

HGFormer: Hierarchical Grouping Transformer for Domain Generalized Semantic Segmentation

Jian Ding^{1,2,3}, Nan Xue¹, Gui-Song Xia^{1,2,*}, Bernt Schiele³, Dengxin Dai³

¹NERCMS, School of Computer Science, Wuhan University, China

²State Key Lab. LIESMARS, Wuhan University, China

³Max Planck Institute for Informatics, Saarland Informatics Campus, Germany

{jian.ding, xuenan, guisong.xia}@whu.edu.cn, {schiele, ddai}@mpi-inf.mpg.de

Abstract

Current semantic segmentation models have achieved great success under the independent and identically distributed (i.i.d.) condition. However, in real-world applications, test data might come from a different domain than training data. Therefore, it is important to improve model robustness against domain differences. This work studies semantic segmentation under the domain generalization setting, where a model is trained only on the source domain and tested on the unseen target domain. Existing works show that Vision Transformers are more robust than CNNs and show that this is related to the visual grouping property of self-attention. In this work, we propose a novel hierarchical grouping transformer (HGFormer) to explicitly group pixels to form part-level masks and then whole-level masks. The masks at different scales aim to segment out both parts and a whole of classes. HGFormer combines mask classification results at both scales for class label prediction. We assemble multiple interesting cross-domain settings by using seven public semantic segmentation datasets. Experiments show that HGFormer yields more robust semantic segmentation results than per-pixel classification methods and flat-grouping transformers, and outperforms previous methods significantly. Code will be available at <https://github.com/dingjiansw101/HGFormer>.

1. Introduction

Research in semantic image segmentation has leaped forward in the past years due to the development of deep neural network. However, most of these models assume that the training and testing data follow the same distribution. In the real-world, we frequently encounter testing data that is out of distribution. The generalization ability of models under distribution shift is crucial for applications related to safety,

*Corresponding author




Figure 1. Semantic segmentation can be considered as partitioning an image into classification units (regions), then classifying the units. The units can range from pixels to large masks. Intuitively, mask classification is more robust than per-pixel classification, as masks allow to aggregate features over large image regions of the same class to predict a ‘global’ label. Despite this promise, the process of grouping pixels into *whole-level* masks directly from pixels is very challenging under the distribution shift (e.g., Gaussian Noise). In order to tackle this problem, we present a hierarchical grouping paradigm to group pixels to *part-level* masks first and then to group *part-level* masks to *whole-level* masks to get reliable masks. Then we combine both part-level and whole-level mask classification for robust semantic segmentation, given that the masks at the two levels capture complementary information.

such as self-driving. In domain generalization setting, models are trained only on source domains and tested on target domains, where the distributions of source domains and target domains are different. Unlike the domain adaptation [25, 56], target data is not accessible / needed during training, making the task challenging but practically useful.

Recently, Vision Transformers have been shown to be significantly more robust than traditional CNNs in the out-of-distribution generalization [21, 25, 42, 58, 60, 70]. Some works interpret self-attention as a kind of *visual group-*

ing [7, 38], and believe that it is related to robustness [70]. However, these works mainly focus on classification. Although FAN [70] and Segformer [60] have been evaluated on segmentation, they do not explicitly introduce visual grouping in their networks. Since grouping is naturally aligned with the task of semantic segmentation, we would like to ask the question: *can we improve the robustness of semantic segmentation by introducing an explicit grouping mechanism into semantic segmentation networks?*

Most deep learning based segmentation models directly conduct per-pixel classification without the process of grouping. Some recent segmentation models introduced *flat grouping* [15, 67] into the segmentation decoder, where pixels are grouped into a set of binary masks directly and classification on masks is used to make label prediction. By using a one-to-one matching similar to DETR [5], the loss between predicted masks and ground truth masks is computed. Therefore the network is trained to directly predict *whole-level masks*, as shown in Fig. 1. Intuitively, if whole-level masks are accurate, mask classification will be more robust than per-pixel classification due to its information aggregation over regions of the same class. But we find that using the *flat grouping* to generate whole-level masks is susceptible to errors, especially under cross-domain settings. This is shown by the example in Fig. 1 - bottom.

Different from the *flat grouping* works [14, 67], we propose a *hierarchical grouping* in the segmentation decoder, where the pixels are first grouped into *part-level masks*, and then grouped into *whole-level masks*. Actually, the hierarchical grouping is inspired by the pioneer works of image segmentation [2, 13, 49] and is further supported by strong psychological evidence that humans parse scenes into part-whole hierarchies [24]. We find that grouping pixels to part-level masks and then to whole-level masks is more robust than grouping pixels directly to whole-level masks. Part-level masks and whole-level masks segment images at different scales such as *parts* and *a whole* of classes. Therefore, part-level and whole-level masks are complementary, and combining mask classification results at those different scales improves the overall robustness.

To instantiate a hierarchical grouping idea, we propose a hierarchical grouping transformer (HGFormer) in the decoder of a segmentation model. The diagram is shown in Fig. 2. We first send the feature maps to the part-level grouping module. In the part-level grouping module, the initialization of cluster centers is down sampled from feature maps. Then we compute the *pixel-center similarities* and assign pixels to cluster centers according to the similarities. To get the part-level masks, we only compute the similarities between each pixel feature and its nearby center features. We then aggregate information of the part-level masks and generate whole-level masks by using cross-attention, similar to how previous methods aggregate pixels

information to generate whole-level masks [14, 67]. Finally, we classify masks at different levels, and average the semantic segmentation results of all the scales.

We evaluate the method under multiple settings, which are assembled by using seven challenging semantic segmentation datasets. In each of the setting, we train the methods on one domain and test them on other domains. Extensive experiments show that our model is significantly better than previous *per-pixel classification based*, and *whole-level mask based* segmentation models for out-of-distribution generalization.

To summarize, our contributions are: 1) We present a hierarchical grouping paradigm for robust semantic segmentation; 2) based on the hierarchical grouping paradigm, we propose a hierarchical grouping transformer (HGFormer), where the pixels are first grouped into part-level masks, and then grouped into whole-level masks. Final semantic segmentation results are obtained by making classifications on all masks; 3) HGFormer outperforms previous semantic segmentation models on domain generalized semantic segmentation across various experimental settings. We also give detailed analyses of the robustness of grouping-based methods under distribution shift.

2. Related Work

2.1. Semantic Segmentation

Semantic segmentation is a classic and fundamental problem in computer vision. It aims to segment the objects and scenes in images and give their classifications. In the deep learning era, semantic segmentation is usually formulated as a pixel-level classification problem [9–12, 36, 60] since FCN [36]. Recently, it is becoming popular to use whole-level mask classification to formulate the semantic segmentation problem [14, 15, 54, 66, 67]. In contrast to pixel-level and mask-level classification, to our best knowledge, there are very few works on learning part-level masks [27, 62], and using part-level mask classification [20, 69] for semantic segmentation in deep learning era. Among them, SSN [27] and super pixel FCN [62] mainly focus on part-level mask learning instead of part-level classification for the semantic segmentation results. BI [20] is not an end-to-end model, which needs extra part-level masks as input. RegProxy [69] is a recent work that closes to our work, which uses convolutions to learn part-level masks, and is only evaluated on the i.i.d. condition. In contrast, we use similarity-based grouping to learn part-level masks, and are the first to validate the effectiveness of using part-level mask classification for domain generalized semantic segmentation. Besides, Regproxy [69] is customized with *plain ViT* [51], while our work is applicable to pyramid transformers [35, 57] and CNNs. Our work is also different for the hierarchical segmentation design.

2.2. Domain Generalization

Domain generalization (DG) assumes that the target data (even unlabelled) is not accessible during training. Methods for DG in classification includes domain alignment [31,37], meta-learning [3,30], data augmentation [52,61], ensemble learning [8,34,63], self-supervised learning [4,6], and regularization strategies [26,53]. The ensemble of mask classification at different levels is related to the ensemble methods for domain generalization. The drawback of the previous ensemble-based methods [29,64] is that they will largely increase the runtime. Some ensemble methods [45,59] focus on the averaging of model weights, which do not increase the runtime, but increase the training time. Our method does not introduce extra FLOPS due to the efficient hierarchical grouping design, and does not introduce extra training time.

While the DG in classification is widely studied in the previous works, there are only several works that study the DG in semantic segmentation. The previous methods for DG in semantic segmentation includes: (1) Domain Randomization [44,68] and (2) Normalization and Whitening [16,40,41,43]. Although not designed specifically for domain generalization tasks, the Vision Transformers [18] have shown their robustness [70] in the out-of-distribution setting. The robustness of Vision Transformer was explained to be related to the grouping property of self-attention [70]. However, there are no works that study the effect of explicit grouping in the segmentation decoder for semantic segmentation in DG. Motivated by these results, we study the different levels of grouping and mask classification for semantic segmentation in DG.

3. Methods

Given an image $I \in \mathbb{R}^{H \times W \times 3}$, an image partition is defined as $S = \{R_1, \dots, R_N\}$, such that $\cup_{i=1}^N R_i = \Omega$ and $R_i \cap R_j = \emptyset$, if $i \neq j$. After mapping each region R_i to a class by L_i , we get $C = \{L_1(R_1), \dots, L_N(R_N)\}$. Semantic segmentation can then be defined as:

$$Y = \{S, C\}. \quad (1)$$

According to the scales, we can roughly divide the image partitions into three levels: pixels, part-level masks, and whole-level masks. The *pixel* partition, where each element in S is a pixel, is widely used by all per-pixel classification methods [10,36]. The *whole-level masks* partition, where each element in S represents a whole mask of a class, is used by a few recent approaches [14,15]. The *part-level mask* partition, where each element aims to cover a class part, is proposed by this work and used along with *whole-level masks* to enhance robustness.

Intuitively, mask classification is more robust than per-pixel classification, as *masks* allow to aggregate features over large image regions of the same class to predict a

Algorithm 1 Part-level grouping

Require: *Pixel* feature map $\mathbf{K} \in \mathbb{R}^{(H \times W) \times d}$, classification feature map $\mathbf{V} \in \mathbb{R}^{(H \times W) \times d}$

- 1: Initialize the cluster *center* features $\mathbf{Q}^1 \in \mathbb{R}^{N_p \times d}$ by down sampling \mathbf{K}
- 2: **for** $t = 1, \dots, L$ **do**
- 3: Compute assignment matrix \mathbf{A}^t by \mathbf{Q}^t and \mathbf{K}
- 4: Update the cluster center features $\mathbf{Q}^{t+1} = \mathbf{A}^t \times \mathbf{K}$
- 5: Update the part-level tokens $\mathbf{Z}^t = \mathbf{A}^t \times \mathbf{V}$
- 6: **end for**

‘global’ label. Despite of this promise, generating whole-level masks directly from pixels is a very challenging task. While SOTA methods [14,15] can generate reasonable *whole-level masks* directly from pixels, the *flat grouping* methods used are not robust to domain changes – when tested on a different domain, the generated masks are of poor quality, leading to low semantic segmentation performance (see Fig. 1). In order to tackle this problem, this work proposes using hierarchical grouping in the segmentation transformer architecture to group pixels to *part-level masks* first and then to group *part-level masks* to *whole-level masks*. The advantages are 1) the grouping task becomes less challenging; 2) mask classification can be performed at both scales and their results can be combined for more robust label prediction, given that the *masks* at the two scales capture complementary information; 3) global self-attention to aggregate long-range context information can be performed directly at both part-level and whole-level now, which is not possible at pixel-level due to the huge computation cost.

3.1. HGFormer

To efficiently implement a model which can predict semantic segmentation at different scales, we adopt a hierarchical grouping process, which consists of two stages. The first is to group pixels into part-level masks by similarity-based local clustering. The second is to group part-level masks into whole-level masks by cross-attention. Then we make classification on partitions at different scales. The framework can be seen in Fig. 2. We will introduce the details in the following.

Part-level grouping. The goal of part-level grouping is to compute a partition $S_m = \{R_1, \dots, R_{N_p}\}$, with N_p the number of part-level masks. S_m can be represented by a *hard assignment* matrix $\tilde{\mathbf{A}} \in \{0, 1\}^{N_p \times (HW)}$ such that $\tilde{\mathbf{A}}_{ij} = 1$ if the j -th pixel is assigned to mask i and 0 otherwise. Since the hard assignment is not differentiable, we compute a *soft assignment* matrix $\mathbf{A} \in [0, 1]^{N_p \times (HW)}$ such that $\sum_{i=1}^{N_p} \mathbf{A}_{ij} = 1$. \mathbf{A}_{ij} represents the probability of assigning the j -th pixel to mask i .

To compute \mathbf{A} , we perform an iterative grouping al-

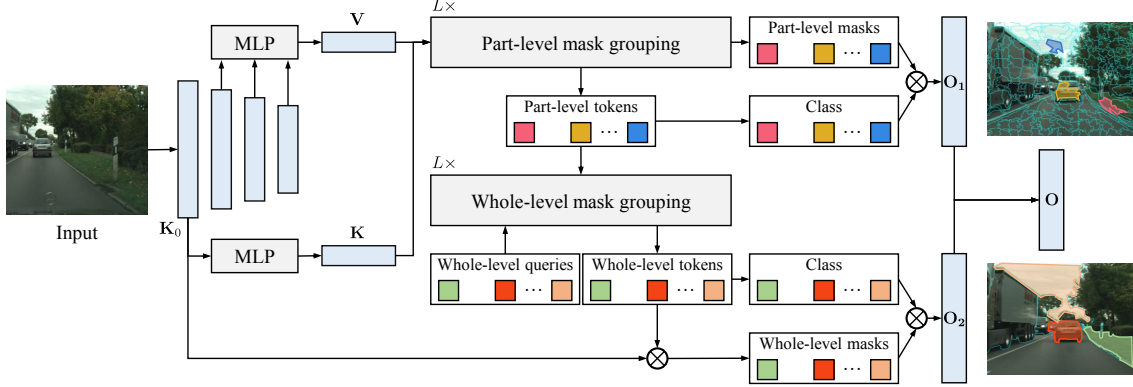


Figure 2. The pipeline of our proposed method. We first pass an image to a backbone network and get feature maps at different resolutions. The largest feature map \mathbf{K}_0 is projected to \mathbf{K} for part-level grouping. The other three feature maps are fused to form a new feature map \mathbf{V} for part-level mask feature extraction used for later classification. The details of part-level grouping can be seen in Algorithm 1. The grouping process is repeated L iterations. At the end of each iteration, there are N_p part-level masks, and their tokens. Combining part-level classifications and part-level masks, we can get the semantic segmentation results \mathbf{O}_1 . The part-level tokens from the last iteration of part-level grouping are aggregated to whole-level masks by whole-level grouping (which are actually cross-attention layers). Similarly, there are also L iterations in the whole-level grouping. At the end of each iteration, there are N_o whole-level tokens. Whole-level masks are computed by a matrix multiplication between \mathbf{K}_0 and projected whole-level mask tokens. Similarly, we can get semantic segmentation results \mathbf{O}_2 by combining whole-level masks and their classifications. The final results \mathbf{O} are the sum of \mathbf{O}_1 and \mathbf{O}_2 .

gorithm (see Algorithm 1). It takes a feature map $\mathbf{K} \in \mathbb{R}^{(H \times W) \times d}$ as input to compute *assignment matrix*. We partition an image feature map \mathbf{K} into regular grid cells with size $(r \times r)$ as the initialization of part-level masks, which is a common strategy in super pixel learning [1, 27, 62]. Then we average the features inside regular grid cells to get the features of part-level masks (or called cluster center) features $\mathbf{Q} \in \mathbb{R}^{N_p \times d}$, where $N_p = H/r \times W/r$. Then we compute the cosine similarities between pixel-center pairs and get $\mathbf{D} \in \mathbb{R}^{N_p \times (HW)}$. For efficiency, we do not compute the similarities between all pixel-center pairs. Instead, we only compute the similarities between pixels and their 9 nearby centers (see Fig. 3). As a result, we get $\mathbf{D}' \in \mathbb{R}^{9 \times (HW)}$. But for the convenience of describing, we still use the \mathbf{D} in the following. Due to the local constraint, each cluster center can only aggregate the nearby pixels, so we can get the *part-level masks*.

The similarities between the i -th center feature and j -th pixel feature are written as:

$$\mathbf{D}_{i,j} = \begin{cases} f(\mathbf{Q}_i, \mathbf{K}_j) & \text{if } i \in N_j \\ -\infty & \text{if } i \notin N_j, \end{cases} \quad (2)$$

where $\mathbf{Q}_i \in \mathbb{R}^d$ is the i -th cluster center feature, and $\mathbf{K}_j \in \mathbb{R}^d$ is the j -th pixel feature. $f(\mathbf{x}, \mathbf{y}) = \frac{1}{\tau} \frac{\mathbf{x} \cdot \mathbf{y}}{|\mathbf{x}| \cdot |\mathbf{y}|}$ computes the cosine similarity between \mathbf{x} and \mathbf{y} , where τ is the temperature to adjust the scale of similarities. N_j is the set of nearby regular grid cells of j -th pixel, which can be viewed in Fig. 3. Then we can compute the soft assignment matrix as:

$$\mathbf{A}_{i,j} = \text{softmax}(\mathbf{D})(i, j) = \frac{\exp(\mathbf{D}_{i,j})}{\sum_{i=1}^{N_p} \exp(\mathbf{D}_{i,j})}, \quad (3)$$




Figure 3. Explanation of the similarities between pixel features and its nearby center features. The grouping process is to assign each pixel to one of N_p center features. However, due to the computation cost of the *global* comparisons, we only compute the similarities between pixels and their nearby center features to perform *local* comparisons. For example, we only assign each pixel in the green box to one of its 9 nearby center features.

Then we can update the cluster center features by

$$\mathbf{Q}_{new} = \mathbf{A} \times \mathbf{K}. \quad (4)$$

After we get the new center features, we can compute the new assignment matrix by using updated center features \mathbf{Q}_{new} and feature map \mathbf{K} . The process is repeated L times, as shown in Algorithm 1. To get the part-level mask tokens for classification, we use the assignment matrix to extract part-level tokens from another feature map \mathbf{V} by

$$\mathbf{Z} = \mathbf{A} \times \mathbf{V}. \quad (5)$$

To strengthen the part-level mask features, we pass \mathbf{Z} to a self-attention layer and a feed forward network (FFN) layer and get \mathbf{Z}' . Then we use a linear classifier layer and a softmax activation layer to map $\mathbf{Z}' \in \mathbb{R}^{N_p \times d}$ to part-level class predictions $\mathbf{P}_m \in \mathbb{R}^{N_p \times K}$, where K is the number of classes.

Note that we use different feature maps for part-level grouping and classification to decouple these two kinds of tasks, since the shallow layers are usually used for localization while the deep layers benefit the classification [33].

Whole-level grouping. The aim of this stage is to group part-level masks into whole-level masks. Our framework is agnostic to the detailed grouping method. But for a fair comparison, we use the transformer decoders for grouping, following [14]. Each transformer decoder layer consists of one multi-head cross-attention, one self-attention, and an FFN layer. Firstly, we perform cross-attention between N_o learnable positional embeddings $\mathbf{E} \in \mathbb{R}^{N_o \times d}$ and part-level tokens $\mathbf{Z}' \in \mathbb{R}^{N_p \times d}$ to get the output features:

$$\mathbf{E}_{out} = \text{Softmax}((\mathbf{E}\mathbf{W}_q) \times (\mathbf{Z}'\mathbf{W}_k)^T) \times (\mathbf{Z}'\mathbf{W}_v), \quad (6)$$

where $\mathbf{W}_q \in \mathbb{R}^{d \times d}$, $\mathbf{W}_k \in \mathbb{R}^{d \times d}$, $\mathbf{W}_v \in \mathbb{R}^{d \times d}$ are projection heads for queries, keys, and values, respectively. For simplicity, the multi-head mechanism is ignored in the formulation. The cross-attention layer is followed by self-attention, and \mathbf{E}_{out} and FFN layers. Similar to the process in part-level mask learning, the transformer decoder operation is also repeated L times. At the end of each transformer decoder layer, there are two MLP layers. The first MLP layer maps the out features $\mathbf{E}_{out} \in \mathbb{R}^{N \times d}$ to mask embedding $\varepsilon \in \mathbb{R}^{N \times d}$. Then the whole-level masks $\mathbf{M} \in [0, 1]^{N \times (H_o \times W_o)}$ can be computed as

$$\mathbf{M} = \sigma(\varepsilon \times \mathbf{K}_0^T), \quad (7)$$

where $\mathbf{K}_0^T \in \mathbb{R}^{(H \times W) \times d}$ is a feature map before \mathbf{K} (see Fig. 2). The second MLP layer maps the out features to class logits. Then we apply a softmax activation on the class logits to get the probability predictions $\mathbf{P}_h \in \mathbb{R}^{N \times (K+1)}$, where K is the number of classes. There is an extra dimension representing the ‘‘no object’’ category (\emptyset).

The difference between our work and the previous counterparts [14, 67] is: the *keys* and *values* in previous works are pixel features, while our *keys* and *values* are features of part-level masks. Our hierarchical grouping design reduces the computation complexity, since the number of part-level masks is much smaller than the pixels.

Multi-scale semantic segmentation. We conduct classification at two levels: part-level mask classification and whole-level mask classification. The semantic segmentation results from the part-level mask classification can be computed as

$$\mathbf{O}_1 = \mathbf{P}_m^T \times \mathbf{A} \quad (8)$$

The semantic segmentation results from the whole-level mask classification can be computed as

$$\mathbf{O}_2 = \mathbf{P}_h^T \times \mathbf{M} \quad (9)$$

The final semantic segmentation result is an ensemble of two results by a simple addition.

Loss design. The challenge in part-level mask learning is that we do not have the ground truth for the part-level partition. The partition at the part-level stages is not unique. Therefore, we design two kinds of losses. First, we directly add a *cross-entropy* loss $\mathcal{L}_{\text{part,cls}}$ on \mathbf{O}_1 . Given that the *cross-entropy* loss is also affected by classification accuracy, which does not have a strong constraint on the mask quality, we also add a *pixel-cluster contrastive loss*. The core idea of the contrastive loss is to learn more discriminative feature maps, which can be used for similarity-based part-level grouping. Given ground truth masks $\mathbf{M}_G \in \mathbb{R}^{g \times (H \times W)}$, where g is the number of masks in an image, we first average the features within each ground truth mask and get $\mathbf{T} \in \mathbb{R}^{g \times d}$. Then the contrastive loss [22, 50, 54] for each pixel in \mathbf{K} is computed as:

$$\mathcal{L}_{\text{contrast}}^i = -\log \frac{\sum_{j=1}^K \mathbf{M}_{G,j,i} \exp(f(\mathbf{K}_i, \mathbf{T}_j))}{\sum_{j=1}^K \exp(f(\mathbf{K}_i, \mathbf{T}_j))}, \quad (10)$$

where $f(\mathbf{x}, \mathbf{y})$ is a function to measure the similarity between two feature vectors. The losses for the whole-level learning mainly follow previous works [14]. There is a one-to-one matching between predictions and ground truths. For the matched prediction, a dice loss $\mathcal{L}_{\text{dice}}$, a mask loss $\mathcal{L}_{\text{mask}}$, and a mask classification loss $\mathcal{L}_{\text{mask,cls}}$ are computed.

3.2. Implementation Details

We set the weights for $\mathcal{L}_{\text{part,cls}}$, $\mathcal{L}_{\text{contrast}}$, $\mathcal{L}_{\text{dice}}$, $\mathcal{L}_{\text{mask}}$, and $\mathcal{L}_{\text{mask,cls}}$ to 2, 6, 5, 5, and 2, respectively. The temperature τ in contrastive loss is 0.1, and the down sample rate r is 4. We set the number of refining stages L in part-level and whole-level grouping to 6. We use the deformable attention Transformer (MSDeformAttn) [71] layers to strengthen the features before we project the feature maps to \mathbf{K} and \mathbf{V} . The output strides for \mathbf{K} and \mathbf{V} are 1/8. We conduct a multi-scale augmentation and then crop a $512 \times 1,024$ patch for training. During testing, we send the images with the original size to models. By default, models are trained with 20k iterations with a batch size of 16. We use the ADAMW as our optimizer with an initial learning rate of 0.0001 and 0.05 weight decay.

4. Experiments

4.1. Datasets

Real-world datasets. Cityscapes contains 5,000 urban scene images collected from 50 cities primarily in Germany. The image size of Cityscapes images is $2,048 \times 1,024$. BDD is another real-world dataset, which contains 7,000 images for training, and 1,000 images for testing. The images of BDD is mainly collected from US. The image size of BDD is $1,280 \times 720$. Mapillary [39] is a large-scale dataset, which contains 18,000 images for training, 2,000

images for validation, and 5,000 images for testing. The images of Mapillary are captured from all over the world, at various conditions regarding weather and season, which makes the dataset very diverse. ACDC [48] collects the images with a resolution of $1,920 \times 1,080$ under adverse conditions, including night, fog, rain, and snow. ACDC contains 1,600 images for training, 406 images for validation, and 2,000 images for testing.

Synthetic datasets. GTAV is a synthetic dataset collected from GTAV game, which contains 12,403, 6,382, and 6,181 images with a resolution of $1,914 \times 1,052$ for training, validation, and testing, respectively. SYNTHIA [47] is another synthetic dataset, which consists of 9,400 photo-realistic images with a size of $1,280 \times 760$.

Common corruption dataset. We follow the previous works [28, 60] to expand the Cityscapes validation set with 16 type of generated corruptions. The corruptions can be divided into 4 categories: noise, blur, weather, and digital. There are 5 severity levels for each kind of corruption.

4.2. Main Results

We evaluate models on 4 kinds of generalization settings.

Normal-to-adverse generalization. In this experimental setting, all the models are trained on Cityscapes [17] (images at normal conditions) and tested on ACDC [48] (images at 4 kinds of adverse conditions). Although not specifically designed for domain generalization, transformer-based methods have been shown to be more robust than traditional CNN methods. Therefore, we compare our proposed HGFormer with two representative transformer-based segmentation methods in Tab. 1. Among them, all CNN-based methods and Segformer [60] are based on *per-pixel* classification. Mask2former [14] is based on *whole-level* classification. We can see that our method outperforms the previous CNN-based methods by a large margin, and also significantly outperforms the competitive transformer-based segmentation models.

Cityscapes-to-other datasets generalization. In this experimental setting, models are trained on Cityscapes [17] and tested on BDD [65], Mapillary [39], GTAV [46], and Synthia [47]. The results are shown in Tab. 2. In the first block of Tab. 2, we compare all the methods with a ResNet-50 [23] backbone. We can see that the grouping-based method Mask2Former [14] is already comparable to the previous domain generalization methods, which indicates the effectiveness of grouping-based model for generalization. Our HGFormer outperforms Mask2Former [14] by 1.5 points, showing that our *hierarchical* grouping-based model is better than the *flat* grouping-based model for domain generalized semantic segmentation.

Mapillary-to-other datasets generalization. Here, models are trained on Mapillary [39] and tested on BDD [65], Mapillary [39], GTAV [46], and Synthia [47]. We can

Table 1. **Cityscapes-to-ACDC generalization.** The models are trained on Cityscapes [17] only, and tested on ACDC [48]. The results of Mask2former [14], Segformer [60], and HGFormer are implemented by us. Others are from ACDC paper [48]. The results of models trained by us are an average of 3 times. The results of Segformer [60] are obtained by their officially released model.

Method	backbone	Fog	Night	Rain	Snow	All
RefineNet [32]	R101	46.4	29	52.6	43.3	43.7
DeepLabv2 [10]	R101	33.5	30.1	44.5	40.2	38
DeepLabv3+ [12]	R101	45.7	25	50	42	41.6
DANet [19]	DA101	34.7	19.1	41.5	33.3	33.1
HRNet [55]	HR-w48	38.4	20.6	44.8	35.1	35.3
Mask2former [14]	R50	54.1	36.5	53.1	50.6	49.8
HGFormer (ours)	R50	56.5	35.8	57.7	56.2	53.0
Mask2former [14]	Swin-T	56.4	39.1	58.9	58.2	54.6
Segformer [60]	B2	59.2	38.9	62.5	58.2	56.2
HGFormer (ours)	Swin-T	58.5	43.3	62.0	58.3	56.7
Segformer [60]	B5	63.2	47.8	66.4	63.7	62.0
Mask2former [14]	Swin-L	69.1	53.1	68.3	65.2	65.0
HGFormer (ours)	Swin-L	69.9	52.7	72.0	68.6	67.2

Table 2. **Cityscapes to other datasets generalization.** Models are trained on Cityscapes and tested on BDD (B), Mapillary (M), GTAV (G), and Synthia (S). Results of IBN, SW, DRPC, GTR, ISW, and SAN-SAW are from paper [43]. Others are implemented by us. Our results are an average of 3 times.

Method	backbone	B	M	G	S	Average
IBN [40]	R50	48.6	57.0	45.1	26.1	44.2
SW [41]	R50	48.5	55.8	44.9	26.1	43.8
DRPC [68]	R50	49.9	56.3	45.6	26.6	44.6
GTR [44]	R50	50.8	57.2	45.8	26.5	45.0
ISW [16]	R50	50.7	58.6	45	26.2	45.1
SAN-SAW [43]	R50	53.0	59.8	47.3	28.3	47.1
Mask2former [14]	R50	46.8	61.6	48.0	31.2	46.9
HGFormer (ours)	R50	51.5	61.6	50.4	30.1	48.4
Mask2former [14]	Swin-T	51.3	65.3	50.6	34	50.3
HGFormer (ours)	Swin-T	53.4	66.9	51.3	33.6	51.3
Mask2former [14]	Swin-L	60.1	72.2	57.8	42.4	58.1
HGFormer (ours)	Swin-L	61.5	72.1	59.4	41.3	58.6

Table 3. **Mapillary-to-other datasets generalization.** The models are trained on Mapillary, and tested on GTAV (G), Synthia (S), Cityscapes (C), and BDD (B).

Method	backbone	G	S	C	B	Average
IBN [40]	R50	30.7	27.0	42.8	31.0	32.9
SW [41]	R50	28.5	27.4	40.7	30.5	31.8
DRPC [68]	R50	33.0	29.6	46.2	32.9	35.4
GTR [44]	R50	32.9	30.3	45.8	32.6	35.4
ISW [16]	R50	33.4	30.2	46.4	32.6	35.6
SAN-SAW [43]	R50	34.0	31.6	48.7	34.6	37.2
Mask2former [14]	R50	55.8	37.7	65.6	56.4	53.9
HGFormer (ours)	R50	59.2	37.4	67.1	59.1	55.7
Mask2former [14]	Swin-T	57.8	40.1	68.2	59.1	56.3
HGFormer (ours)	Swin-T	60.1	39.5	69.3	61.0	57.5
Mask2former [14]	Swin-L	64.8	48.4	77.9	64.7	63.9
HGFormer (ours)	Swin-L	66.5	47.7	78.2	66.3	64.7

see that HGFormer is consistently better than Mask2former with all backbones, as shown in Tab. 3.

Normal-to-corruption generalization. In this set-

Table 4. Cityscapes-to-Cityscapes-C generalization (level 5).

Method	Average	Blur				Noise				Digital				Weather			
		Motion	Defoc	Glass	Gauss	Gauss	Impul	Shot	Speck	Bright	Contr	Satur	JPEG	Snow	Spatt	Fog	Frost
Mask2former-Swin-T [14]	41.6	51.5	49.4	38.2	46.2	9.6	9.8	13.5	44.4	74.2	60.0	70.0	23.3	23.7	59.4	65.4	27.3
HGFormer-Swin-T (ours)	43.9	52.9	53.9	39.0	49.5	12.1	12.3	18.2	46.3	75.0	60.0	71.2	27.2	29.4	60.6	65.0	29.1
Mask2former-Swin-L [14]	58.7	63.5	66.6	62.1	62.3	26.2	35.9	33.2	62.9	80.0	72.6	77.3	52.5	50.5	75.3	75.1	43.0
HGFormer-Swin-L (ours)	59.4	64.1	67.2	61.5	63.6	27.2	35.7	32.9	63.1	79.9	72.9	78.0	53.6	55.4	75.8	75.5	43.2

Table 5. Ablation of iterations in part-level mask classification. In this ablation, HGFormer with Swin-T [35] is trained on Cityscapes (C), and tested on Cityscapes, ACDC all (A), GTAV (G), BDD (B), Synthia (S), and Mapillary (M).

Iter	C	A	G	B	S	M	Avg
1	76.8	56.1	51.3	52.1	32.1	65.8	55.7
2	77.6	56.1	51.4	52.0	32.3	65.9	55.9
3	77.9	56.2	51.8	52.6	32.8	66.2	56.2
4	77.9	56.5	52.0	52.6	32.6	66.3	56.3
5	77.8	56.4	51.7	52.6	32.5	66.3	56.2
6	77.4	55.4	50.5	52.2	32.3	65.6	55.6




Figure 4. Visualization of results on adverse conditions. The models are only trained on Cityscapes, and tested on images with adverse conditions. Our method is significantly better than the Mask2former [14] under adverse conditions.

ting, models are trained on Cityscapes [17] and tested on Cityscapes-C [28] level 5, which includes 16 types of artificial corruptions at an extreme level. We compare HGFormer with Mask2former in Tab. 4, showing that HGFormer is significantly better than Mask2former when generalizing to extremely corrupted images.

4.3. Ablation Studies

Ablation of iterations in part-level mask classification. We test the results of different iterations of HGFormer and show the results in Tab. 5. It shows that the first iteration is much lower than later iterations on in-domain performance, and slightly lower than the later stages on out-of-distribution performance. As the iteration increases, the performance gradually increases. The performance is saturated at iteration 4. The last stage is lower than the second-last stage. We hypothesize that the last stage is influenced

by the gradients from whole-level grouping since only the part-level tokens of the last stage are taken as the input of whole-level grouping. When we remove the whole-level grouping during training, the last stage is slightly higher than the second-last stage, which verifies our hypothesis.

Individual performance of part-level and whole-level masks. We report the individual generalization performance of part-level and whole-level mask classification in Tab. 6. It shows that the part-level classification is significantly better than the whole-level classification in HGFormer. And the ensemble of part-level and whole-level classification can further improve the performance, which indicates that the part-level and whole-level mask classification are complementary to each other.

4.4. Visualization Analyses

Visualization comparisons with Mask2former. We present the visualization results of Mask2former and HGFormer, both with Swin-Tiny on ACDC (see Fig. 4) and Cityscapes-C (see Fig. 5) to demonstrate the performance of models for real-world adverse conditions and for synthetic corruptions. We choose impulse noise and defocus blur at level 5 for visualization. The results on both of the datasets show that HGFormer makes fewer errors than Mask2former in adverse conditions.

Masks at different levels of corruption. We visualize the part-level and whole-level masks at different levels of corruption to show how they change as the severity level increases (see Fig. 6). We can see that the whole-level masks are not stable with the increasing of severity levels, and totally failed at level 5. In contrast, the part-level masks are more stable, and can achieve high recall for the boundaries between classes.

Part-level masks with different model weights. To provide more insights about our method, we visualize the part-level masks with model weights from random initialization, ImageNet pre-trained weights and Cityscapes trained weights. The results are shown in Fig. 7. We can see that even using the randomly initialized weights and ImageNet pre-trained weights, our model can produce reasonable part-level masks, which indicates that our model has the potential for unsupervised segmentation and weakly-supervised segmentation. The results indicate that the *part-level grouping structure itself can provide a good prior*, which can explain the generalization from the grouping side. For the Cityscapes trained model weights, the *boundaries between*

Table 6. **Comparison of part-level classification and whole-level classification, and their combination.** We train HGFormer with Swin-T on Cityscapes, then test the model on other datasets.

whole-level mask	part-level mask	ACDC (all)	GTAV	BDD	Synthia	Mapillary	Average
✓		54.5	49.5	51.5	33.8	66.3	51.1
	✓	56.2	51.3	53.1	33.3	66.5	52.1
✓	✓	56.6	51.3	53.4	33.6	66.9	52.4



Figure 5. **Visualization of results on corruptions.** We choose two kinds of corruption at level 5 for this visualization: impulse noise and defocus blur. The models are trained on Cityscapes.




Figure 6. **Visualization of part-level and whole-level masks at different levels of Gaussian noise.** In the first row, we visualize the whole-level masks from Mask2former. In the second row, we visualize the part-level masks from our method. We can see that our part-level masks are more robust than the whole-level masks in Mask2Former as the increasing severity level of Gaussian noise.




Figure 7. **Visualization of part-level masks with different weights.** We find that the randomly initialized weights can also produce some reasonable part-level masks.

different categories are more accurate than the randomly initialized and ImageNet pre-trained weights. It is worthwhile noting that the *boundaries between the same class* is not unique, due to no ground truths being used for part-level masks. But with a part-level classification, all feasible part-level partitioning can be transformed to correct semantic segmentation results, if the *boundaries between different classes* are correct.

5. Conclusion and Future work

In this paper, we propose a hierarchical semantic segmentation model, which can efficiently generate image partitions in a hierarchical structure. Then we perform both part-level mask classification and whole-level mask classification.

The final semantic segmentation result is an ensemble of two results. Our method is verified to be robust in out-of-distribution images. We can explore more complicated fusion methods of classification at different scales. We leave them as future works. Since our model can be considered a kind of multi-task learning, how to automatically balance the loss weights of classification at different scales can be studied in the future.

Acknowledgement. This work was supported by the National Nature Science Foundation of China under grants U22B2011 and 62101390. Jian Ding was also supported by the China Scholarship Council. We would like to thank Ahmed Abbas for the insightful discussions.

References

- [1] Radhakrishna Achanta, Appu Shaji, Kevin Smith, Aurelien Lucchi, Pascal Fua, and Sabine Süsstrunk. Slic superpixels compared to state-of-the-art superpixel methods. *IEEE PAMI*, 34(11):2274–2282, 2012. 4
- [2] Pablo Arbelaez, Michael Maire, Charless Fowlkes, and Jitendra Malik. Contour detection and hierarchical image segmentation. *IEEE TPAMI*, 33(5):898–916, 2010. 2
- [3] Yogesh Balaji, Swami Sankaranarayanan, and Rama Chellappa. Metareg: Towards domain generalization using meta-regularization. *NeurIPS*, 31, 2018. 3
- [4] Silvia Bucci, Antonio D’Innocente, Yujun Liao, Fabio M Carlucci, Barbara Caputo, and Tatiana Tommasi. Self-supervised learning across domains. *IEEE PAMI*, 44(9):5516–5528, 2021. 3
- [5] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, pages 213–229. Springer, 2020. 2
- [6] Fabio M Carlucci, Antonio D’Innocente, Silvia Bucci, Barbara Caputo, and Tatiana Tommasi. Domain generalization by solving jigsaw puzzles. In *CVPR*, pages 2229–2238, 2019. 3
- [7] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, pages 9650–9660, 2021. 2
- [8] Junbum Cha, Hanchool Cho, Kyungjae Lee, Seunghyun Park, Yunsung Lee, and Sungrae Park. Domain generalization needs stochastic weight averaging for robustness on domain shifts. *arXiv:2102.08604*, 3, 2021. 3
- [9] 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 PAMI*, 40(4):834–848, 2017. 2
- [10] 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 PAMI*, 40(4):834–848, 2017. 2, 3, 6
- [11] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv:1706.05587*, 2017. 2
- [12] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *ECCV*, pages 801–818, 2018. 2, 6
- [13] Yuhua Chen, Dengxin Dai, Jordi Pont-Tuset, and Luc Van Gool. Scale-aware alignment of hierarchical image segmentation. In *CVPR*, 2016. 2
- [14] Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *CVPR*, pages 1290–1299, 2022. 2, 3, 5, 6, 7
- [15] Bowen Cheng, Alex Schwing, and Alexander Kirillov. Per-pixel classification is not all you need for semantic segmentation. *NeurIPS*, 34:17864–17875, 2021. 2, 3
- [16] Sungha Choi, Sanghun Jung, Huiwon Yun, Joanne T Kim, Seungryong Kim, and Jaegul Choo. Robustnet: Improving domain generalization in urban-scene segmentation via instance selective whitening. In *CVPR*, pages 11580–11590, 2021. 3, 6
- [17] 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 *CVPR*, pages 3213–3223, 2016. 6, 7
- [18] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv:2010.11929*, 2020. 3
- [19] Jun Fu, Jing Liu, Haijie Tian, Yong Li, Yongjun Bao, Zhiwei Fang, and Hanqing Lu. Dual attention network for scene segmentation. In *CVPR*, pages 3146–3154, 2019. 6
- [20] Raghudeep Gadde, Varun Jampani, Martin Kiefel, Daniel Kappler, and Peter V Gehler. Superpixel convolutional networks using bilateral inceptions. In *ECCV*, pages 597–613. Springer, 2016. 2
- [21] Yong Guo, David Stutz, and Schiele Bernt. Improving robustness of vision transformers by reducing sensitivity to patch corruptions. In *CVPR*, 2023. 1
- [22] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, pages 9729–9738, 2020. 5
- [23] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 6
- [24] Geoffrey Hinton. How to represent part-whole hierarchies in a neural network. *arXiv:2102.12627*, 2021. 2
- [25] Lukas Hoyer, Dengxin Dai, and Luc Van Gool. DAFormer: Improving network architectures and training strategies for domain-adaptive semantic segmentation. In *CVPR*, 2022. 1
- [26] Zeyi Huang, Haohan Wang, Eric P Xing, and Dong Huang. Self-challenging improves cross-domain generalization. In *ECCV*, pages 124–140. Springer, 2020. 3
- [27] Varun Jampani, Deqing Sun, Ming-Yu Liu, Ming-Hsuan Yang, and Jan Kautz. Superpixel sampling networks. In *ECCV*, pages 352–368, 2018. 2, 4
- [28] Christoph Kamann and Carsten Rother. Benchmarking the robustness of semantic segmentation models. In *CVPR*, pages 8828–8838, 2020. 6, 7
- [29] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. *NeurIPS*, 30, 2017. 3
- [30] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy Hospedales. Learning to generalize: Meta-learning for domain generalization. In *AAAI*, volume 32, 2018. 3
- [31] Haoliang Li, Sinno Jialin Pan, Shiqi Wang, and Alex C Kot. Domain generalization with adversarial feature learning. In *CVPR*, pages 5400–5409, 2018. 3

- [32] Guosheng Lin, Anton Milan, Chunhua Shen, and Ian Reid. Refinenet: Multi-path refinement networks for high-resolution semantic segmentation. In *CVPR*, pages 1925–1934, 2017. 6
- [33] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, pages 2117–2125, 2017. 5
- [34] Quande Liu, Qi Dou, Lequan Yu, and Pheng Ann Heng. Ms-net: multi-site network for improving prostate segmentation with heterogeneous mri data. *T-MI*, 39(9):2713–2724, 2020. 3
- [35] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, pages 10012–10022, 2021. 2, 7
- [36] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, pages 3431–3440, 2015. 2, 3
- [37] Krikamol Muandet, David Balduzzi, and Bernhard Schölkopf. Domain generalization via invariant feature representation. In *International Conference on Machine Learning*, pages 10–18. PMLR, 2013. 3
- [38] Muhammad Muzammal Naseer, Kanchana Ranasinghe, Salman H Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Intriguing properties of vision transformers. *NeurIPS*, 34:23296–23308, 2021. 2
- [39] Gerhard Neuhof, Tobias Ollmann, Samuel Rota Buló, and Peter Kotschieder. The mapillary vistas dataset for semantic understanding of street scenes. In *ICCV*, pages 4990–4999, 2017. 5, 6
- [40] Xingang Pan, Ping Luo, Jianping Shi, and Xiaoou Tang. Two at once: Enhancing learning and generalization capacities via ibn-net. In *ECCV*, pages 464–479, 2018. 3, 6
- [41] Xingang Pan, Xiaohang Zhan, Jianping Shi, Xiaoou Tang, and Ping Luo. Switchable whitening for deep representation learning. In *ICCV*, pages 1863–1871, 2019. 3, 6
- [42] Sayak Paul and Pin-Yu Chen. Vision transformers are robust learners. In *AAAI*, volume 36, pages 2071–2081, 2022. 1
- [43] Duo Peng, Yinjie Lei, Munawar Hayat, Yulan Guo, and Wen Li. Semantic-aware domain generalized segmentation. In *CVPR*, pages 2594–2605, 2022. 3, 6
- [44] Duo Peng, Yinjie Lei, Lingqiao Liu, Pingping Zhang, and Jun Liu. Global and local texture randomization for synthetic-to-real semantic segmentation. *IEEE TIP*, 30:6594–6608, 2021. 3, 6
- [45] Alexandre Rame, Matthieu Kirchmeyer, Thibaud Rahier, Alain Rakotomamonjy, Patrick Gallinari, and Matthieu Cord. Diverse weight averaging for out-of-distribution generalization. *arXiv:2205.09739*, 2022. 3
- [46] Stephan R Richter, Vibhav Vineet, Stefan Roth, and Vladlen Koltun. Playing for data: Ground truth from computer games. In *ECCV*, pages 102–118. Springer, 2016. 6
- [47] German Ros, Laura Sellart, Joanna Materzynska, David Vazquez, and Antonio M Lopez. The synthia dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In *CVPR*, pages 3234–3243, 2016. 6
- [48] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Acdc: The adverse conditions dataset with correspondences for semantic driving scene understanding. In *ICCV*, pages 10765–10775, 2021. 6
- [49] Zhuowen Tu, Xiangrong Chen, Alan L Yuille, and Song-Chun Zhu. Image parsing: Unifying segmentation, detection, and recognition. *IJCV*, 63(2):113–140, 2005. 2
- [50] Wouter Van Gansbeke, Simon Vandenhende, Stamatios Georgoulis, and Luc Van Gool. Unsupervised semantic segmentation by contrasting object mask proposals. In *ICCV*, pages 10052–10062, 2021. 5
- [51] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *NeurIPS*, 30, 2017. 2
- [52] Riccardo Volpi and Vittorio Murino. Addressing model vulnerability to distributional shifts over image transformation sets. In *ICCV*, pages 7980–7989, 2019. 3
- [53] Haohan Wang, Zexue He, Zachary C Lipton, and Eric P Xing. Learning robust representations by projecting superficial statistics out. *arXiv:1903.06256*, 2019. 3
- [54] Huiyu Wang, Yukun Zhu, Hartwig Adam, Alan Yuille, and Liang-Chieh Chen. Max-deeplab: End-to-end panoptic segmentation with mask transformers. In *CVPR*, pages 5463–5474, 2021. 2, 5
- [55] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Mingkui Tan, Xingang Wang, et al. Deep high-resolution representation learning for visual recognition. *IEEE PAMI*, 43(10):3349–3364, 2020. 6
- [56] Mei Wang and Weihong Deng. Deep visual domain adaptation: A survey. *Neurocomputing*, 312:135–153, 2018. 1
- [57] Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In *ICCV*, pages 568–578, 2021. 2
- [58] Florian Wenzel, Andrea Dittadi, Peter Vincent Gehler, Carl-Johann Simon-Gabriel, Max Horn, Dominik Zietlow, David Kernert, Chris Russell, Thomas Brox, Bernt Schiele, et al. Assaying out-of-distribution generalization in transfer learning. *arXiv:2207.09239*, 2022. 1
- [59] Mitchell Wortsman, Gabriel Ilharco, Samir Ya Gadre, Rebecca Roelofs, Raphael Gontijo-Lopes, Ari S Morcos, Hongseok Namkoong, Ali Farhadi, Yair Carmon, Simon Kornblith, et al. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In *International Conference on Machine Learning*, pages 23965–23998. PMLR, 2022. 3
- [60] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. *NeurIPS*, 34:12077–12090, 2021. 1, 2, 6
- [61] Zhenlin Xu, Deyi Liu, Junlin Yang, Colin Raffel, and Marc Niethammer. Robust and generalizable visual representation learning via random convolutions. *arXiv:2007.13003*, 2020. 3

- [62] Fengting Yang, Qian Sun, Hailin Jin, and Zihan Zhou. Superpixel segmentation with fully convolutional networks. In *CVPR*, pages 13964–13973, 2020. [2](#), [4](#)
- [63] Teresa Yeo, Oğuzhan Fatih Kar, and Amir Zamir. Robustness via cross-domain ensembles. In *ICCV*, pages 12189–12199, 2021. [3](#)
- [64] Teresa Yeo, Oğuzhan Fatih Kar, and Amir Zamir. Robustness via cross-domain ensembles. In *ICCV*, pages 12189–12199, 2021. [3](#)
- [65] Fisher Yu, Haofeng Chen, Xin Wang, Wenqi Xian, Yingying Chen, Fangchen Liu, Vashisht Madhavan, and Trevor Darrell. Bdd100k: A diverse driving dataset for heterogeneous multitask learning. In *CVPR*, pages 2636–2645, 2020. [6](#)
- [66] Qihang Yu, Huiyu Wang, Dahun Kim, Siyuan Qiao, Maxwell Collins, Yukun Zhu, Hartwig Adam, Alan Yuille, and Liang-Chieh Chen. Cmt-deeplab: Clustering mask transformers for panoptic segmentation. In *CVPR*, pages 2560–2570, 2022. [2](#)
- [67] Qihang Yu, Huiyu Wang, Siyuan Qiao, Maxwell Collins, Yukun Zhu, Hartwig Adam, Alan Yuille, and Liang-Chieh Chen. k-means mask transformer. In *ECCV*, pages 288–307. Springer, 2022. [2](#), [5](#)
- [68] Xiangyu Yue, Yang Zhang, Sicheng Zhao, Alberto Sangiovanni-Vincentelli, Kurt Keutzer, and Boqing Gong. Domain randomization and pyramid consistency: Simulation-to-real generalization without accessing target domain data. In *ICCV*, pages 2100–2110, 2019. [3](#), [6](#)
- [69] Yifan Zhang, Bo Pang, and Cewu Lu. Semantic segmentation by early region proxy. In *CVPR*, pages 1258–1268, 2022. [2](#)
- [70] Daquan Zhou, Zhiding Yu, Enze Xie, Chaowei Xiao, Animesh Anandkumar, Jiashi Feng, and Jose M Alvarez. Understanding the robustness in vision transformers. In *International Conference on Machine Learning*, pages 27378–27394. PMLR, 2022. [1](#), [2](#), [3](#)
- [71] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv:2010.04159*, 2020. [5](#)