

Structural and Statistical Texture Knowledge Distillation for Semantic Segmentation

Deyi Ji^{1,2,*}, Haoran Wang^{1,*}, Mingyuan Tao¹, Jianqiang Huang¹, Xian-Sheng Hua^{1,†}, Hongtao Lu^{2,†}

¹Alibaba Cloud Computing Ltd.

²Department of Computer Science and Engineering, Shanghai Jiao Tong University

{jideyi.jdy, jinglai.whr, juchen.tmy, jianqiang.hjq, xiansheng.hxs}@alibaba-inc.com

htlu@sjtu.edu.cn

Abstract

Existing knowledge distillation works for semantic segmentation mainly focus on transferring high-level contextual knowledge from teacher to student. However, low-level texture knowledge is also of vital importance for characterizing the local structural pattern and global statistical property, such as boundary, smoothness, regularity and color contrast, which may not be well addressed by high-level deep features. In this paper, we are intended to take full advantage of both structural and statistical texture knowledge and propose a novel Structural and Statistical Texture Knowledge Distillation (SSTKD) framework for Semantic Segmentation. Specifically, for structural texture knowledge, we introduce a Contourlet Decomposition Module (CDM) that decomposes low-level features with iterative laplacian pyramid and directional filter bank to mine the structural texture knowledge. For statistical knowledge, we propose a Denoised Texture Intensity Equalization Module (DTIEM) to adaptively extract and enhance statistical texture knowledge through heuristics iterative quantization and denoised operation. Finally, each knowledge learning is supervised by an individual loss function, forcing the student network to mimic the teacher better from a broader perspective. Experiments show that the proposed method achieves state-of-the-art performance on Cityscapes, Pascal VOC 2012 and ADE20K datasets.

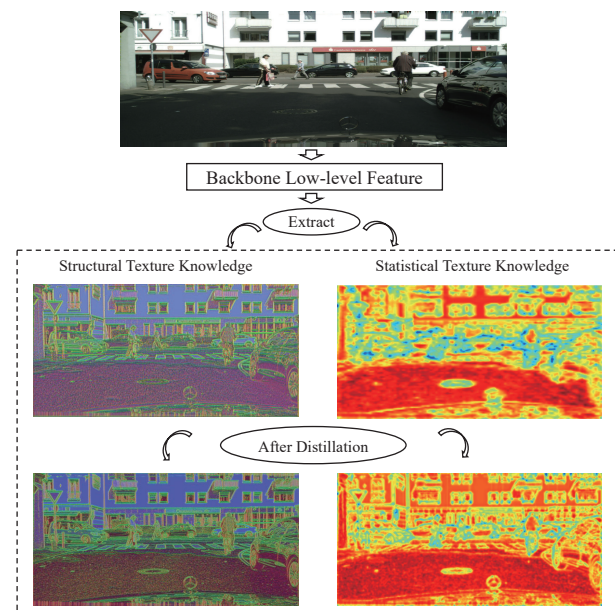


Figure 1. The overview of the structural and statistical texture knowledge distillation of an example image. Two kinds of the texture knowledge are extracted from the low-level feature of a CNN backbone. The original structural and statistical texture are fuzzy and in low-contrast. After distillation, the contour is clearer and the intensity contrast is more equalized, showing that two kinds of the texture are both enhanced.

1. Introduction

Semantic segmentation, which aims to assign each pixel a unique category label for the input image, is a crucial and challenging task in computer vision. Recently, deep fully convolution network [32] based methods have

achieved remarkable results in semantic segmentation, and extensive methods have been investigated to improve the segmentation accuracy by introducing sophisticated models [6, 8, 23, 27, 44, 46, 47, 50, 52]. However, these methods are usually based on a large model, which contains tremendous parameters. Since semantic segmentation has shown great potential in many applications like autonomous driving, video surveillance, robot sensing and so on, how to keep efficient inference speed and high accuracy with high-

*Equal Contribution

†Corresponding Authors. Hongtao Lu is also with MOE Key Lab of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University, China.

resolution images is a critical problem.

The focus of this paper is knowledge distillation, which is introduced by Hinton et al. [22] based on a teacher-student framework, and has received increasing attention in semantic segmentation community [20, 30, 31, 37, 42]. Previous works mainly focus on the high-level contextual knowledge [30, 37] or the final response knowledge [20, 31, 42], which are appropriate to capture the global context and long range relation dependencies among pixels, but will also result in coarse and inaccurate segmentation results, since they are usually extracted with a large receptive field and miss many low-level texture details. In this paper, we concentrate on exploiting the texture knowledge from the teacher to enrich the low-level information of the student. According to the digital image processing [18], texture is a region descriptor which can provide measures for both local structural property and global statistical property of an image. The structural property can also be viewed as spectral domain analysis and often refer to some local patterns, such as boundary, smoothness and coarseness. While the statistical property pay more attention to the global distribution analysis, such as histogram of intensity.

Based on the above analysis, we propose a novel Structural and Statistical Texture Knowledge Distillation (SSTKD) framework to effectively distillate two kinds of the texture knowledge from the teacher model to the student model, as shown in Figure 1. More comprehensively, we introduce a Contourlet Decomposition Module (CDM) that decomposes low-level features to mine the structural texture knowledge with iterative laplacian pyramid and directional filter bank. The contourlet decomposition is a kind of multiscale geometric analysis tool and can enable the neural network the ability of geometric transformations, thus is naturally suitable for describing the structural properties. Moreover, we propose a Denoised Texture Intensity Equalization module (DTIEM) to adaptively extract and enhance the statistical knowledge, cooperated with an Anchor-Based Adaptive Importance Sampler. The DTIEM can effectively describe the statistical texture intensity in a statistical manner in deep neural networks, as well as suppress the noise produced by the amplification effect in near-constant regions during the texture equalization. Overall, our contribution is threefold:

- To our knowledge, it is the first work to introduce both the structural and statistical texture to knowledge distillation for semantic segmentation. We propose a novel Structural and Statistical Texture Knowledge Distillation (SSTKD) framework to effectively extract and enhance the unified texture knowledge and apply them to teacher-student distillation.
- More comprehensively, we introduce the Contourlet Decomposition Module (CDM) and propose the Denoised Texture Intensity Equalization Module

(DTIEM) to describe the structural and statistical texture, respectively. Moreover, DTIEM utilizes an adaptive importance sampler and a denoised operation for efficient and accurate characterization.

- Experimental results show that the proposed framework achieves the state-of-the-art performance on three popular benchmark datasets in spite of the choice of student backbones.

2. Related Work

Semantic Segmentation. Existing works for semantic segmentation mainly focused on contextual information and elaborate network [16, 21, 26, 33]. Many works [6–8, 23, 28, 50] try to take advantage of rich contextual information in deep features, while the improvement of network design has also significantly driven the performance, such as graph modules [23, 24, 39] and attention techniques [7, 17, 48]. How to find a better trade-off between accuracy and efficiency have been discussed for a long time. Real-time semantic segmentation algorithms aim to produce high-quality prediction under limited calculation [2, 34, 41, 45, 49]. BiSeNet [45] introduced spatial path and semantic path to reduce computational cost. However, though their efficiency, there is a large gap in the performance to be considered.

Knowledge Distillation. In a typical knowledge distillation framework, the logits are used as the knowledge from the teacher model [4, 22]. Many works followed this paradigm in other visual applications, such as object detection and human pose estimation. In semantic segmentation, existing methods also utilized this idea to produce fundamental results. They transferred the output maps to the student model, distilling the class probabilities for each pixel separately. Based on this, they further proposed to extract the different knowledge from the task-specific view. He et al. [20] proposed an affinity distill module to transfer the long-range dependencies among widely separated spatial regions from a teacher model to a student model. SKD [30] proposed structured knowledge distillation to transfer pairwise relations and holistic knowledge with the help of adversarial learning. IFVD [40] forced the student model to mimic the intra-class feature variation of the teacher model. CWD [37] minimized the Kullback–Leibler divergence between the channel-wise probability map of the teacher and student networks. Different from them, we first introduce the texture knowledge to semantic segmentation, showing an effective framework for this task.

Texture in Semantic Segmentation. In the perspective of digital image processing [18], texture is a kind of descriptor that provides measures of properties such as smoothness, coarseness, regularity and so on. Image texture is not only about the local structural patterns, but also global statistical

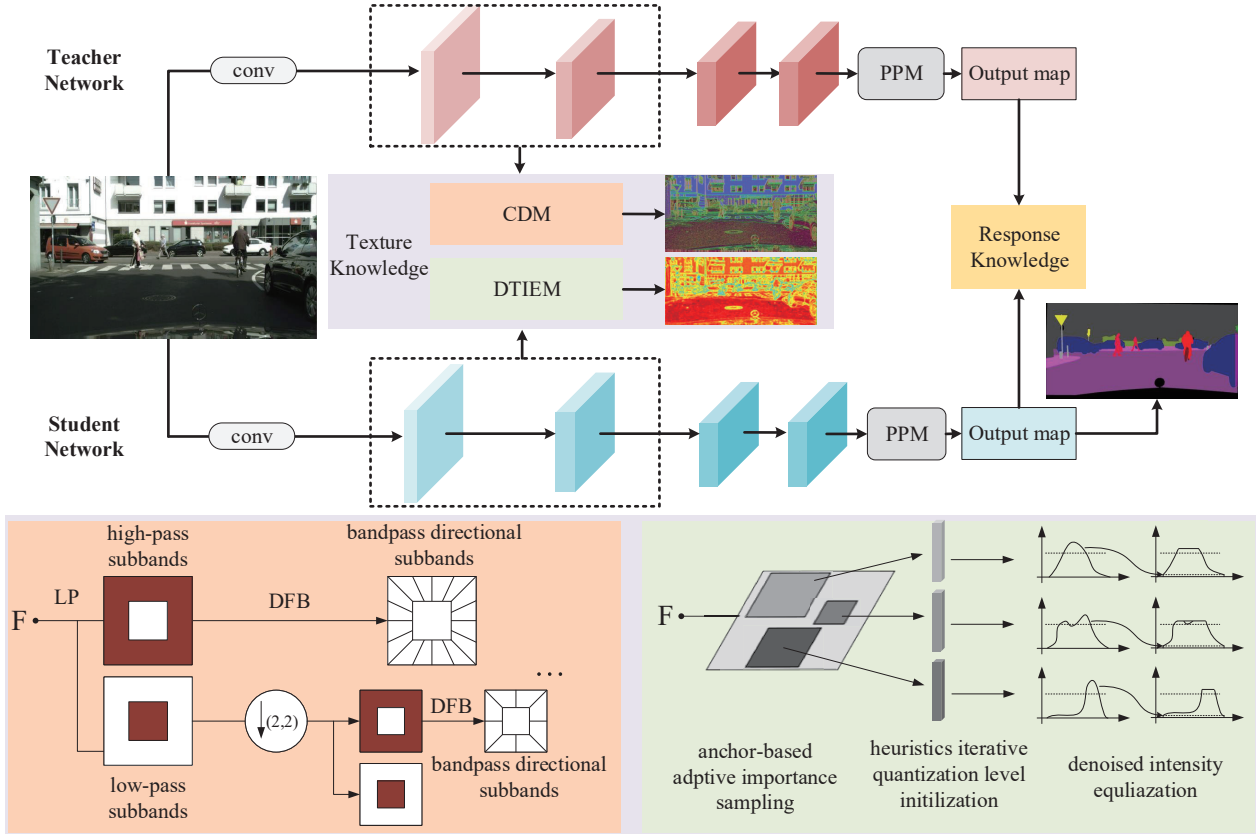


Figure 2. An overview of our proposed framework. PSPNet [50] is used as the model architecture for both teacher and student network, which consists of the backbone network, pyramid pooling module (PPM) and the final output map. Apart from the response knowledge, we further propose to extract the texture knowledge from low-level features. The corresponding parts of two kinds of the texture knowledge are depicted in the light red [1, 3, 29] and light green below the network pipeline, respectively.

property. Zhu et al. [52] firstly introduced the statistical texture to semantic segmentation, and proposed a Quantization and Count Operator (QCO) to extract the low-level statistical texture feature, which is then aggregated with high-level contextual feature to obtain a more precise segmentation map. However the QCO performs as a global manner and sparse quantization, as well as neglects the noise produced by the amplification effect. Moreover, structural texture information is not well addressed in their work. In this paper, we focus on the unified texture information including both structural and statistical, and propose to improve the statistical characterization with anchor-based importance sampling, heuristics iterative quantization and denoised operation, as well as propose the CDM to re-emphasize the structural texture.

3. Method

In this section, we introduce the proposed Structural and Statistical Texture Knowledge Distillation (SSTKD) framework in detail. Firstly, we introduce the overall Structure in

section 3.1. Subsequently, we introduce the Structural Texture Knowledge Distillation and Statistical Texture Distillation in section 3.2 and 3.3, respectively. Finally, we provide the optimization process in section 3.4.

3.1. Overview

The overall framework of the proposed method is illustrated in Figure 2. The upper network is the teacher network while the lower one is the student network. Following the previous works [30, 37, 40], the PSPNet [50] architecture is used for both the teacher and student, and ResNet-101 and ResNet-18 [19] are used as their backbone respectively, which can be also changed to any other backbone networks. Firstly, we adopt the same basic idea in knowledge distillation to align the response-based knowledge between the teacher and student as previous works [30, 37, 40]. Specifically, we use the KL divergence to supervise the pixel-wise probability distribution and adversarial learning to strengthen the output segmentation maps. Furthermore, we propose to extract two kinds of the texture knowledge




Figure 3. LP decomposition [11, 12, 29]. The low-pass subbands a is generated from the input x with a low-pass analysis filters H and a sampling matrix S . The high-pass subbands b are then computed as the difference between x and the prediction of a , with a sampling matrix S followed by a low-pass synthesis filters G .

from the first two layers of ResNet backbone for both the teacher and student model, as the texture information is more reflected on low-level features. For structural texture knowledge, we introduce a Contourlet Decomposition Module (CDM) which exploits structural information in the spectral space. For statistical texture knowledge, we introduce a Denoised Texture Intensity Equalization Module (DTIEM) to adaptively extract the statistical texture intensity histogram with an adaptive importance sampler, then enhance it with a denoised operation and graph reasoning. Finally we optimize two kinds of the knowledge between the teacher and student model with two individual Mean Squared (L2) loss.

3.2. Structural Texture Knowledge Distillation

Traditional filters have inherent advantages for texture representation with different scales and directions in the spectral domain, and here we consider to utilizing the contourlet decomposition, a kind of multiscale geometric analysis tool which has substantial advantages in locality and directionality [11–14, 29, 36], and can enhance the ability of geometric transformations in CNN. Based on the advantages, we introduce a Contourlet Decomposition Module (CDM) [11, 29] to mine the texture knowledge in the spectral space. The light red part in Figure 2 shows the details of CDM. Specifically, it adopts a Laplacian Pyramid (LP) [5] and a directional filter bank (DFB) [1, 3, 9] iteratively on the low-pass image. The LP is aimed to obtain multiscale decomposition. As shown in Figure 3, given the input feature x , a low-pass analysis filters H and a sampling matrix S are used to generate downsampled low-pass subbands, then the high-pass subbands are obtained by the difference between the original x and the intermediate result, which is computed by a sampling matrix S and a low-pass synthesis filters G [5, 29]. Next, DFB is utilized to reconstruct the original signal with a minimum sample representation, which is generated by m -level binary tree decomposition in the two dimensional frequency domain, resulting in 2^m directional subbands [3]. For example, the frequency domain is divided into $2^3 = 8$ directional subbands when $m = 3$, and the subbands 0-3 and 4-7 correspond to the vertical and

horizontal details, respectively. Finally, the output of the contourlet decomposition in level n can be described by the following equations:

$$\begin{aligned} F_{l,n+1}, F_{h,n+1} &= LP(F_{l,n}) \downarrow p \\ F_{bds,n+1} &= DFB(F_{h,n+1}) \quad n \in [1, m] \end{aligned} \quad (1)$$

where the symbol \downarrow is downsampling operator, p denotes the interlaced downsampling factor, l and h represent the low-pass and high-pass components respectively, bds denotes the bandpass directional subbands.

For a richer expression, we stack multiple contourlet decomposition layers iteratively in the CDM. In this way, abundant bandpass directional features are obtained via the contourlet decomposition from the low-level features, which are used as the structural texture knowledge F^{str} for distillation. We apply the CDM to the teacher and student networks respectively and use the conventional Mean Squared (L2) loss to formulate the texture distillation loss:

$$L_{str}(S) = \frac{1}{(W \times H)} \sum_{i \in R} (F_i^{str;T} - F_i^{str;S})^2 \quad (2)$$

where $F_i^{str;T}$ and $F_i^{str;S}$ denote the i th pixel in texture features produced from the teacher network T and the student network S , and $i \in R = W \times H$ represents the feature size.

3.3. Statistical Texture Distillation

The statistical texture is usually of a wide variety and a continuous distribution in spectral domain, which is difficult to be extracted and optimized in deep neural networks. Previous work [52] firstly proposed a Quantization and Count Operator (QCO) to describe the statistical texture, which quantized the whole input feature into multiple uniform levels, then counted the number of features belonging to each level to get the quantization encoding matrix, followed by a graph module to perform quantization levels redistribution and texture enhancement. However, there are three limitations in the ordinary QCO. Firstly, it quantizes the input feature in a global manner, resulting in very sparse and discrete quantization levels distribution, which can not well balance the trade-off between accurate texture and computation burden. Secondly, amounts of pixels will never be quantized to any level as it generates the initial quantization levels by a uniform distribution, meanwhile using several narrow-peak functions to quantize the input feature. Thirdly, it may come across quantization noise produced by the over-amplify effect in near-constant areas during quantization process.

We propose to describe the statistical knowledge based on the ordinary QCO. Moreover, we propose three improvements in response to the three limitations as shown in the light green part in Figure 2. Firstly, for the sake of accurate characterization and efficient computation, we exploit an

Anchor-Based Adaptive Importance Sampler to only select the important regions for extraction. Secondly, we design a heuristics iterative method for a more equalized quantization level initialization. Thirdly, we utilize a denoised operation to suppress the over-amplify effect. Based on above analysis, we propose the Denoised Texture Intensity Equalization Module (DTIEM) to extract and enhance the statistical texture adaptively. In the following sections, we illustrate the above modules in detail.

Anchor-Based Adaptive Importance Sampler. Extracting the statistical texture of the whole input feature is straightforward, but it lacks the attention on discriminative regions and only works well when the feature intensities of pixels are near-uniformly distributed. However, real-world scenes are usually offended with chaotic conditions and the pixel intensities tend to be heavily imbalanced distributed. Moreover, the global operation on a whole image is not able to describe local intensity contrast accurately and always means a large computation burden. Thus we consider using an importance sampling method to mine the hard-to-classify areas, where the feature intensity distribution is discrete and varying, while the statistical texture is rich and diverse. Following the typical paradigm [25], we design an Anchor-Based Adaptive Importance Sampler. It is designed to bias selection towards most uncertain regions, while retaining some degree of uniform coverage, by the following steps [25]. (i) *Over Generation*: Aiming at sample M points, to guarantee the variety and recall, we over-generate candidate points by randomly selecting the kM ($k > 1$) points with a uniform distribution. (ii) *Importance sampling*: In the kM points, we are intended to choose out the most uncertain βM ($\beta \in [0, 1]$) ones with an anchor-based adaptive importance sampling strategy. (iii) *Coverage*: to balance the distribution, we select the remaining $(1 - \beta)M$ ones from the rest of points with a uniform distribution.

As found that the most essential step lies in *importance sampling*, we formulate this process as follows. For each sample $s_i \in kM$, we set several anchors with various scales and aspect ratios at the location. In this way, we generate ξ region proposals R_i for each s_i . For each $r_{ij} \in R_i$ ($i \in [1, kM], j \in [1, \xi]$), we calculate its sample probability by:

$$prob_{ij} = std(r_{ij}) \quad (3)$$

where $std(\cdot)$ means the variance function. It shows that the region with larger variance will be more likely sampled, as its intensity distribution is diverse and may have more rich statistical textures that need to be enhanced. Noted that the student utilizes the same importance sampling results as teacher to select region proposals.

Texture Extraction and Intensity Equalization. QCO is inspired of histogram quantification [18] and describes texture in a statistical manner. Given the input feature $\mathbf{A} \in \mathbb{R}^{C \times H \times W}$, it firstly calculates the self-similarity matrix \mathbf{S} ,

and quantizes it into N uniform levels \mathbf{L} . Next it counts the value for each quantization level and gets the quantization encoding matrix $\mathbf{E} \in \mathbb{R}^{N \times HW}$. Finally it utilizes a simple fully-connected graph to perform quantization levels redistribution and texture enhancement. The quantization can be formulated as:

$$\begin{aligned} \mathbf{L}_n &= \frac{\max(\mathbf{S}) - \min(\mathbf{S})}{N} \cdot n + \min(\mathbf{S}), \quad n \in [1, N] \\ \mathbf{E}_i^n &= \begin{cases} 1 - |L_n - \mathbf{S}_i|, & \text{if } -\frac{0.5}{N} \leq L_n - \mathbf{S}_i < \frac{0.5}{N} \\ 0 & \text{else} \end{cases} \end{aligned} \quad (4)$$

where $i \in [1, HW]$. Based on the above discussion, we propose the other two improvements here.

For the quantization level, a heuristics iterative method is adopted instead of a uniform quantization. Specifically, we use a simple t -step uniform sampling to formulate an attention sampled results. Firstly, we over-quantize the input into $2N$ levels to guarantee that most points can be quantized to one of the levels, obtaining the quantization encoding matrix with Eq. 4. Based on the count values of all quantization levels, we set an intensity ratio threshold δ to divide the quantization levels into two groups $G_{<\delta}$ and $G_{\geq\delta}$, which are then re-quantized into αN ($\alpha \in [0, 1]$) and $(1 - \alpha)N$ levels to obtain final N levels. The above process can be performed iteratively. In this way, the imbalanced distribution problem among quantization levels can be weakened. Moreover, for each quantization function in Eq. 4, we widen its peak adaptively by replacing the N with its current group levels number, as each intermediate group level number is much smaller than N , thus the cover range of each quantization function can be broadened.

Intensity-Limited Denoised Strategy. For the noise over-amplity problem, inspired by the Contrast-Limited Adaptive Histogram Equalization [35], we propose an intensity-limited denoised strategy to constrain the intensity peak and redistribute the extra peak to all quantization levels dynamically. More comprehensively, for each selected region and N quantization levels, we get the initial quantization encoding matrix \mathbf{E} by Eq. 4. Due to the near-constance sub-area, the count value of some levels may show in a very high peak [35], we perform an intensity-clip operation to limit these peaks with a given ratio threshold θ , and then redistribute the sum of extra peak values \mathbf{E}_{extra} equally among all the quantization levels. Finally, the denoised quantization encoding matrix \mathbf{DE} can be calculated:

$$\mathbf{E}_{extra} = \sum_n^N [\max(\mathbf{E}^n - \theta \cdot \max(\mathbf{E}), 0)] \quad (5)$$

$$\mathbf{DE}^n = \begin{cases} \theta \cdot \max(\mathbf{E}) + \frac{\mathbf{E}_{extra}}{N} & \text{if } \mathbf{E}^n > \theta \cdot \max(\mathbf{E}) \\ \mathbf{E}^n + \frac{\mathbf{E}_{extra}}{N} & \text{else} \end{cases} \quad (6)$$

where $n \in [1, N]$, $\max(\cdot)$ is the maximum function. Then **DE** is also enhanced with the graph reasoning as the previous work [52] to get the statistical texture knowledge F^{sta} . Similarly as the structural texture knowledge distillation, we apply the DTIEM to both teacher and student respectively and use the L2 loss to formulate the distillation loss:

$$L_{sta}(S) = \frac{1}{(W \times H)} \sum_{i \in R} (F_i^{sta;T} - F_i^{sta;S})^2 \quad (7)$$

where $F_i^{sta;T}$ and $F_i^{sta;S}$ denote the statistical texture knowledge of the teacher and student respectively.

3.4. Optimization

Following the common practice and previous knowledge distillation works for semantic segmentation [30, 40], we also add the fundamental response-based distillation loss L_{re} and adversarial loss L_{adv} for stable gradient descent optimization:

$$L_{re} = \frac{1}{(W_{re} \times H_{re})} \sum_{i \in R} KL(P_i^{re;T} || P_i^{re;S}) \quad (8)$$

where $P_i^{re;T}$ and $P_i^{re;S}$ denote the class probabilities of i -th pixel produced by the teacher and student model respectively, and $i \in R = W_{re} \times H_{re}$ represents the output size. The adversarial training is aimed to formulate the holistic distillation problem [30, 43] and denoted as follows:

$$L_{adv} = \mathbb{E}_{S \sim p(S)} [D(S|I)] \quad (9)$$

where $\mathbb{E}(\cdot)$ is the expectation operator, and $D(\cdot)$ is the discriminator. I and S are the input image and corresponding segmentation map respectively.

Therefore, for the overall optimization, the whole objective function consists of a conventional cross-entropy loss L_{seg} for semantic segmentation and the abovementioned distillation loss:

$$L = L_{seg} + \lambda_1 L_{str} + \lambda_2 L_{sta} + \lambda_3 L_{re} - \lambda_4 L_{adv} \quad (10)$$

where $\lambda_1, \lambda_2, \lambda_3, \lambda_4$ are set to 0.9, 1.15, 5, 0.01, respectively.

4. Experiments

4.1. Datasets and Evaluation Metrics

To verify the effectiveness of the proposed method, we conduct experiments on the following large-scale datasets. **Cityscapes.** The Cityscapes dataset [10] has 5,000 images captured from 50 different cities, and contains 19 semantic classes. Each image has 2048×1024 pixels, which have high quality pixel-level labels of 19 semantic classes. There are 2,979/500/1,525 images for training, validation and testing.

Method	mIOU(%)
T: PSPNet-R101	78.56
S: PSPNet-R18	69.10
+Response Knowledge	72.47
+Response+Structural Texture Knowledge	74.10
+Response+Statistical Texture Knowledge	74.69
+Response+Structural+Statistical Texture Knowledge	75.15

Table 1. Efficacy of two kinds of the texture knowledge.

ADE20K. The ADE20K dataset has 20K/2K/3K images for training, validation, and testing, and contains 150 classes of diverse scenes.

Pascal VOC 2012. The Pascal VOC 2012 dataset [15] is a segmentation benchmark of 10,582/1,449/1,456 images for training, validation and testing, which involves 20 foreground object classes and one background class.

Evaluation Metrics. In all experiments, we adopt the mean Intersection-over-Union (mIoU) to study the distillation effectiveness. The model size is represented by the number of network parameters, and the Complexity is evaluated by the sum of floating point operations (FLOPs) in one forward propagation on a fixed input size.

4.2. Implementation Details

Following [30, 37, 40], we adopt PSPNet [50] with ResNet101 [19] backbone as the teacher network, and use PSPNet with different compact backbones as the student networks, including Resnet18 [19] and EfficientNet-B1 [38], which also validates the effectiveness when the teacher model and the student model are of different architectural types. In this paper, we adopt an offline distillation method, first train the teacher model and then keep the parameters frozen during the distillation progress. In the training process of the student network, random scaling (from 0.5 to 2.1) and random horizontal flipping (with the probability of 0.5) are applied as the data augmentation. We implement two-level contourlet decomposition iteratively in the CDM, where m is set to 4 and 3 respectively. We set $N = 50, \alpha = 0.3, \theta = 0.9$. Stochastic Gradient Descent with momentum is deployed as the optimizer, where the momentum is 0.9 and weight decay rate is $1e-5$. The base learning rate is 0.015 and multiplied by $(1 - \frac{iter}{max-iter})^{0.9}$. We train the model for 80000 iterations with the batch size of 16.

4.3. Ablation Study

In all the ablation studies, we use Cityscapes validation dataset, and ResNet-18 pretrained from ImageNet as the backbone of the student network.

Efficacy of Two Kinds of the Texture Knowledge. Table 1 shows the effectiveness of two kinds of the texture knowl-

Levels number in CDM	val mIoU(%)
baseline	72.47
1	73.44
2	74.10
3	74.11

Table 2. The impact of layer number of comtourlet decomposition.

Method	Params (M)	FLOPs (G)
PSPNet	70.43	574.9
CDM	1.24	10.90
DTIEM	2.80	23.73

Table 3. The FLOPs and Parameters of the proposed texture modules.

Components in DTIEM	mIOU(%)
baseline	72.47
+Global	73.57
+Adap. Samp.	74.06
+Adap. Samp. + Heuristics Init.	74.25
+Adap. Samp. + Heuristics Init. + Denoised	74.69

Table 4. Abalation Study of Statistiacal Knowledge. “Global” means global operation without sampling, “Adap. Samp.” means anchor-based adaptive importance sampling, “Heuristics Init.” means quantization levels heuristics initialization.

edge. The student network without distillation achieves the result of 69.10%, and the response knowledge improves it to 72.47%. Then we further add two kinds of the texture knowledge successively to validate the effect of each one. Concretely, it shows that the structural texture knowledge brings the improvement to 5.0%, and the statistical texture knowledge boosts the improvement to 5.59%. Finally, when we add both the texture knowledge, the performance is promoted to 75.15% with a large increase of 6.05%. The gap between student and teacher is finally reduced, providing a closer result to the teacher network.

Analysis of Structural Texture Knowledge. We conduct experiments to verify the effectiveness of the components in the structural texture module, and show the impact of level number of comtourlet decomposition in the CDM. As shown in Table 2, “baseline” means the results of student network with response knowledge, with the level number gradually increasing, mIOU is gradually increasing and stays still around 74.1%, which indicates that the texture knowledge is almost saturated when the level number gets to 2.

Analysis of Statistical Texture Knowledge. In Table 4, we show the effectness of different components in the statistical texture module, “baseline” means the results of student network with response knowledge. Firstly, we perform a global

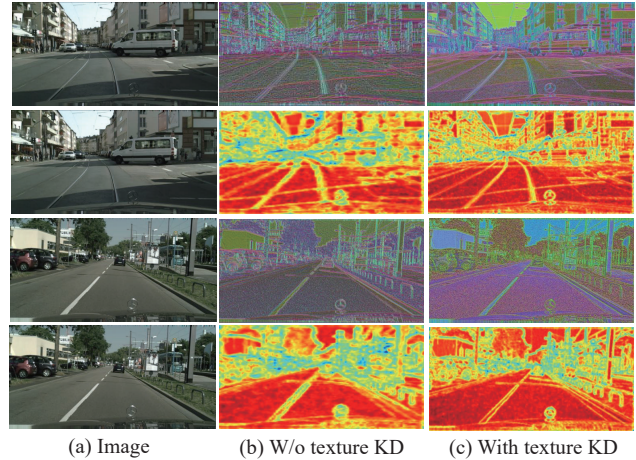


Figure 4. Comparison of visualization of the low-level feature from stage 1 of the backbone. KD means knowledge distillation. (a) is the original image. (b) is from the student network without texture knowledge distillation, (c) shows the changes after applying it in our method. Line 1 and 3 show the structural texture, while line 2 and 4 show the statistical texture.

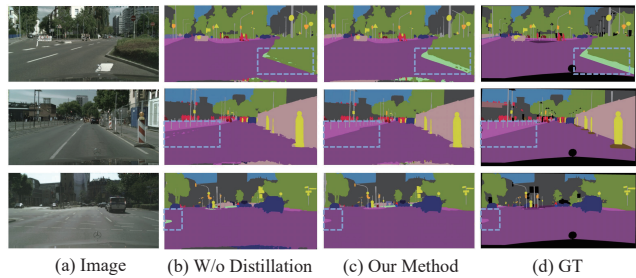


Figure 5. Visual improvements on Cityscapes dataset: (a) original images, (b) w/o distillation, (c) Our distillation method, (d) ground truth. Our method improves the student network w/o distillation to produce more accurate and detailed results, which are circled by dotted lines.

feature intensity equalization on the whole image and get a limited improvement. Then we add the adaptive sampling, quantization level heuristics initialization and the denoised strategy successively, and the performance is gradually improved in this case.

Complexity of Texture Extraction Modules. We show that the proposed texture modules are lightweight in Table 3, which are estimated with the fixed input size. It shows that the CDM and DTIEM only bring very little extra cost compared to the PSPNet.

Visualization. Figure 4 shows the visualization of the low-level feature from the student network with ResNet18 backbone. For comparison, we give the results of the student network w/ and w/o the texture knowledge distillation in our method. Specifically, Figure 4 (b) shows the results which are produced without the structural/statistical texture

Method	Cityscapes mIOU(%)		Params (M)	Flops (G)
	<i>val</i>	<i>test</i>		
ENet [34]	-	58.3	0.358	3.612
ICNet [49]	-	69.5	26.50	28.30
FCN [32]	-	62.7	134.5	333.9
RefineNet [27]	-	73.6	118.1	525.7
OCNet [47]	-	80.1	62.58	548.5
STLNet [52]	82.3	82.3	81.39	293.03
Results w/ and w/o distillation schemes				
T: PSPNet-R101 [50]	78.56	78.4	70.43	574.9
S: PSPNet-R18	69.10	67.60	13.07	125.8
+ SKDS [30]	72.70	71.40	13.07	125.8
+ SKDD [31]	74.08	-	13.07	125.8
+ IFVD [40]	74.54	72.74	13.07	125.8
+ CWD [37]	74.87	-	13.07	125.8
+ SSTKD	75.15	74.39	13.07	125.8
S: Deeplab-R18	73.37	72.39	12.62	123.9
+ SKDS [30]	73.87	72.63	12.62	123.9
+ IFVD [40]	74.09	72.97	12.62	123.9
+ CWD [37]	75.91	74.32	12.62	123.9
+ SSTKD	76.13	75.01	12.62	123.9
S: EfficientNet-B1	60.40	59.91	6.70	9.896
+ SKDS [30]	63.13	62.59	6.70	9.896
+ IFVD [40]	66.50	64.42	6.70	9.896
+ CWD [37]	-	-	6.70	9.896
+ SSTKD	68.26	65.77	6.70	9.896

Table 5. Quantitative results on Cityscapes. “R18”(“R101”) means ResNet-18(ResNet-101).

knowledge. As we can see, the texture details are very fuzzy and in low-contrast, the contours of the objects are also incomplete. By contrast, they are clearer in Figure 4 (c) when incorporating with two kinds of the texture knowledge. In this case, the contours of the auto logo and lane lines can offer more correct cues for semantic segmentation. The phenomenon illustrates the validity of the texture knowledge, providing a better understanding of our method. Besides, we also visualize the results of the different methods in Figure 5 for a better comparison.

4.4. Comparison with State-of-the-arts

Cityscapes. Table 5 shows that the proposed SSTKD framework achieves state-of-the-art results with different backbones in Cityscapes validation and test datasets. More comprehensively, SSTKD improves the student model(PSPNet) built on ResNet-18 to 75.15% and 74.39% on validation and test dataset respectively. Moreover, we also change the student backbone to Deeplab and EfficientNet-B1, which shows the generality of SSTKD. Besides, the experimental results also show that we improve the baseline by a large margin.

Method	Pascal VOC mIOU(%)	ADE20K mIOU(%)	Params (M)
FCN [32]	69.6	39.91	134.5
RefineNet [27]	82.4	40.7	118.1
Deeplab V3 [6]	77.9	44.99	87.1
PSANet [51]	77.9	43.47	78.13
OCRNet [46]	80.3	43.7	70.37
Results w/ and w/o distillation schemes			
T: PSPNet-R101 [50]	78.52	44.39	70.43
S: PSPNet-R18	65.42	24.65	13.07
+SKDS [30]	67.73	25.11	13.07
+IFVD [40]	68.04	25.72	13.07
+CWD [31]	69.25	26.80	13.07
+SSTKD	70.98	29.19	13.07
S: Deeplab-R18	66.81	24.89	12.62
+SKDS [30]	68.13	25.52	12.62
+IFVD [40]	68.42	26.53	12.62
+CWD [31]	69.97	27.37	12.62
+SSTKD	71.45	29.79	12.62

Table 6. Quantitative results on Pascal VOC 2012 and ADE20K. “R18”(“R101”) means ResNet-18(ResNet-101)

Pascal VOC 2012 and ADE20K. Table 6 shows the comparisons with state-of-the-art methods on PASCAL VOC 2012 and ADE20K validation dataset. Experimental results show that the proposed method improves the performance of the student network without distillation, meanwhile surpasses previous works in all cases in spite of the choice of architectures and backbones for student networks.

5. Conclusion

In this paper, we focus on the low-level structural and statistical knowledge in distillation for semantic segmentation. Specifically, we introduce the Contourlet Decomposition Module to effectively extract the structural texture knowledge, and the Denoised Texture Intensity Equalization Module to describe and enhance statistical texture knowledge, respectively. Under different supervisions, we force the student network to mimic the teacher network better from a broader perspective. Experimental results show that we achieve new state-of-the-art results on three semantic segmentation datasets, which proves the effectiveness and superiority of our method.

Acknowledgements

This work was partially supported by the National Key R&D Program of China under Grant 2020AAA0103902, NSFC(No. 62176155 and No. 61772330), Science and Technology Commission of Shanghai Municipality under Grant 2021SHZDZX0102.

References

- [1] Vincent Andrearczyk and Paul F Whelan. Using filter banks in convolutional neural networks for texture classification. *Pattern Recognition Letters*, 84:63–69, 2016.
- [2] Vijay Badrinarayanan, Alex Kendall, and Roberto Cipolla. Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 39(12):2481–2495, 2017.
- [3] Roberto H Bamberger and Mark JT Smith. A filter bank for the directional decomposition of images: Theory and design. *IEEE transactions on signal processing*, 40(4):882–893, 1992.
- [4] Cristian Bucilu, Rich Caruana, and Alexandru Niculescu-Mizil. Model compression. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 535–541, 2006.
- [5] Peter Burt and Edward Adelson. The laplacian pyramid as a compact image code. *IEEE Transactions on communications*, 31(4):532–540, 1983.
- [6] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2017.
- [7] Liang-Chieh Chen, Yi Yang, Jiang Wang, Wei Xu, and Alan L Yuille. Attention to scale: Scale-aware semantic image segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3640–3649, 2016.
- [8] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 801–818, 2018.
- [9] Mircea Cimpoi, Subhansu Maji, and Andrea Vedaldi. Deep filter banks for texture recognition and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3828–3836, 2015.
- [10] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016.
- [11] Minh N Do and Martin Vetterli. Contourlets: a directional multiresolution image representation. In *Proceedings. International Conference on Image Processing*, volume 1, pages I–I. IEEE, 2002.
- [12] Minh N Do and Martin Vetterli. Contourlets, beyond wavelets. *New York: Academic*, 2003.
- [13] Minh N Do and Martin Vetterli. The contourlet transform: an efficient directional multiresolution image representation. *IEEE Transactions on image processing*, 14(12):2091–2106, 2005.
- [14] David Leigh Donoho and Ana Georgina Flesia. Can recent innovations in harmonic analysis explain key findings in natural image statistics? *Network: computation in neural systems*, 12(3):371–393, 2001.
- [15] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88(2):303–338, 2010.
- [16] Clement Farabet, Camille Couprie, Laurent Najman, and Yann LeCun. Learning hierarchical features for scene labeling. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1915–1929, 2012.
- [17] Jun Fu, Jing Liu, Haijie Tian, Yong Li, Yongjun Bao, Zhiwei Fang, and Hanqing Lu. Dual attention network for scene segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3146–3154, 2019.
- [18] Rafael C Gonzalez, Richard E Woods, et al. Digital image processing, 2002.
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [20] Tong He, Chunhua Shen, Zhi Tian, Dong Gong, Changming Sun, and Youliang Yan. Knowledge adaptation for efficient semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 578–587, 2019.
- [21] Xuming He, Richard S Zemel, and Miguel Á Carreira-Perpiñán. Multiscale conditional random fields for image labeling. In *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004.*, volume 2, pages II–II. IEEE, 2004.
- [22] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.
- [23] Hanzhe Hu, Deyi Ji, Weihao Gan, Shuai Bai, Wei Wu, and Junjie Yan. Class-wise dynamic graph convolution for semantic segmentation. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVII 16*, pages 1–17. Springer, 2020.
- [24] Deyi Ji, Haoran Wang, Hanzhe Hu, Weihao Gan, Wei Wu, and Junjie Yan. Context-aware graph convolution network for target re-identification. *arXiv preprint arXiv:2012.04298*, 2020.
- [25] Alexander Kirillov, Yuxin Wu, Kaiming He, and Ross Girshick. Pointrend: Image segmentation as rendering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9799–9808, 2020.
- [26] L’ubor Ladický, Chris Russell, Pushmeet Kohli, and Philip HS Torr. Associative hierarchical crfs for object class image segmentation. In *2009 IEEE 12th International Conference on Computer Vision*, pages 739–746. IEEE, 2009.
- [27] Guosheng Lin, Anton Milan, Chunhua Shen, and Ian Reid. Refinenet: Multi-path refinement networks for high-resolution semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1925–1934, 2017.
- [28] Guosheng Lin, Chunhua Shen, Anton Van Den Hengel, and Ian Reid. Exploring context with deep structured models for

- semantic segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 40(6):1352–1366, 2017.
- [29] Mengkun Liu, Licheng Jiao, Xu Liu, Lingling Li, Fang Liu, and Shuyuan Yang. C-cnn: Contourlet convolutional neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 32(6):2636–2649, 2020.
- [30] Yifan Liu, Ke Chen, Chris Liu, Zengchang Qin, Zhenbo Luo, and Jingdong Wang. Structured knowledge distillation for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2604–2613, 2019.
- [31] Yifan Liu, Changyong Shu, Jingdong Wang, and Chunhua Shen. Structured knowledge distillation for dense prediction. *IEEE transactions on pattern analysis and machine intelligence*, 2020.
- [32] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015.
- [33] Mohammadreza Mostajabi, Payman Yadollahpour, and Gregory Shakhnarovich. Feedforward semantic segmentation with zoom-out features. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3376–3385, 2015.
- [34] Adam Paszke, Abhishek Chaurasia, Sangpil Kim, and Eugenio Culurciello. Enet: A deep neural network architecture for real-time semantic segmentation. *arXiv preprint arXiv:1606.02147*, 2016.
- [35] Stephen M Pizer, E Philip Amburn, John D Austin, Robert Cromartie, Ari Geselowitz, Trey Greer, Bart ter Haar Romeny, John B Zimmerman, and Karel Zuiderveld. Adaptive histogram equalization and its variations. *Computer vision, graphics, and image processing*, 39(3):355–368, 1987.
- [36] Yuheng Sha, Lin Cong, Qiang Sun, and Licheng Jiao. Unsupervised image segmentation using contourlet domain hidden markov trees model. In *International Conference Image Analysis and Recognition*, pages 32–39. Springer, 2005.
- [37] Changyong Shu, Yifan Liu, Jianfei Gao, Zheng Yan, and Chunhua Shen. Channel-wise knowledge distillation for dense prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5311–5320, 2021.
- [38] Mingxing Tan and Quoc V Le. Efficientnet: Rethinking model scaling for convolutional neural networks. *arXiv preprint arXiv:1905.11946*, 2019.
- [39] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.
- [40] Yukang Wang, Wei Zhou, Tao Jiang, Xiang Bai, and Yongchao Xu. Intra-class feature variation distillation for semantic segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, 2020.
- [41] Shinji Watanabe, Takaaki Hori, Shigeki Karita, Tomoki Hayashi, Jiro Nishitoba, Yuya Unno, Nelson Enrique Yalta Soplín, Jahn Heymann, Matthew Wiesner, Nanxin Chen, et al. Espnet: End-to-end speech processing toolkit. *arXiv preprint arXiv:1804.00015*, 2018.
- [42] Jiafeng Xie, Bing Shuai, Jian-Fang Hu, Jingyang Lin, and Wei-Shi Zheng. Improving fast segmentation with teacher-student learning. *arXiv preprint arXiv:1810.08476*, 2018.
- [43] Zheng Xu, Yen-Chang Hsu, and Jiawei Huang. Training shallow and thin networks for acceleration via knowledge distillation with conditional adversarial networks. *International Conference on Learning Representations Workshop*, 2018.
- [44] Maoke Yang, Kun Yu, Chi Zhang, Zhiwei Li, and Kuiyuan Yang. Denseaspp for semantic segmentation in street scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3684–3692, 2018.
- [45] Changqian Yu, Jingbo Wang, Chao Peng, Changxin Gao, Gang Yu, and Nong Sang. Bisenet: Bilateral segmentation network for real-time semantic segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 325–341, 2018.
- [46] Yuhui Yuan, Xilin Chen, and Jingdong Wang. Object-contextual representations for semantic segmentation. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VI 16*, pages 173–190. Springer, 2020.
- [47] Yuhui Yuan and Jingdong Wang. Ocnet: Object context network for scene parsing. *arXiv preprint arXiv:1809.00916*, 2018.
- [48] Sergey Zagoruyko and Nikos Komodakis. Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. *arXiv preprint arXiv:1612.03928*, 2016.
- [49] Hengshuang Zhao, Xiaojuan Qi, Xiaoyong Shen, Jianping Shi, and Jiaya Jia. Icnets for real-time semantic segmentation on high-resolution images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 405–420, 2018.
- [50] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2881–2890, 2017.
- [51] Hengshuang Zhao, Yi Zhang, Shu Liu, Jianping Shi, Chen Change Loy, Dahua Lin, and Jiaya Jia. Psanet: Point-wise spatial attention network for scene parsing. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018.
- [52] Lanyun Zhu, Deyi Ji, Shiping Zhu, Weihao Gan, Wei Wu, and Junjie Yan. Learning statistical texture for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12537–12546, June 2021.