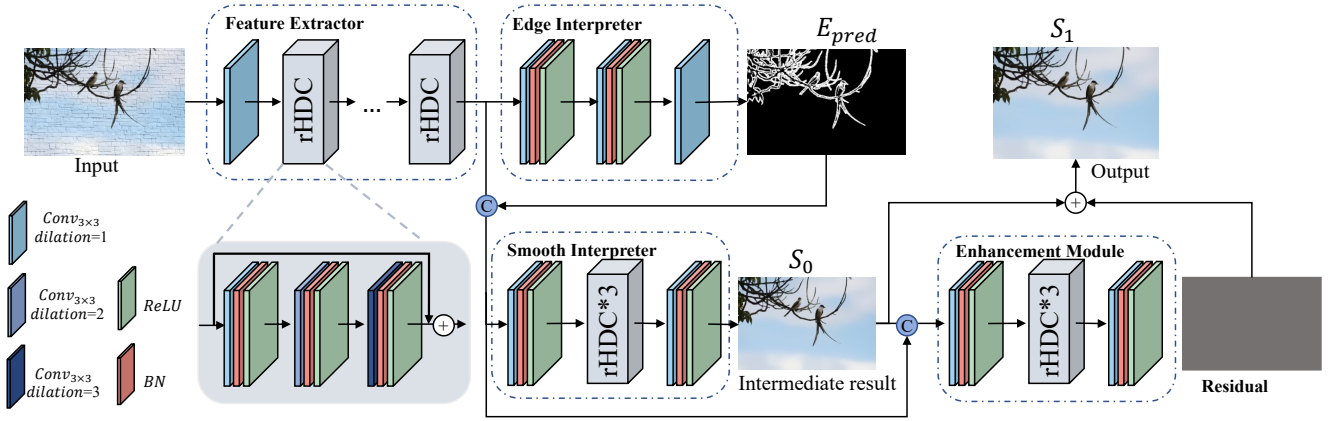


Figure 2: The architecture of the smoothing network.

3.2. Image Smoothing Network

Our smoothing network comprises a feature extractor, two interpreters (i.e., image smoothing and edge detection), and a feature enhancement module inspired by [DWW*20] (see Fig. 2). Given an image I , we first apply the feature extractor to extract the image-level feature maps. After that, the extracted features are fed into the edge interpreter for edge prediction, which is formulated as:

$$\begin{aligned} F &= \text{feature_extractor}(I) \\ E_{pred} &= \text{edge_interpreter}(F) \end{aligned} \quad (2)$$

where the feature extractor consists of a convolutional layer and series residual hybrid dilated convolution (rHDC) blocks [FDY*21]. Note that, the ground-truth edge map generated from smoothing is regarded as semantic knowledge. Then the smooth interpreter receives the edge prediction E_{pred} and the global features F and fuses them in a concatenation way, it can be expressed as:

$$S_0 = \text{smooth_interpreter}(F || E_{pred}) \quad (3)$$

where $\text{smooth_interpreter}$ denotes rHDC operations, $||$ denotes the concatenation. The outputs of the aforementioned modules will be delivered to the enhancement module, which is constructed by stacking multiple rHDC blocks. The final result is obtained as:

$$S_1 = S_0 + \text{enhance_module}(F || E_{pred} || S_0) \quad (4)$$

where S_0 denotes the intermediate result and S_1 denotes the final smoothing prediction.

The rHDC block uses the dilated convolution [CPK*18] in a cascaded way. It gains a dense receptive field to learn the high-level image semantics, which highly improves the performance. In addition, the two interpreters perform in both parallel and cascaded manners, thereby the output of the edge interpreter will affect the parameter updating process of the smooth interpreter. Furthermore, the intermediate network parameters of the edge interpreter would not be ignored.

The loss function consists of an edge detection term, a smoothing term, and a regularization term. For the edge term, it is formulated

via the weighted cross entropy:

$$\begin{aligned} L_{edge}(E_{pred}, E_{gt}) &= -\alpha \sum_{i \in E^+} \log(p) - (1 - \alpha) \sum_{i \in E^-} \log(1 - p) \\ \alpha &= \frac{|E^-|}{|E|}, p = \Pr(e_j = 1 | E) \end{aligned} \quad (5)$$

where E^+ (E^-) denotes the (non-)edge pixel set. And the smoothing term can be written as:

$$L_{re}(S, S_{gt}) = L_1(S, S_{gt}) + \sum_i \text{SSIM}(P_{2^i}(S), P_{2^i}(S_{gt})) \quad (6)$$

where P_n denotes the max pooling with $n \times n$ aiming to measure the SSIM in a multi-scale manner. The edge-guided regularization term is expressed as:

$$L_{DTV}(S, S_{gt}) = L_1(S, S_{gt}, E^+) + TV(S, S_{gt}, E^-) \quad (7)$$

The TV loss encourages smoothness in the non-edge domain, and the L1 term contributes to the structure consistency.

3.3. Contrastive Constraint

Prior networks supervised by limited smoothing ground-truths often result in the domain shift and low generalization towards real-world domain. They neglect the over/under-smoothed images, such extra supervision would help the network to perceive negative representations and meet a performance gain. Inspired by the success of contrastive learning in low-level vision tasks, which aims to learn a representation that pulls the predicted image together with positive samples and pushes it far away from the negative ones, we tend to augment the supervision by fully exploring the information behind the negative samples. Therefore, three problems are proposed: 1) How to define the positive/negative samples? 2) How to construct a latent feature space? 3) How do measure the distance in the feature space?

In our work, we present the contrastive module to facilitate the training process, see from Fig. 3. For structure preservation, we argue that the over-smoothed images usually lose the semantic structure and suffer from artifacts and halos. For the image smooth-

ing, an input without any operation is full of complicated textures, which is opposite to our target. Therefore, as illustrated in Fig. 3, we regard the ground-truth S_{gt} as a positive sample, the original input I and over-smoothed image I_{over} as negative ones, and the prediction S_1 as an anchor. That is, $I_P = \{S\}$, $I_N = \{I, I_{over}\}$. We then adopt the pretrained VGG16 to construct the latent space. Due to the high sparsity of the deep feature space, we measure the distance based on the channel correlation. Technically, for the feature map of a convolutional layer $f \in \mathbb{R}^{C \times H \times W}$, a channel-wise Gram matrix $G \in \mathbb{R}^{C \times C}$, $G_{ij}^f = \sum_k f_{ik}^f f_{jk}^f$ is obtained firstly. In a sense, the Gram matrix represents the relationship between channels: the diagonal elements focus on self-correlation while others tend to reveal cross-correlations, which is a quantitative description of the latent feature. Therefore, our goal can be formulated as follows:

$$\begin{aligned} d(G(\tilde{S}_1), G(\tilde{I}_P)) &\ll d(G(\tilde{S}_1), G(\tilde{I}_N)) \\ d(E(\tilde{S}_1), E(\tilde{I}_P)) &\ll d(E(\tilde{S}_1), E(\tilde{I}_N)) \end{aligned} \quad (8)$$

where S_1, I_P, I_N denote the anchor, positive sample, and negative sample respectively. $\tilde{(\cdot)}$ denotes the feature map in VGG16, $E(\cdot)$ denotes the expectation. Inspired by the triplet loss, our loss function can be formulated as follows:

$$\begin{aligned} L_c = \max \{ &d(G(\tilde{S}_1), G(\tilde{I}_P)) - d(G(\tilde{S}_1), G(\tilde{I}_N)) + \alpha, 0 \} \\ &+ \max \{ d(E(\tilde{S}_1), E(\tilde{I}_P)) - d(E(\tilde{S}_1), E(\tilde{I}_N)) + \beta, 0 \} \end{aligned} \quad (9)$$

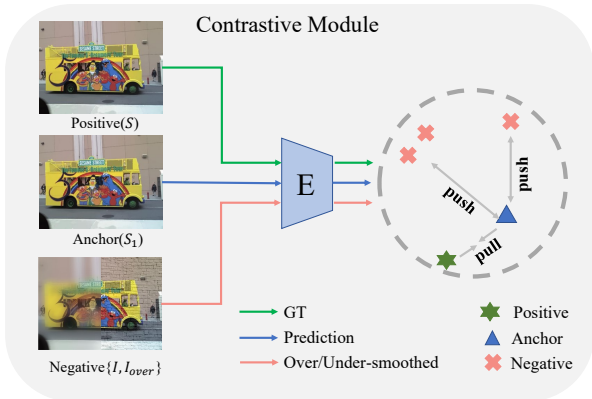


Figure 3: The architecture of the contrastive module.

3.4. Semantic Structure Preservation

Although the introduction of the contrastive module has improved the performance of the network, smoothing images from a semantic perspective is still a challenge since strong details (non-semantic) and weak yet semantic structures are often intertwined in real data.

To further improve the model's capacity for structure retention and encourage semantic distinctiveness, we introduce a semantic segmentation module to explore the high-level semantic information to benefit our CSGIS-Net. In fact, experiments have demonstrated that well-smoothed images can boost the performance of semantic segmentation tasks. That is, if the image is filled with non-semantic textures, the performance of the segmentation will be

significantly dropped. It means that applying semantic supervision would meet a performance gain. Additionally, for non-synthetic images, since textures vary greatly across different object semantics, segmentation maps would guide the smoothing network to be semantic adaptive. We consider employing the mainstream network Deeplabv3+ [CZP*18] to perform the segmentation task and adopt the cross entropy loss to ensure the consistency at the semantic level. The loss function can be defined as:

$$L_{seg} = CrossEntropy(\phi(S_1), Seg_{gt}) \quad (10)$$

where $\phi(\cdot)$ denotes the mapping of Deeplabv3+, Seg_{gt} denotes the segmentation labels of PASCAL VOC 2012.

Generally, the total loss is represented as follows:

$$\begin{aligned} L_{total} = \lambda_e L_{edge}(E_{pred}, E_{gt}) &+ L_{re}(S_0, S_{gt}) + L_{re}(S_1, S_{gt}) \\ &+ L_{DTV}(S_0, S_{gt}) + L_{DTV}(S_1, S_{gt}) + \lambda_c L_c + \lambda_{seg} L_{seg} \end{aligned} \quad (11)$$

We set $\lambda_e = 0.001, \lambda_c = \lambda_{seg} = 1$.

4. Experiments

In this section, we first establish a dataset with smoothing, edge, and segmentation ground-truths for subsequent experiments. Then, we evaluate the proposed method against several smoothing methods on two synthetic datasets quantitatively and qualitatively. Finally, several ablation studies are conducted to verify the key contribution of our method and reveal how the established ground-truth domain affects the smoothing results.

4.1. Dataset

Inspired by [XLH*20], structure-preserve smoothing can be regarded as a structure restoration task. However, it is difficult to determine the ground-truths of nature images. To our knowledge, most existing datasets on detection and segmentation (e.g., Cityscapes [COR*16], MS COCO [LMB*14], PASCAL VOC 2012 [EEVG*15]) lack the smoothing ground-truth, which is necessary for our supervised method. To realize the designed network, a synthetic dataset is established that first bridges image smoothing and semantic segmentation tasks together. It contains smoothing ground-truths and segmentation labels, where the former is also used to generate the texture-enhanced images (i.e input of the network). For pixel-wise segmentation labels, since the manual annotation is extremely laborious, we build the dataset based on the PASCAL VOC 2012 benchmark [EEVG*15]. After that, our main task can be converted into two subtasks: 1) How to generate the smoothing ground-truth images? 2) How to design the blending strategy to generate the input?

Ground-truths Generation: For the ground-truths of semantic smoothing, the following properties should be satisfied: 1) no severe distortion; 2) be salient in semantic structure domain; 3) smooth out non-semantic textures.

Previous methods usually fail to perform well since most of them do not consider the distinctive needed in structure and non-structure areas. Intuitively, these methods will smooth "easy" images (salient structure and small details) with better efficiency. Hence, we first collect a set of images from PASCAL VOC 2012,

which is relatively easy to handle. More than 800 images are selected roughly, and we feed them into different smoothing operators such as BF [TM98], L0 [XLXJ11], RGF [ZSXJ14], RTV [XYXJ12], SDF [HCP17] to generate a series of candidates. Due to the duality of our task, that is structure preservation and detail smoothing, the evaluations of smoothness and structure similarity are performed in different regions.

Smoothness Evaluation: We argue that a visual-pleasing smoothed patch should have the following properties: 1) no extreme or sharp color oscillations (such as edges); 2) the transition of the pixel value is relatively smooth (e.g. no strong noise). Technically, we first set a window with a specific size, and slide it to search the region without edges. Note that, the edge label is obtained through a Sobel operator. After that, previous work has proved that the 1st derivative (gradient) of an image would reflect the variance rate of the gray level, and the 2nd derivative is more sensitive to the noise. To perform patch-level measurements, we combine the 2nd derivative and the standard variance, which can be formulated as follows:

$$L(x, y) = \left(\partial^2 I / \partial x^2 \right) + \left(\partial^2 I / \partial y^2 \right) \quad (12)$$

$$\text{smooth_value}_i = \text{std}(\text{patch}_i) + L(x, y)_i$$

Where $\text{patch}_i \in [0, 1]$ denotes the gray value of the i -th patch. Experiments show that it is nearly close to zero. Hence, we set 0.05 as a threshold value, the candidate would pass the smooth test if the smooth value is less than the threshold.

Structure Consistency Evaluation: The non-smoothed patches are obtained via a voting strategy (i.e. more than 90% of the volunteers agree that the patch contains semantic structure). Once a volunteer reports a non-ignorable distortion, the corresponding image would be directly abandoned. It is notable that volunteers only make binary judgments, which is very simple and fast. Their judgment is only used for excluding ineligible candidates rather than being the evidence for final decisions as in [ZLJ*19]. To measure the structural integrity, we adopt the SSIM [WBSS04] formulated as:

$$\text{SSIM}(x, y) = [l(x, y)]^\alpha \times [c(x, y)]^\beta \times [s(x, y)]^\gamma \quad (13)$$

where $l(x, y) = \frac{2\mu_x\mu_y + c_1}{\mu_x^2 + \mu_y^2 + c_1}$, $c(x, y) = \frac{2\sigma_x\sigma_y + c_2}{\sigma_x^2 + \sigma_y^2 + c_2}$, $s(x, y) = \frac{\sigma_{xy} + c_3}{\sigma_x\sigma_y + c_3}$ represent brightness, contrast and structure. α , β , γ represent the importance of three items respectively, μ_i , σ_i , σ_{xy} denote the mean value, standard deviation, and covariance of the gray images respectively. The candidate passing the smooth test with the largest SSIM value will be the smooth ground-truth.

Generally, we select N easy samples and 5 smoothing operators, the smoothing result of the i^{th} sample is represented as $O_i = \{f_1(I_i), f_2(I_i), \dots, f_p(I_i)\}$, where $f_i(\cdot)$ denotes the i^{th} smoothing operation. The final ground-truth images obtained by the screening test can be represented as $F = \{S_1, S_2, \dots, S_n\}$, where $S_i \in O_j, n \leq N$. Note that, no more than one ground-truth for a single sample.

Texture Blending: To remove texture details and preserve semantic structures, based on the well-smoothed ground-truth, we naturally introduce the texture blending process and apply the algorithm presented by [EF01]. Such a strategy easily distinguishes the texture and structure of an input image. Technically speaking, we

collect 100 texture images from [FDY*21], which cover different materials (such as grass, brick, fiber, etc.) and shapes (point, line, etc.). The algorithm is summarized as follows:

- 1) Convert the texture image to grayscale:

$$I_{\text{gray}} = 0.299 \times I_R + 0.587 \times I_G + 0.114 \times I_B \quad (14)$$

where I_R, I_G, I_B denote the pixel value in R, G, B channels respectively.

- 2) Calculate the mean value:

$$\mu = \frac{\sum_{i=0}^H \sum_{j=0}^W I_{\text{gray}}(i, j)}{H \times W} \quad (15)$$

where H, W denote the height and weight of the input image.

- 3) Texture blending:

$$I = I_{GT} + (I_{\text{gray}} - \mu) \quad (16)$$

where $(I_{\text{gray}} - \mu)$ represents the texture layer, and I_{GT} is the ground-truth image. The blending result I would be fed into the network as an input. Fig. 4 shows some examples.




Figure 4: Some examples of our established dataset. The left is ground-truth images, the right is the inputs of our network.

4.2. Implementation Details

The established VOC-smooth dataset consists of 15k images, which blends 100 texture layers and 150 smoothing ground-truths. Among them, 12k images are chosen randomly for training, 2600 images for validation, and 400 images for testing. Images for testing are not overlapped with the training images. In order to better generalize our method to various cases, we adopt a series of data augmentations: 1) random horizontal flip; 2) random image rotation; 3) color disturbances. The proposed CSGIS-Net is trained in a fully supervised fashion.

Our networks are implemented on the Pytorch platform and trained on an NVIDIA GeForce GTX 1080Ti GPU. During the training phase, the patch size is set to 128×128 and the batch size is set to 8. For accelerating the training process, the ADAM

Table 1: Average PSNR/SSIM comparison on the synthetic VOC-smooth and NKS datasets. The highest rank in comparison is highlighted in bold.

	VOC-smooth		NKS	
	SSIM	PSNR	SSIM	PSNR
Original	0.3882	20.9701	0.5736	28.0418
GF [HST12]	0.5452	23.7863	0.8809	32.0883
RGF [ZSXJ14]	0.7076	22.3704	0.9033	31.0875
L0 [XLXJ11]	0.4684	21.9665	0.7854	28.4811
RTV [XYXJ12]	0.7775	24.8952	0.9221	31.8124
SDF [HCP17]	0.5312	22.8657	0.9077	33.1725
ILS [LZH*20]	0.5474	22.3584	0.9050	29.6929
PNLS [XLH*20]	0.6622	23.6922	0.9341	33.2073
JESS [FDY*21]	0.8788	28.3701	0.9435	33.0475
Ours	0.8993	29.3366	0.9486	34.5098

optimizer [KB14] is adopted with $\beta_1 = 0.9, \beta_2 = 0.999$. The initial learning rate $1e-3$ decayed by multiplying 0.5 after every 30 epochs and we train 250 epochs in total. We set the number of rHDC blocks $N = 3$ and $N = 6$ in the smooth interpreter and the enhancement module, respectively.

4.3. Comparison Results

In the testing phase, we adopt two synthetic datasets (VOC-smooth and NKS [XLH*20]) to quantitatively evaluate different methods. All images in the two datasets are blended by texture layers and diverse structure images (i.e. ground-truths). Note that the NKS dataset is blended by cartoon images, which is quite simple compared to real-world images. We apply the Peak Signal to Noise Ratio (PSNR) and the Structural Similarity (SSIM) for evaluation, and higher values indicate better performance.

Our proposed method is compared with eight representative smoothing algorithms, including two filters (i.e. GF [HST12], RGF [ZSXJ14]), five optimization-based methods (i.e. L0 [XLXJ11], RTV [XYXJ12], SDF [HCP17]), ILS [LZH*20], PNLS [XLH*20] and a data-driven network JESS [FDY*21]. We determine the model with the best performance in the validating phase for further analysis. The results are summarized in Tab. 1, by averaging the PSNR and SSIM of the test set. Note that, for traditional methods, we choose the best result from a set of manual parameters. It can be observed that our CSGIS-Net outperforms the others on the two synthetic datasets. Specifically, compared to the second best methods, the PSNRs of VOC-smooth and NKS are 0.9665 and 1.3373 higher, respectively.

We also exhibit visual comparisons where the parameters of traditional algorithms are manually determined via the visual quality. As shown in Fig. 5, 6, 7, GF and SDF can hardly remove the texture layer. L0 and RTV remove texture with appropriate parameters, but some salient structures are also overblurred, such as eyes, windows. Those methods fail to achieve a trade-off between texture removal and structure preservation since the large-scale smoothing strength will introduce distortion and overblurring, while the small-scale one is disabled to smooth the texture. However, our proposed

network can achieve notable performance compared with other algorithms.

In fact, we expect that our network would be generalized to the real domain. In the testing phase, the chosen real-world image contains strong textures or complex details, such as grass, road, etc. Since the absence of smoothing ground-truth, we only exhibit the visual comparisons of these two categories, i.e. 1) natural images; 2) real textural images.

As depicted in Fig. 8,9, our approach generates considerably better results while others suffer blurring or distortion to varying degrees. For the natural images, the smooth strength of GF is small, leading to the retention of small textures such as leaves; L0 loses salient edges and introduces artifacts; SDF fails to smooth the clouds. Our method produces a more natural mapping in the texture layer. For the real textural image, GF, L0, and SDF have limited ability for texture removal. RGF and RTV gain better smoothing performance but neglect semantic structures such as red dots and white lines (see from Fig. 9). Our proposed network possesses a better generalization ability to real data, due to the exploration of the contrastive and semantic priors.

4.4. Ablation Study

In this subsection, extensive ablation studies are conducted to verify the effectiveness of the proposed components in our framework, including 1) two introduced modules of contrastive learning and image segmentation; 2) the established dataset (i.e. the training domain). Considering the training process on the entire dataset would take about 3-4 days with 250 epochs, all the ablation studies are organized on the subset of our datasets (VOC-smooth). Specifically, we randomly choose 2600 images for training, 400 images for validating and 400 images for testing.

Ablation in Network Architecture

We discuss the contributions of the introduced modules and evaluate the performance on the two datasets, i.e. VOC-smooth and NKS, the quantitative comparison is illustrated in Tab. 2. One can notice that the network with all proposed modules achieves the best performance. Specifically, compared to the backbone, the PSNR and SSIM are improved by 1.8088 and 0.0266 respectively. Moreover, simply adding a contrastive module also meets a performance gain. A typical example is shown in Fig. 10. It can be observed that (b) and (c) overblur some edges on the cloud and butterfly's wings (see from the first row). Furthermore, the second row demonstrates that our network generates visual-pleasing results, especially for semantic structure preservation. It is inferred that contrastive learning with negatives may prevent the prediction from the over/under-smoothed domain.

Note that using only the contrastive loss to train the proposed network may not yield satisfactory results. This is due to the sparsity of the contrastive module with the high-dimensional feature space, where small changes may have a large impact on the low-dimensional smoothing results. Furthermore, our recent experiments demonstrate that the model fails to converge when only constrained by the contrastive loss, implying that a low-dimensional loss is necessary.




Figure 5: Visual Results on NKS. (b) GF ($\sigma = 2.5$); (c) RGF ($\sigma_c = 0.3, \sigma_s = 6, iter = 4$); (d) L0 ($\lambda = 0.01$); (e) RTV ($\epsilon_s = 0.02, iter = 3$); (f) SDF ($\lambda = 3$); (g) PNLs ($thr = 0.4, iter = 9$); (h) ILS ($\lambda = 30, \gamma = 10/255, iter = 10$).




Figure 6: Visual Results on VOC-smooth. (b) L0 ($\lambda = 0.1$); (c) RTV ($\epsilon_s = 0.01, iter = 2$); (d) SDF ($\lambda = 3$); (e) PNLs ($thr = 0.4, iter = 9$); (f) ILS ($\lambda = 30, \gamma = 13/255, iter = 10$).




Figure 7: Visual Results on VOC-smooth. (b) GF ($\sigma = 2.5$); (c) RGF ($\sigma_c = 0.3, \sigma_s = 6, iter = 4$); (d) L0 ($\lambda = 0.08$); (e) RTV ($\epsilon_s = 0.02, iter = 3$); (f) SDF ($\lambda = 3$).




Figure 8: Visual results on a natural image. (b) GF ($\sigma = 2.5$); (c) L0 ($\lambda = 0.01$); (d) SDF ($\lambda = 3$).

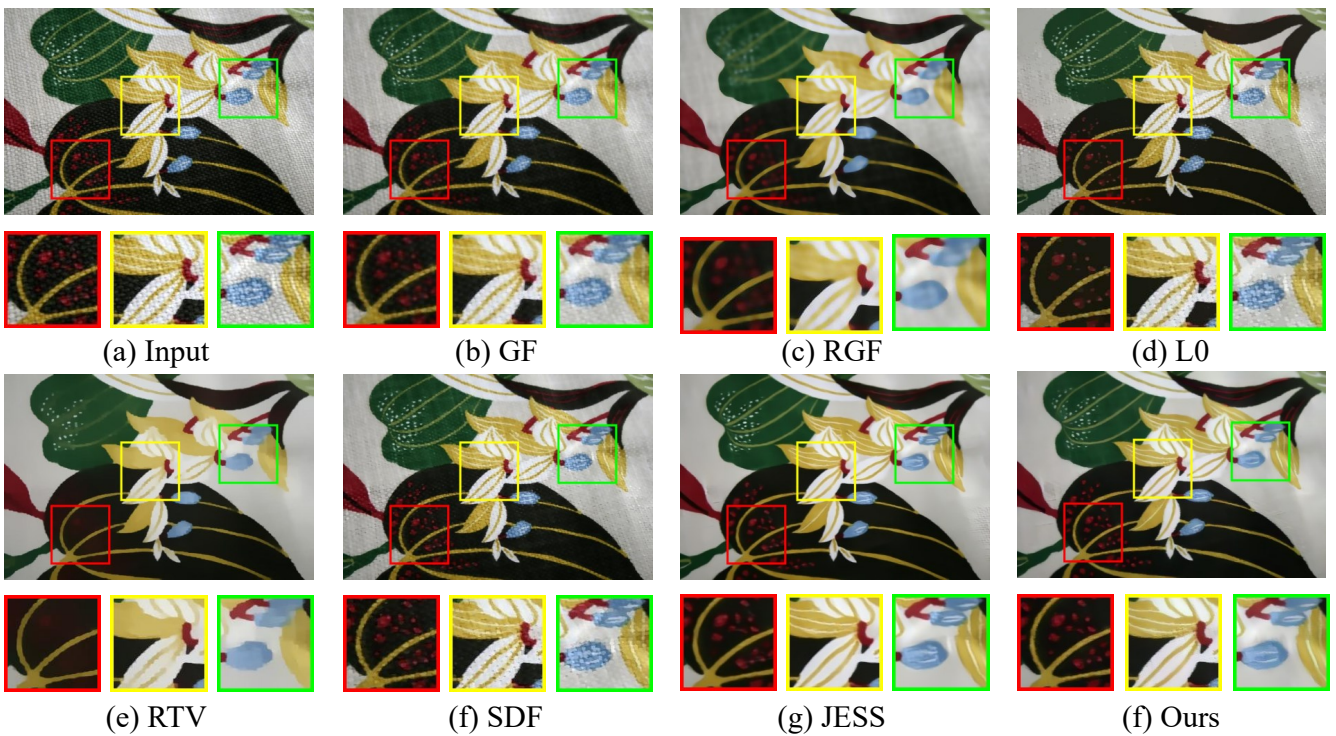


Figure 9: Visual results on a real texture image. (b) GF ($\sigma = 2.5$); (c) RGF ($\sigma_c = 0.3, \sigma_s = 6, \text{iter} = 4$); (d) L0 ($\lambda = 0.01$); (e) RTV ($\epsilon_s = 0.02, \text{iter} = 3$); (f) SDF ($\lambda = 3$).

Table 2: Quantitative comparison for the network ablation study.

	Contrastive Module	Semantic Module	VOC-smooth		NKS	
			SSIM	PSNR	SSIM	PSNR
CSGIS-baseline			0.8381	25.6451	0.9276	30.0736
CSGIS-CM	✓		0.8478	26.4537	0.9310	30.3613
CSGIS	✓	✓	0.8647	27.4539	0.9340	32.2314

Ablation in Dataset

We explore the potential impact of the established dataset (VOC-

smooth). Firstly, over/under-smoothed domains are constructed separately for training. The former takes over-smoothed images as




Figure 10: Visual results for ablation study. (a) Input images; (b) Baseline, no Contrastive module (CM) and Semantic segmentation module (SM); (c) CSGIS-CM, only adding CM; (d) CSGIS-Net, the network of this paper; (e) GT, the ground-truth. The first row is from the NKS dataset and the second row is from the VOC-smooth dataset.




Figure 11: The training datasets of three different domains. The left is ground-truths and the right is input images. Over/under-smoothed means the ground-truth images are over/under-smoothed.

Table 3: Quantitative comparison for the dataset ablation study.

Training Dataset	NKS	
	SSIM	PSNR
Under-smoothed	0.9191	29.2022
Over-smoothed	0.9076	29.8602
VOC-smooth	0.9340	32.2314

smoothing ground-truths, which are obtained by a Bilateral Filter with a large smooth strength ($\sigma_c = 200, \sigma_s = 1000$). The latter uses the original images from PASCAL VOC 2012 as ground-truths, it is interwoven with the real texture. We then apply the aforementioned blending strategy to synthesize the input images. Retrain our model in three different domains, for fairness, we only test on the NKS dataset quantitatively.

As shown in Tab. 3, the network trained on VOC-smooth gains the highest PSNR and SSIM. The visual results also illustrate that the VOC-smooth domain performs favorably against others (see Fig. 12). We argue that for the over-smoothed one, the ground-truth itself loses semantic structures and clear edges. Therefore, the network fails to learn the representation of structures, leading

to the problem of blurring. Moreover, the network trained on the under-smoothed domain has a limited ability in texture removal, since the mixture of synthetic and real texture confuses the network as to what to remove/preserve. However, benefiting from the three properties in sec. 4.1, our dataset contributes to smoothing the non-semantic details under strong texture.

5. Conclusions

In this work, we propose a novel network (CSGIS-Net) that aims to preserve the salient structure of the image while smoothing the insignificant details by introducing two important prior knowledge. CSGIS-Net augments the supervision from unexplored over/under-smoothed samples and semantic information, which promote the network to be more robust and semantic adaptive. Based on the introduced contrastive module and segmentation module, the network is regarded as a multi-task framework with three parts of loss function for semantic structure preservation and texture removal. Additionally, we establish a new dataset (VOC-smooth) that mixes versatile structure images with natural textures. It also makes the first trial to bridge image smoothing and semantic segmentation. Extensive experiments demonstrate that the proposed CSGIS-Net outperforms the state-of-the-art smoothing methods both visually

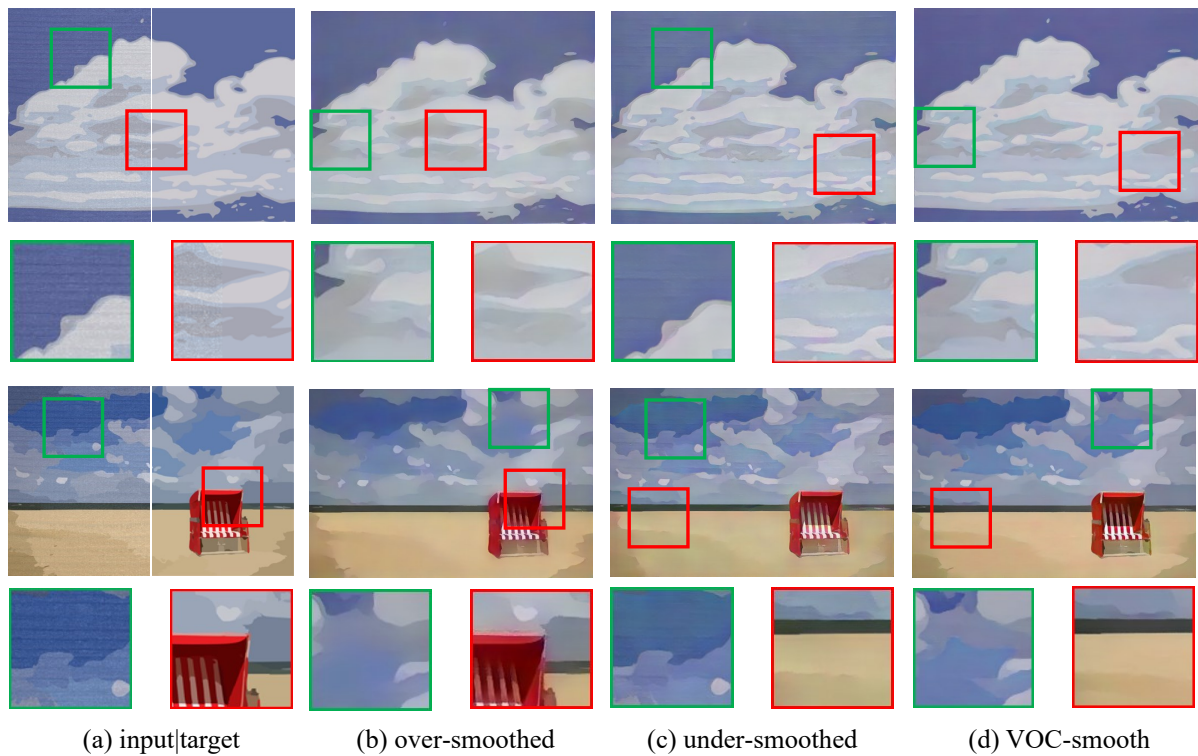


Figure 12: Visual results for dataset ablation study.

and quantitatively. In the future, we will extend our method to other applications such as detail enhancement, and further explore other strategies for exploiting semantic information.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (No. 62172218), the Free Exploration of Basic Research Project, Local Science and Technology Development Fund Guided by the Central Government of China (No. 2021Szvup060), the Natural Science Foundation of Guangdong Province (No. 2022A1515010170), the Direct Grant (No. DR22A2) and the Research Grant entitled "Self-Supervised Learning for Medical Images" (No. 871228) of Lingnan University, Hong Kong.

References

- [AT17] AGUSTSSON E., TIMOFTE R.: Ntire 2017 challenge on single image super-resolution: Dataset and study. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops* (2017), pp. 126–135.
- [BPCD11] BYCHKOVSKY V., PARIS S., CHAN E., DURAND F.: Learning photographic global tonal adjustment with a database of input/output image pairs. In *CVPR 2011* (2011), IEEE, pp. 97–104.
- [BZCC10] BHAT P., ZITNICK C. L., COHEN M., CURLESS B.: Gradientshop: A gradient-domain optimization framework for image and video filtering. *ACM Transactions on Graphics (TOG)* 29, 2 (2010), 1–14.
- [COR*16] CORDTS M., OMRAN M., RAMOS S., REHFELD T., ENZWEILER M., BENENSON R., FRANKE U., ROTH S., SCHIELE B.: The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (2016), pp. 3213–3223.
- [CPK*18] CHEN L.-C., PAPANDREOU G., KOKKINOS I., MURPHY K., YUILLE A. L.: Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 40, 4 (2018), 834–848.
- [CSBL16] CONNEAU A., SCHWENK H., BARRAULT L., LECUN Y.: Very deep convolutional networks for text classification. *arXiv preprint arXiv:1606.01781* (2016).
- [CZP*18] CHEN L.-C., ZHU Y., PAPANDREOU G., SCHROFF F., ADAM H.: Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European conference on computer vision (ECCV)* (2018), pp. 801–818.
- [DDS*09] DENG J., DONG W., SOCHER R., LI L.-J., LI K., FEI-FEI L.: Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition* (2009), Ieee, pp. 248–255.
- [DWW*20] DENG S., WEI M., WANG J., FENG Y., LIANG L., XIE H., WANG F. L., WANG M.: Detail-recovery image deraining via context aggregation networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (2020), pp. 14560–14569.
- [EEVG*15] EVERINGHAM M., ESLAMI S., VAN GOOL L., WILLIAMS C. K., WINN J., ZISSERMAN A.: The pascal visual object classes challenge: A retrospective. *International journal of computer vision* 111, 1 (2015), 98–136.
- [EF01] EFROS A. A., FREEMAN W. T.: Image quilting for texture syn-

- thesis and transfer. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques* (2001), pp. 341–346.
- [FAR07] FATTAL R., AGRAWALA M., RUSINKIEWICZ S.: Multiscale shape and detail enhancement from multi-light image collections. *ACM Trans. Graph.* 26, 3 (2007), 51.
- [FDY*21] FENG Y., DENG S., YAN X., YANG X., WEI M., LIU L.: Easy2hard: Learning to solve the intractables from a synthetic dataset for structure-preserving image smoothing. *IEEE Transactions on Neural Networks and Learning Systems* (2021).
- [FFLS08] FARBMAN Z., FATTAL R., LISCHINSKI D., SZELISKI R.: Edge-preserving decompositions for multi-scale tone and detail manipulation. *ACM Transactions on Graphics (TOG)* 27, 3 (2008), 1–10.
- [FYH*17] FAN Q., YANG J., HUA G., CHEN B., WIPF D.: A generic deep architecture for single image reflection removal and image smoothing. In *Proceedings of the IEEE International Conference on Computer Vision* (2017), pp. 3238–3247.
- [FYW*18] FAN Q., YANG J., WIPF D., CHEN B., TONG X.: Image smoothing via unsupervised learning. *ACM Transactions on Graphics (TOG)* 37, 6 (2018), 1–14.
- [GH10] GUTMANN M., HYVÄRINEN A.: Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics* (2010), JMLR Workshop and Conference Proceedings, pp. 297–304.
- [GO11] GASTAL E. S., OLIVEIRA M. M.: Domain transform for edge-aware image and video processing. In *ACM SIGGRAPH 2011 papers*. 2011, pp. 1–12.
- [GO12] GASTAL E. S., OLIVEIRA M. M.: Adaptive manifolds for real-time high-dimensional filtering. *ACM Transactions on Graphics (TOG)* 31, 4 (2012), 1–13.
- [HBL17] HERMANS A., BEYER L., LEIBE B.: In defense of the triplet loss for person re-identification. *arXiv preprint arXiv:1703.07737* (2017).
- [HCP17] HAM B., CHO M., PONCE J.: Robust guided image filtering using nonconvex potentials. *IEEE transactions on pattern analysis and machine intelligence* 40, 1 (2017), 192–207.
- [Hen20] HENAFF O.: Data-efficient image recognition with contrastive predictive coding. In *International Conference on Machine Learning* (2020), PMLR, pp. 4182–4192.
- [HST12] HE K., SUN J., TANG X.: Guided image filtering. *IEEE transactions on pattern analysis and machine intelligence* 35, 6 (2012), 1397–1409.
- [HW77] HOLLAND P. W., WELSCH R. E.: Robust regression using iteratively reweighted least-squares. *Communications in Statistics-theory and Methods* 6, 9 (1977), 813–827.
- [HZRS16] HE K., ZHANG X., REN S., SUN J.: Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (2016), pp. 770–778.
- [KB14] KINGMA D. P., BA J.: Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [LJKL17] LEE H., JEON J., KIM J., LEE S.: Structure-texture decomposition of images with interval gradient. In *Computer graphics forum* (2017), vol. 36, Wiley Online Library, pp. 262–274.
- [LMB*14] LIN T.-Y., MAIRE M., BELONGIE S., HAYS J., PERONA P., RAMANAN D., DOLLÁR P., ZITNICK C. L.: Microsoft coco: Common objects in context. In *European conference on computer vision* (2014), Springer, pp. 740–755.
- [LYB18] LU K., YOU S., BARNES N.: Deep texture and structure aware filtering network for image smoothing. In *Proceedings of the European conference on computer vision (ECCV)* (2018), pp. 217–233.
- [LZH*20] LIU W., ZHANG P., HUANG X., YANG J., SHEN C., REID I.: Real-time image smoothing via iterative least squares. *ACM Transactions on Graphics (TOG)* 39, 3 (2020), 1–24.
- [LZL*21] LIU W., ZHANG P., LEI Y., HUANG X., YANG J., NG M. K.-P.: A generalized framework for edge-preserving and structure-preserving image smoothing. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2021).
- [MFTM01] MARTIN D., FOWLKES C., TAL D., MALIK J.: A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proceedings Eighth IEEE International Conference on Computer Vision. ICCV 2001* (2001), vol. 2, IEEE, pp. 416–423.
- [PEZZ20] PARK T., EFROS A. A., ZHANG R., ZHU J.-Y.: Contrastive learning for unpaired image-to-image translation. In *European Conference on Computer Vision* (2020), Springer, pp. 319–345.
- [PM90] PERONA P., MALIK J.: Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on pattern analysis and machine intelligence* 12, 7 (1990), 629–639.
- [ROF92] RUDIN L. I., OSHER S., FATEMI E.: Nonlinear total variation based noise removal algorithms. *Physica D: nonlinear phenomena* 60, 1-4 (1992), 259–268.
- [SCTJ17] SHEN X., CHEN Y.-C., TAO X., JIA J.: Convolutional neural pyramid for image processing. *arXiv preprint arXiv:1704.02071* (2017).
- [SSW17] SUN Y., SCHAEFER S., WANG W.: Image structure retrieval via l_0 minimization. *IEEE transactions on visualization and computer graphics* 24, 7 (2017), 2129–2139.
- [SZ14] SIMONYAN K., ZISSERMAN A.: Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).
- [TKI20] TIAN Y., KRISHNAN D., ISOLA P.: Contrastive multiview coding. In *European conference on computer vision* (2020), Springer, pp. 776–794.
- [TM98] TOMASI C., MANDUCHI R.: Bilateral filtering for gray and color images. In *Sixth international conference on computer vision (IEEE Cat. No. 98CH36271)* (1998), IEEE, pp. 839–846.
- [VdOLV18] VAN DEN OORD A., LI Y., VINYALS O.: Representation learning with contrastive predictive coding. *arXiv e-prints* (2018), arXiv:1807.
- [WBSS04] WANG Z., BOVIK A. C., SHEIKH H. R., SIMONCELLI E. P.: Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* 13, 4 (2004), 600–612.
- [WWH06] WANG G., WONG T.-T., HENG P.-A.: Deringing cartoons by image analogies. *ACM Transactions on Graphics (TOG)* 25, 4 (2006), 1360–1379.
- [XLH*20] XU J., LIU Z.-A., HOU Y.-K., ZHEN X.-T., SHAO L., CHENG M.-M.: Pixel-level non-local image smoothing with objective evaluation. *IEEE Transactions on Multimedia* 23 (2020), 4065–4078.
- [XLXJ11] XU L., LU C., XU Y., JIA J.: Image smoothing via l_0 gradient minimization. In *Proceedings of the 2011 SIGGRAPH Asia conference* (2011), pp. 1–12.
- [XRY*15] XU L., REN J., YAN Q., LIAO R., JIA J.: Deep edge-aware filters. In *International Conference on Machine Learning* (2015), PMLR, pp. 1669–1678.
- [XYXJ12] XU L., YAN Q., XIA Y., JIA J.: Structure extraction from texture via relative total variation. *ACM transactions on graphics (TOG)* 31, 6 (2012), 1–10.
- [ZLJ*19] ZHU F., LIANG Z., JIA X., ZHANG L., YU Y.: A benchmark for edge-preserving image smoothing. *IEEE Transactions on Image Processing* 28, 7 (2019), 3556–3570.
- [ZSXJ14] ZHANG Q., SHEN X., XU L., JIA J.: Rolling guidance filter. In *European conference on computer vision* (2014), Springer, pp. 815–830.