

# Rethinking Semantic Segmentation: A Prototype View

Tianfei Zhou<sup>1</sup>, Wenguan Wang<sup>2,1\*</sup>, Ender Konukoglu<sup>1</sup>, Luc Van Gool<sup>1</sup>

<sup>1</sup> Computer Vision Lab, ETH Zurich <sup>2</sup> ReLER, AAIL, University of Technology Sydney

<https://github.com/tfzhou/ProtoSeg>

## Abstract

Prevalent semantic segmentation solutions, despite their different network designs (FCN based or attention based) and mask decoding strategies (parametric softmax based or pixel-query based), can be placed in one category, by considering the softmax weights or query vectors as learnable class prototypes. In light of this prototype view, this study uncovers several limitations of such parametric segmentation regime, and proposes a nonparametric alternative based on non-learnable prototypes. Instead of prior methods learning a single weight/query vector for each class in a fully parametric manner, our model represents each class as a set of non-learnable prototypes, relying solely on the mean features of several training pixels within that class. The dense prediction is thus achieved by nonparametric nearest prototype retrieving. This allows our model to directly shape the pixel embedding space, by optimizing the arrangement between embedded pixels and anchored prototypes. It is able to handle arbitrary number of classes with a constant amount of learnable parameters. We empirically show that, with FCN based and attention based segmentation models (i.e., HR-Net, Swin, SegFormer) and backbones (i.e., ResNet, HRNet, Swin, MiT), our nonparametric framework yields compelling results over several datasets (i.e., ADE20K, Cityscapes, COCO-Stuff), and performs well in the large-vocabulary situation. We expect this work will provoke a rethink of the current *de facto* semantic segmentation model design.

## 1. Introduction

With the renaissance of connectionism, rapid progress has been made in semantic segmentation. Till now, most of state-of-the-art segmentation models [15, 34, 49, 135] were built upon Fully Convolutional Networks (FCNs) [79]. Despite their diversified model designs and impressive results, existing *FCN based* methods commonly apply **parametric softmax** (●●●) over pixel-wise features for dense prediction (Fig. 1(a)). Very recently, the vast success of Transformer [105] stimulates the emergence of *attention based*

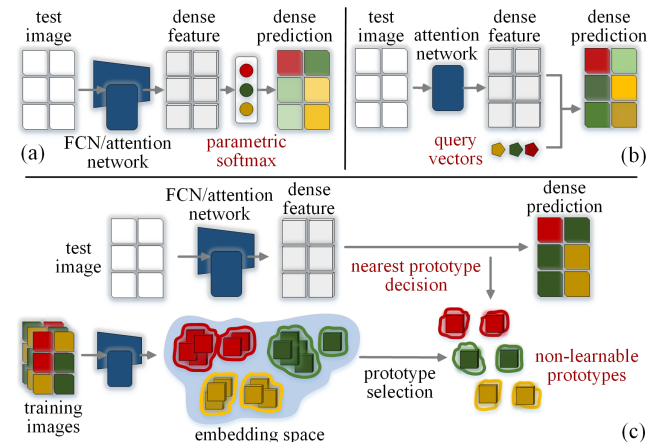


Figure 1. Different semantic segmentation paradigms: (a-b) **parametric** vs (c) **nonparametric**. Modern segmentation solutions, no matter using (a) parametric softmax or (b) query vectors for mask decoding, can be viewed as learnable prototype based methods that learn class-wise prototypes in a fully parametric manner. We instead propose a nonparametric scheme (c) that directly selects sub-cluster centers of embedded pixels as prototypes, and achieves per-pixel prediction via nonparametric nearest prototype retrieving.

segmentation solutions. Many of these ‘non-FCN’ models, like [118, 139], directly follow the standard mask decoding regime, *i.e.*, estimate softmax distributions over dense visual embeddings (extracted from patch token sequences). Interestingly, the others [20, 100] follow the good practice of Transformer in other fields [11, 82, 113] and adopt a **pixel-query** strategy (Fig. 1(b)): utilize a set of learnable vectors (●●●) to query the dense embeddings for mask prediction. They speculate the learned query vectors can capture class-wise properties, however, lacking in-depth analysis.

Noticing there exist two different mask decoding strategies, the following questions naturally arise: ① *What are the relation and difference between them?* and ② *If the learnable query vectors indeed implicitly capture some intrinsic properties of data, is there any better way to achieve this?*

Tackling these two issues can provide insights into modern segmentation model design, and motivate us to rethink the task from a *prototype view*. The idea of prototype based classification [31] is classical and intuitive (which can date back to the nearest neighbors algorithm [23] and find evi-

\*Corresponding author: Wenguan Wang.

dence in cognitive science [60, 91]): data samples are classified based on their proximity to *representative* prototypes of classes. With this perspective, in §2, we first answer question ❶ by pointing out most modern segmentation methods, from softmax based to pixel-query based, from FCN based to attention based, fall into one grand category: *parametric* models based on *learnable prototypes*. Consider a segmentation task with  $C$  semantic classes. Most existing efforts seek to directly *learn*  $C$  class-wise prototypes – softmax weights or query vectors – for parametric, pixel-wise classification. Hence question ❷ becomes more fundamental: ❸ *What are the limitations of this learnable prototype based parametric paradigm?* and ❹ *How to address these limitations?*

Driven by question ❸, we find there are three critical limitations: *First*, usually only one single prototype is learned per class, insufficient to describe rich intra-class variance. The prototypes are simply learned in a fully parametric manner, without considering their representative ability. *Second*, to map a  $H \times W \times D$  image feature tensor into a  $H \times W \times C$  semantic mask, at least  $D \times C$  parameters are needed for prototype learning. This hurts generalizability [115], especially in the large-vocabulary case; for instance, if there are 800 classes and  $D = 512$ , we need 0.4M learnable prototype parameters *alone*. *Third*, with the cross-entropy loss, only the relative relations between intra-class and inter-class distances are optimized [89, 111, 134]; the actual distances between pixels and prototypes, *i.e.*, intra-class compactness, are ignored.

As a response to question ❹, in §3, we develop a *nonparametric* segmentation framework, based on *non-learnable prototypes*. Specifically, building upon the ideas of prototype learning [116, 133] and metric learning [40, 64], it is fully aware of the limitations of its parametric counterpart. Independent of specific backbone architectures (FCN based or attention based), our method is general and brings insights into segmentation model design and training. For model design, our method explicitly sets *sub-class centers*, in the pixel embedding space, as the prototypes. Each pixel data is predicted to be in the same class as the nearest prototype, without relying on extra learnable parameters. For training, as the prototypes are representative of the dataset, we can directly pose known inductive biases (*e.g.*, intra-class compactness, inter-class separation) as extra optimization criteria and efficiently shape the whole embedding space, instead of optimizing the prediction accuracy only. Our model has three appealing advantages: *First*, each class is abstracted by a set of prototypes, well capturing class-wise characteristics and intra-class variance. With the clear meaning of the prototypes, the interpretability is also enhanced – the prediction of each pixel can be intuitively understood as the reference of its closest class center in the embedding space [3, 7]. *Second*, due to the nonparametric nature, the generalizability is improved. Large-vocabulary semantic segmentation can also be handled efficiently, as the amount of learnable

prototype parameters is no longer constrained to the number of classes (*i.e.*, 0 vs  $D \times C$ ). *Third*, via prototype-anchored metric learning, the pixel embedding space is shaped as well-structured, benefiting segmentation prediction eventually.

By answering questions ❶-❹, we formalize prior methods within a learnable prototype based, parametric framework, and link this field to prototype learning and metric learning. We provide literature review and related discussions in §4.

In §5.2, we show our method achieves impressive results over famous datasets (*i.e.*, ADE20K [140], Cityscapes [22], COCO-Stuff [10]) with top-leading FCN based and attention based segmentation models (*i.e.*, HRNet [108], Swin [78], SegFormer [118]) and backbones (*i.e.*, ResNet [45], HRNet [108], Swin [78], MiT [118]). Compared with the parametric counterparts, our method does not cause any extra computational overhead during testing while reduces the amount of learnable parameters. In §5.3, we demonstrate our method consistently performs well when increasing the number of semantic classes from 150 to 847. Accompanied with a set of ablative studies in §5.4, our extensive experiments verify the power of our idea and the efficacy of our algorithm.

Finally, we draw conclusions in §6. This work is expected to open a new venue for future research in this field.

## 2. Existing Semantic Segmentation Models as Parametric Prototype Learning

Next we first formalize the existing two mask decoding strategies mentioned in §1, and then answer question ❶ from a unified view of parametric prototype learning.

**Parametric Softmax Projection.** Almost all FCN-like and many attention-based segmentation models adopt this strategy. Their models comprise two learnable parts: i) an encoder  $\phi$  for dense visual feature extraction, and ii) a classifier  $\rho$  (*i.e.*, projection head) that projects pixel features into the semantic label space. For each pixel example  $i$ , its embedding  $\mathbf{i} \in \mathbb{R}^D$ , extracted from  $\phi$ , is fed into  $\rho$  for  $C$ -way classification:

$$p(c|i) = \frac{\exp(\mathbf{w}_c^\top \mathbf{i})}{\sum_{c'=1}^C \exp(\mathbf{w}_{c'}^\top \mathbf{i})}, \quad (1)$$

where  $p(c|i) \in [0, 1]$  is the probability that  $i$  being assigned to class  $c$ .  $\rho$  is a pixel-wise linear layer, parameterized by  $\mathbf{W} = [\mathbf{w}_1, \dots, \mathbf{w}_C] \in \mathbb{R}^{C \times D}$ ;  $\mathbf{w}_c \in \mathbb{R}^D$  is a learnable projection vector for  $c$ -th class; the bias term is omitted for brevity.

**Parametric Pixel-Query.** A few attention-based segmentation networks [118, 139] work in a more ‘Transformer-like’ manner: given the pixel embedding  $\mathbf{i} \in \mathbb{R}^D$ , a set of  $C$  query vectors, *i.e.*,  $\mathbf{E} = [\mathbf{e}_1, \dots, \mathbf{e}_C] \in \mathbb{R}^{C \times D}$ , are learned to generate a probability distribution over the  $C$  classes:

$$p(c|i) = \frac{\exp(\mathbf{e}_c * \mathbf{i})}{\sum_{c'=1}^C \exp(\mathbf{e}_{c'} * \mathbf{i})}, \quad (2)$$

where ‘\*’ is inner product between  $\ell_2$ -normalized inputs.




Figure 2. **Architecture illustration** of our non-learnable prototype based nonparametric segmentation model during the training phase.

**Prototype-based Classification.** Prototype-based classification [31, 33] has been studied for a long time, dating back to the nearest neighbors algorithm [23] in machine learning and *prototype theory* [60, 91] in cognitive science. Its prevalence stems from its intuitive idea: represent classes by prototypes, and refer to prototypes for classification. Let  $\{p_m\}_{m=1}^M$  be a set of prototypes that are representative of their corresponding classes  $\{c_{p_m} \in \{1, \dots, C\}\}_m$ . For a data sample  $i$ , prediction is made by comparing  $i$  with  $\{p_m\}_m$ , and taking the class of the *winning* prototype as response:

$$\hat{c}_i = c_{p_{m^*}}, \quad \text{with } m^* = \arg \min_m \{\langle i, p_m \rangle\}_{m=1}^M, \quad (3)$$

where  $i$  and  $\{p_m\}_m$  are embeddings of the data sample and prototypes in a feature space, and  $\langle \cdot, \cdot \rangle$  stands for the distance measure, which is typically set as  $\ell_2$  distance (i.e.,  $\|i - p_m\|$ ) [123], yet other proximities can be applied.

Further, Eqs. 1-2 can be formulated in a unified form:

$$p(c|i) = \frac{\exp(-\langle i, g_c \rangle)}{\sum_{c'=1}^C \exp(-\langle i, g_{c'} \rangle)}, \quad (4)$$

where  $g_c \in \mathbb{R}^D$  can be either  $w_c$  in Eq. 1 or  $e_c$  in Eq. 2.

With Eqs. 3-4, we are ready to answer questions 1-2. Both the two types of methods are based on *learnable* prototypes; they are *parametric* models in the sense that they learn one prototype  $g_c$ , i.e., linear weight  $w_c$  or query vector  $e_c$ , for each class  $c$  (i.e.,  $M=C$ ). Thus one can consider softmax projection based methods ‘secretly’ learn the query vectors. As for the difference, in addition to different distance measures (i.e., inner product vs cosine similarity), pixel-query based methods [118, 139] can feed the queries into cross-attention decoder layers for cross-class context exchanging, rather than softmax projection based counterparts only leveraging the learned class weights within the softmax layer.

With the unified view of parametric prototype learning, a few intrinsic yet long ignored issues in this field unfold:

First, *prototype selection* [36] is a vital aspect in the design of a prototype based learner – prototypes should be *typical* for their classes. Nevertheless, existing semantic segmentation algorithms often describe each class by only one prototype, bearing no intra-class variation. Moreover, the prototypes are directly learned in a fully parametric manner, without accounting for their representative ability.

Second, the amount of the learnable prototype parameters, i.e.,  $\{g_c \in \mathbb{R}^D\}_{c=1}^C$ , grows with the number of classes. This may hinder the scalability, especially when a large number of classes are present. For example, if there are 800 classes and the pixel feature dimensionality is 512, at least 0.4M parameters are needed for prototype learning *alone*, making large-vocabulary segmentation a hard task. Moreover, if we want to represent each class by ten prototypes, instead of only one, we need to learn 4M prototype parameters.

Third, Eq. 3 intuitively shows that prototype based learners make metric comparisons of data [8]. However, existing algorithms often supervise dense segmentation representation by directly optimizing the accuracy of pixel-wise prediction (e.g., cross-entropy loss), ignoring known inductive biases [83, 84], e.g., intra-class compactness, about the feature distribution. This will hinder the discrimination potential of the learned segmentation features, as suggested by many literature in representation learning [76, 95, 114].

After tackling question 3, in the next section we will detail our non-learnable prototype based nonparametric segmentation method, which serves as a solid response to question 4.

### 3. Non-Learnable Prototype based Nonparametric Semantic Segmentation

We build a nonparametric segmentation framework that conducts dense prediction by a set of non-learnable class prototypes, and directly supervises the pixel embedding space via a prototype-anchored metric learning scheme (Fig. 2).

**Non-Learnable Prototype based Pixel Classification.** As normal, an encoder network (FCN based or attention based), i.e.,  $\phi$ , is first adopted to map the input image  $I \in \mathbb{R}^{h \times w \times 3}$ , to a 3D feature tensor  $I \in \mathbb{R}^{H \times W \times D}$ . For pixel-wise  $C$ -way classification, rather than prior semantic segmentation models that automatically learn  $C$  class weights  $\{w_c \in \mathbb{R}^D\}_{c=1}^C$  (cf. Eq. 1) or  $C$  queries vectors  $\{e_c \in \mathbb{R}^D\}_{c=1}^C$  (cf. Eq. 2), we refer to a group of  $CK$  non-learnable prototypes, i.e.,  $\{p_{c,k} \in \mathbb{R}^D\}_{c,k=1}^{C,K}$ , which are based solely on class data sub-centers. More specifically, each class  $c \in \{1, \dots, C\}$  is represented by a total of  $K$  prototypes  $\{p_{c,k}\}_{k=1}^K$ , and prototype  $p_{c,k}$  is determined as the center of  $k$ -th sub-cluster of training pixel samples belonging to class  $c$ , in the embedding space  $\phi$ . In this way, the prototypes can comprehen-

sively capture characteristic properties of the corresponding classes, without introducing extra learnable parameters outside  $\phi$ . Analogous to Eq. 3, the category prediction of each pixel  $i \in I$  is achieved by a winner-take-all classification:

$$\hat{c}_i = c^*, \text{ with } (c^*, k^*) = \arg \min_{(c,k)} \{\langle \mathbf{i}, \mathbf{p}_{c,k} \rangle\}_{c,k=1}^{C,K}, \quad (5)$$

where  $\mathbf{i} \in \mathbb{R}^D$  stands for the  $\ell_2$ -normalized embedding of pixel  $i$ , i.e.,  $\mathbf{i} \in I$ , and the distance measure  $\langle \cdot, \cdot \rangle$  is defined as the negative cosine similarity, i.e.,  $\langle \mathbf{i}, \mathbf{p} \rangle = -\mathbf{i}^\top \mathbf{p}$ .

With this exemplar-based reasoning mode, we first define the probability distribution of pixel  $i$  over the  $C$  classes:

$$p(c|i) = \frac{\exp(-s_{i,c})}{\sum_{c'=1}^C \exp(-s_{i,c'})}, \text{ with } s_{i,c} = \min_{k=1}^K \{\langle \mathbf{i}, \mathbf{p}_{c,k} \rangle\}_{k=1}^K, \quad (6)$$

where the *pixel-class distance*  $s_{i,c} \in [-1, 1]$  is computed as the distance to the closest prototype of class  $c$ . Given the groundtruth class of each pixel  $i$ , i.e.,  $c_i \in \{1, \dots, C\}$ , the cross-entropy loss can be therefore used for training:

$$\begin{aligned} \mathcal{L}_{CE} &= -\log p(c_i|i) \\ &= -\log \frac{\exp(-s_{i,c_i})}{\exp(-s_{i,c_i}) + \sum_{c' \neq c_i} \exp(-s_{i,c'})}. \end{aligned} \quad (7)$$

In our case, Eq. 7 can be viewed as pushing pixel  $i$  closer to the nearest prototype of its corresponding class, i.e.,  $c_i$ , and further from other close prototypes of irrelevant classes, i.e.,  $c' \neq c_i$ . However, only adopting such training objective is not enough, due to two reasons. First, Eq. 7 only considers pixel-class distances, e.g.,  $s_{i,c}$ , without addressing within-class pixel-prototype relations, e.g.,  $\langle \mathbf{i}, \mathbf{p}_{c_i,k} \rangle$ . For example, for discriminative representation learning, pixel  $i$  is expected to be pushed further close to a certain prototype (i.e., a particularly suitable pattern) of class  $c_i$ , and, distant from other prototypes (i.e., other irrelevant but within-class patterns) of class  $c_i$ . Eq. 7 cannot capture this nature. Second, as the pixel-class distances are normalized across all classes (cf. Eq. 6), Eq. 7 only optimizes the *relative* relations between intra-class (i.e.,  $s_{i,c_i}$ ) and inter-class (i.e.,  $\{s_{i,c'}\}_{c' \neq c_i}$ ) distances, instead of directly regularizing the cosine distances between pixels and classes. For example, when the intra-class distance  $s_{i,c_i}$  of pixel  $i$  is relatively smaller than other inter-class distances  $\{s_{i,c'}\}_{c' \neq c_i}$ , the penalty from Eq. 7 will be small, but the intra-class distance  $s_{i,c_i}$  might still be large [89, 134]. Next we first elaborate on our within-class online clustering strategy and then detail our two extra training objectives which rely on prototype assignments (i.e., clustering results) and address the above two issues respectively.

**Within-Class Online Clustering.** We approach online clustering for prototype selection and assignment: pixel samples within the same class are assigned to the prototypes belonging to that class, and the prototypes are then updated according to the assignments. Clustering imposes a natural bottleneck [55] that forces the model to discover intra-class discriminative patterns yet discard instance-specific details.

Thus the prototypes, selected as the sub-cluster centers, are typical of the corresponding classes. Conducting clustering *online* makes our method scalable to large amounts of data, instead of offline clustering requiring multiple passes over the entire dataset for feature computation [13].

Formally, given pixels  $\mathcal{I}^c = \{i_n\}_{n=1}^N$  in a training batch that belong to class  $c$  (i.e.,  $c_{i_n} = c$ ), our goal is to map the pixels  $\mathcal{I}^c$  to the  $K$  prototypes  $\{\mathbf{p}_{c,k}\}_{k=1}^K$  of class  $c$ . We denote this pixel-to-prototype mapping as  $\mathbf{L}^c = [\mathbf{l}_{i_n}]_{n=1}^N \in \{0, 1\}^{K \times N}$ , where  $\mathbf{l}_{i_n} = [l_{i_n,k}]_{k=1}^K \in \{0, 1\}^K$  is the one-hot assignment vector of pixel  $i_n$  over the  $K$  prototypes. The optimization of  $\mathbf{L}^c$  is achieved by maximizing the similarity between pixel embeddings, i.e.,  $\mathbf{X}^c = [\mathbf{i}_n]_{n=1}^N \in \mathbb{R}^{D \times N}$ , and the prototypes, i.e.,  $\mathbf{P}^c = [\mathbf{p}_{c,k}]_{k=1}^K \in \mathbb{R}^{D \times K}$ :

$$\begin{aligned} &\max_{\mathbf{L}^c} \text{Tr}(\mathbf{L}^{c\top} \mathbf{P}^{c\top} \mathbf{X}^c), \\ \text{s.t. } &\mathbf{L}^c \in \{0, 1\}^{K \times N}, \mathbf{L}^{c\top} \mathbf{1}^K = \mathbf{1}^N, \mathbf{L}^c \mathbf{1}^N = \frac{N}{K} \mathbf{1}^K, \end{aligned} \quad (8)$$

where  $\mathbf{1}^K$  denotes the vector of all ones of  $K$  dimensions. The unique assignment constraint, i.e.,  $\mathbf{L}^{c\top} \mathbf{1}^K = \mathbf{1}^N$ , ensures that each pixel is assigned to one and only one prototype. The equipartition constraint, i.e.,  $\mathbf{L}^c \mathbf{1}^N = \frac{N}{K} \mathbf{1}^K$ , enforces that on average each prototype is selected at least  $\frac{N}{K}$  times in the batch [13]. This prevents the trivial solution: all pixel samples are assigned to a single prototype, and eventually benefits the representative ability of the prototypes. To solve Eq. 8, one can relax  $\mathbf{L}^c$  to be an element of the *transportation polytope* [2, 24]:

$$\begin{aligned} &\max_{\mathbf{L}^c} \text{Tr}(\mathbf{L}^{c\top} \mathbf{P}^{c\top} \mathbf{X}^c) + \kappa h(\mathbf{L}^c), \\ \text{s.t. } &\mathbf{L}^c \in \mathbb{R}_+^{K \times N}, \mathbf{L}^{c\top} \mathbf{1}^K = \mathbf{1}^N, \mathbf{L}^c \mathbf{1}^N = \frac{N}{K} \mathbf{1}^K, \end{aligned} \quad (9)$$

where  $h(\mathbf{L}^c) = \sum_{n,k} -l_{i_n,k} \log l_{i_n,k}$  is an entropy, and  $\kappa > 0$  is a parameter that controls the smoothness of distribution. With the soft assignment relaxation and the extra regularization term  $h(\mathbf{L}^c)$ , the solver of Eq. 9 can be given as [24]:

$$\mathbf{L}^c = \text{diag}(\mathbf{u}) \exp\left(\frac{\mathbf{P}^{c\top} \mathbf{X}^c}{\kappa}\right) \text{diag}(\mathbf{v}), \quad (10)$$

where  $\mathbf{u} \in \mathbb{R}^K$  and  $\mathbf{v} \in \mathbb{R}^N$  are renormalization vectors, computed by few steps of *Sinkhorn-Knopp iteration* [24]. Our online clustering is highly efficient on GPU, as it only involves a couple of matrix multiplications; in practice, clustering 10K pixels into 10 prototypes takes only 2.5 ms.

**Pixel-Prototype Contrastive Learning.** With the assignment probability matrix  $\mathbf{L}^c = [\mathbf{l}_{i_n}]_{n=1}^N \in [0, 1]^{K \times N}$ , we online group the training pixels  $\mathcal{I}^c = \{i_n\}_{n=1}^N$  into  $K$  prototypes  $\{\mathbf{p}_{c,k}\}_{k=1}^K$  within class  $c$ . After all the samples in current batch are processed, each pixel  $i$  is assigned to  $k_i$ -th prototype of class  $c_i$ , where  $k_i = \arg \max_k \{l_{i,k}\}_{k=1}^K$  and  $l_{i,k} \in \mathbf{l}_i$ . It is natural to derive a training objective for prototype assignment prediction, i.e., maximize the prototype assignment posterior probability. This can be viewed as a pixel-prototype contrastive learning strategy, and addresses

the first limitation of Eq. 7:

$$\mathcal{L}_{\text{PPC}} = -\log \frac{\exp(\mathbf{i}^\top \mathbf{p}_{c_i, k_i} / \tau)}{\exp(\mathbf{i}^\top \mathbf{p}_{c_i, k_i} / \tau) + \sum_{\mathbf{p}^- \in \mathcal{P}^-} \exp(\mathbf{i}^\top \mathbf{p}^- / \tau)}, \quad (11)$$

where  $\mathcal{P}^- = \{\mathbf{p}_{c, k}\}_{c, k=1}^{C, K} / \mathbf{p}_{c_i, k_i}$ , and the temperature  $\tau$  controls the concentration level of representations. Intuitively, Eq. 11 enforces each pixel embedding  $\mathbf{i}$  to be similar with its assigned (‘positive’) prototype  $\mathbf{p}_{c_i, k_i}$ , and dissimilar with other  $CK - 1$  irrelevant (‘negative’) prototypes  $\mathcal{P}^-$ . Compared with prior pixel-wise metric learning based segmentation models [111], which consume numerous negative pixel samples, our method only needs  $CK$  prototypes for pixel-prototype contrast computation, *neither* causing large memory cost *nor* requiring heavy pixel pair-wise comparison.

**Pixel-Prototype Distance Optimization.** Building upon the relative comparison over pixel-class/-prototype distances, Eq. 7 and Eq. 11 inspire inter-class/-cluster discriminativeness, but less consider reducing the intra-cluster variation, *i.e.*, making pixel features of the same prototype compact. Thus a compactness-aware loss is used for further regularizing representations by directly minimizing the distance between each embedded pixel and its assigned prototype:

$$\mathcal{L}_{\text{PPD}} = (1 - \mathbf{i}^\top \mathbf{p}_{c_i, k_i})^2. \quad (12)$$

Note that both  $\mathbf{i}$  and  $\mathbf{p}_{c_i, k_i}$  are  $\ell_2$ -normalized. This training objective minimizes intra-cluster variations while maintaining separation between features with different prototype assignments, making our model more robust against outliers.

**Network Learning and Prototype Update.** Our model is a nonparametric approach that learns semantic segmentation by directly optimizing the pixel embedding space  $\phi$ . It is called nonparametric because it constructs prototype hypotheses directly from the training pixel samples themselves. Thus the parameters of the feature extractor  $\phi$  are learned through stochastic gradient descent, by minimizing the combinatorial loss over all the training pixel samples:

$$\mathcal{L}_{\text{SEG}} = \mathcal{L}_{\text{CE}} + \lambda_1 \mathcal{L}_{\text{PPC}} + \lambda_2 \mathcal{L}_{\text{PPD}}. \quad (13)$$

Meanwhile, the non-learnable prototypes  $\{\mathbf{p}_{c, k}\}_{c, k=1}^{C, K}$  are *not* learned by stochastic gradient descent, but are computed as the centers of the corresponding embedded pixel samples. To do so, we let the prototypes evolve continuously by accounting for the online clustering results. Particularly, after each training iteration, each prototype is updated as:

$$\mathbf{p}_{c, k} \leftarrow \mu \mathbf{p}_{c, k} + (1 - \mu) \bar{\mathbf{i}}_{c, k}, \quad (14)$$

where  $\mu \in [0, 1]$  is a momentum coefficient, and  $\bar{\mathbf{i}}_{c, k}$  indicates the  $\ell_2$ -normalized, mean vector of the embedded training pixels, which are assigned to prototype  $\mathbf{p}_{c, k}$  by online clustering. With the clear meaning of the prototypes, our segmentation procedure can be intuitively understood as retrieving the most similar prototypes (sub-class centers). Fig. 3 provides prototype retrieval results for *person* and *car* with  $K = 3$  prototypes for each. The prototypes are associ-



Figure 3. **Visualization of pixel-prototype similarity** for *person* (top) and *car* (bottom) classes. Please refer to §3 for details.

ated with different colors (*i.e.*, red, green, and blue). For each pixel, its distance to the closest prototype is visualized using the corresponding prototype color. As can be seen, the prototypes well correspond to meaningful patterns within classes, validating their representativeness.

## 4. Related Work

In this section, we review representative work in semantic segmentation, prototype learning and metric learning.

**Semantic Segmentation.** Recent years have witnessed remarkable progress in semantic segmentation, due to the fast evolution of backbone architectures – from CNN-based (*e.g.*, VGG [97], ResNet [45]) to Transformer-like [105] (*e.g.*, ViT [30], Swin [78]), and segmentation models – from FCNs [79] to attention networks (*e.g.*, SegFormer [118]). Specifically, FCN [79] is a milestone; it learns dense prediction efficiently. Since it was proposed, numerous efforts have been devoted to improving FCN, by, for example, enlarging the receptive field [15, 16, 25, 124, 128, 135]; strengthening context cues [4, 43, 47, 48, 56, 57, 70, 75, 77, 81, 90, 126, 128, 129, 132, 138, 141]; leveraging boundary information [6, 14, 27, 66, 127, 131, 137]; incorporating neural attention [34, 41, 42, 49, 50, 63, 68, 101, 110, 112, 136]; or automating network engineering [18, 69, 72, 85]. Lately, Transformer based solutions [20, 100, 118, 139] attained growing attention; enjoying the flexibility in long-range dependency modeling, fully attentive solutions yield impressive results.

Different from current approaches that are typically built upon learnable prototypes, in pre-deep era, many segmentation systems are nonparametric [32, 73, 74, 80, 102, 103]. By absorbing their case-based reasoning ideas, we build a nonparametric segmentation network, which explicitly derives prototypes from sample clusters and hence directly optimizes the embedding space with distance metric constraints. In [62, 111], while cluster-/pixel-level metric loss is adopted to regularize representation, the pixel class is still inferred via parametric softmax. [53] purely relies on class embeddings, which, however, are fully trainable. Thus [53, 62, 111] are all parametric methods. As far as we know, [52] is the only non-learnable prototype, deep learning based semantic segmentation model. But [52] treats image regions as prototypes, incurring huge memory and computational demand. Besides, [52] only considers the relative difference between inter-and intra-class sample-prototype distances like the parametric counterparts. Our method is more principled with fewer heuristic designs. Unlike [52], we rep-

resent prototypes as sub-cluster centers and obtain online assignments, allowing our method to scale gracefully to any dataset size. We encourage a sparse distance distribution with compactness-awareness, reinforcing the embedding discrimination. With a broader view, a few embedding based instance segmentation approaches [26, 86] can be viewed as nonparametric, *i.e.*, treat instance centroids as prototypes.

**Prototype Learning.** Cognitive psychological studies evidence that people use past cases as models when learning to solve problems [1, 87, 125]. Among various machine learning algorithms, ranging from classical statistics based methods to Support Vector Machine to Multilayer Perceptrons [9, 31, 33, 96], prototype based classification gains particular interest, due to its exemplar-driven nature and intuitive interpretation: observations are directly compared with representative examples. Based on the nearest neighbors rule – the earliest prototype learning method [23], many famous, nonparametric classifiers are proposed [36], such as Learning Vector Quantization (LVQ) [61], generalized LVQ [94], and Neighborhood Component Analysis [37, 93]. There has been a recent surge of interest to integrate deep learning into prototype learning, showing good potential in few-shot [98], zero-shot [54], and unsupervised learning [116, 120], as well as supervised classification [38, 83, 115, 123] and interpretable networks [65]. Remarkably, as many few-shot segmentation models can be viewed as prototype-based networks [29, 106, 109], our work sheds light on the possibility of closer collaboration between the two segmentation fields.

**Metric Learning.** The selection of proper distance measure impacts the success of prototype based learners [8]; metric learning and prototype learning are naturally related. As the literature on metric learning is vast [58], only the most relevant ones are discussed. The goal of metric learning is to learn a distance metric/embedding such that similar samples are pulled together and dissimilar samples are pushed away. It has shown a significant benefit by learning deep representation using metric loss functions (*e.g.*, contrastive loss [40], triplet loss [95],  $n$ -pair loss [99]) for applications (*e.g.*, image retrieval [107], face recognition [95]). Recently, metric learning showed good potential in unsupervised representation learning. Specifically, many *instance-based* approaches use the contrastive loss [39, 88] to explicitly compare pairs of image representations, so as to push away features from different images while pulling together those from transformations of the same image [17, 19, 44, 46, 88]. Since computing all the pairwise comparisons on a large dataset is challenging, some *clustering-based* methods turn to discriminate between groups of images with similar features instead of individual images [2, 5, 12, 13, 64, 104, 119, 121, 122]. Our prototype-anchored metric learning strategy shares a similar spirit of posing metric constraints over prototype (cluster) assignments, but it is to reshape the pixel segmentation embedding space with explicit supervision.

## 5. Experiment

### 5.1. Experimental Setup

**Datasets.** Our experiments are conducted on three datasets:

- **ADE20K** [140] is a large-scale scene parsing benchmark that covers 150 stuff/object categories. The dataset is divided into 20k/2k/3k images for `train/val/test`.
- **Cityscapes** [22] has 5k finely annotated urban scene images, with 2,975/500/1,524 for `train/val/test`. The segmentation performance is evaluated over 19 challenging categories, such as rider, bicycle, and traffic light.
- **COCO-Stuff** [10] has 10k images gathered from COCO [71], with 9k and 1k for `train` and `test`, respectively. There are 172 semantic categories in total, including 80 objects, 91 stuffs and 1 unlabeled.

**Training.** Our method is implemented on MMSegmentation [21], following default training settings. In particular, all backbones are initialized using corresponding weights pre-trained on ImageNet-1K [92], while remaining layers are randomly initialized. We use standard data augmentation techniques, including random scale jittering with a factor in [0.5, 2], random horizontal flipping, random cropping as well as random color jittering. We train models using SGD/AdamW for FCN/attention-based models, respectively. The learning rate is scheduled following the polynomial annealing policy. In addition, for Cityscapes, we use a batch size of 8, and a training crop size of 768×768. For ADE20K and COCO-Stuff, we use a crop size of 512×512 and train the models with batch size 16. The models are trained for 160k, 160k, and 80k iterations on Cityscapes, ADE20K and COCO-Stuff, respectively. Exceptionally, for ablation study, we train models for 40K iterations. The hyper-parameters are empirically set to:  $K=10$ ,  $m=0.999$ ,  $\tau=0.1$ ,  $\kappa=0.05$ ,  $\lambda_1=0.01$ ,  $\lambda_2=0.01$ . **Testing.** For ADE20K and COCO-Stuff, we rescale the short scale of the image to training crop size, with the aspect ratio kept unchanged. For Cityscapes, we adopt sliding window inference with the window size 768×768. For simplicity, we do not apply *any* test-time data augmentation. Our model is implemented in PyTorch and trained on eight Tesla V100 GPUs with a 32GB memory per-card. Testing is conducted on the same machine.

**Baselines.** We mainly compare with four widely recognized segmentation models, *i.e.*, two FCN based (*i.e.*, FCN [79], HRNet [108]) and two attention based (*i.e.*, Swin [78] and SegFormer [118]). For fair comparison, all the models are based on our reproduction, following the hyper-parameter and augmentation recipes used in MMSegmentation [21].

**Evaluation Metric.** Following conventions [15, 79], mean intersection-over-union (mIoU) is adopted for evaluation.

### 5.2. Comparison to State-of-the-Arts

**ADE20K [140] val.** Table 1 reports comparisons with representative models on ADE20K `val`. Our nonparametric




Figure 4. **Qualitative results** of Segformer [118] and our approach (from left to right: ADE20K [140], Cityscapes [22], COCO-Stuff [10]).

Method	Backbone	# Param (M)	mIoU (%)
DeepLabV3+ [ECCV18] [16]	ResNet-101 [45]	62.7	44.1
OCR [ECCV20] [129]	HRNetV2-W48 [108]	70.3	45.6
MaskFormer [NeurIPS21] [20]	ResNet-101 [45]	60.0	46.0
UperNet [ECCV20] [117]	Swin-Base [78]	121.0	48.4
OCR [ECCV20] [129]	HRFormer-B [130]	70.3	48.7
SETR [CVPR21] [139]	ViT-Large [30]	318.3	50.2
Segmenter [ICCV21] [100]	ViT-Large [30]	334.0	51.8
†MaskFormer [NeurIPS21] [20]	Swin-Base [78]	102.0	52.7
FCN [CVPR15] [79]	ResNet-101 [45]	68.6	39.9
<b>Ours</b>		68.5	<b>41.1 ↑ 1.2</b>
HRNet [PAMI20] [108]	HRNetV2-W48 [108]	65.9	42.0
<b>Ours</b>		65.8	<b>43.0 ↑ 1.0</b>
Swin [CVPR21] [78]	Swin-Base [78]	90.6	48.0
<b>Ours</b>		90.5	<b>48.6 ↑ 0.6</b>
SegFormer [NeurIPS21] [118]	MiT-B4 [118]	64.1	50.9
<b>Ours</b>		64.0	<b>51.7 ↑ 0.8</b>

†: backbone is pre-trained on ImageNet-22K.

Table 1. **Quantitative results** (§5.2) on ADE20K [140] val.

Method	Backbone	# Param (M)	mIoU (%)
PSPNet [CVPR17] [135]	ResNet-101 [45]	65.9	78.4
PSANet [ECCV18] [136]	ResNet-101 [45]	-	78.6
AAF [ECCV18] [59]	ResNet-101 [45]	-	79.1
Segmenter [ICCV21] [100]	ViT-Large [30]	322.0	79.1
ContrastiveSeg [ICCV21] [111]	ResNet-101 [45]	58.0	79.2
MaskFormer [NeurIPS21] [20]	ResNet-101 [45]	60.0	80.3
DeepLabV3+ [ECCV18] [16]	ResNet-101 [45]	62.7	80.9
OCR [ECCV20] [129]	HRNetV2-W48 [108]	70.3	81.1
FCN [CVPR15] [79]	ResNet-101 [45]	68.6	78.1
<b>Ours</b>		68.5	<b>79.1 ↑ 1.0</b>
HRNet [PAMI20] [108]	HRNetV2-W48 [108]	65.9	80.4
<b>Ours</b>		65.8	<b>81.1 ↑ 0.7</b>
Swin [CVPR21] [78]	Swin-Base [78]	90.6	79.8
<b>Ours</b>		90.5	<b>80.6 ↑ 0.8</b>
SegFormer [NeurIPS21] [118]	MiT-B4 [118]	64.1	80.7
<b>Ours</b>		64.0	<b>81.3 ↑ 0.6</b>

Table 2. **Quantitative results** (§5.2) on Cityscapes [22] val.

scheme obtains consistent improvements over the baselines, with fewer learnable parameters. In particular, it yields **1.2%** and **1.0%** mIoU improvements over the FCN-based counterparts, *i.e.*, FCN [79] and HRNet [108]. Similar performance gains (**0.6%** and **0.8%**) are obtained over recent attention-based models, *i.e.*, Swin [78] and SegFormer [118], manifesting the high versatility of our approach.

**Cityscapes [22] val.** Table 2 shows again our compelling performance on Cityscapes val. Specifically, our approach surpasses all the competitors, *i.e.*, **1.0%** over FCN, **0.7%** over HRNet, **0.8%** over Swin, and **0.6%** over Segformer.

**COCO-Stuff [10] test.** As listed in Table 3, our approach also demonstrates promising performance on COCO-Stuff

Method	Backbone	# Param (M)	mIoU (%)
SVCNet [CVPR19] [28]	ResNet-101 [45]	-	39.6
DANet [CVPR19] [34]	ResNet-101 [45]	69.1	39.7
SpyGR [CVPR20] [67]	ResNet-101 [45]	-	39.9
MaskFormer [NeurIPS21] [20]	ResNet-101 [45]	60.0	39.8
ACNet [ICCV19] [35]	ResNet-101 [45]	-	40.1
OCR [ECCV20] [129]	HRNetV2-W48 [108]	70.3	40.5
FCN [CVPR15] [79]	ResNet-101 [45]	68.6	32.5
<b>Ours</b>		68.5	<b>34.0 ↑ 1.5</b>
HRNet [PAMI21] [108]	HRNetV2-W48 [108]	65.9	38.7
<b>Ours</b>		65.8	<b>39.9 ↑ 1.2</b>
Swin [CVPR21] [78]	Swin-Base [78]	90.6	41.5
<b>Ours</b>		90.5	<b>42.4 ↑ 0.9</b>
SegFormer [NeurIPS21] [118]	MiT-B4 [118]	64.1	42.5
<b>Ours</b>		64.0	<b>43.3 ↑ 0.8</b>

Table 3. **Quantitative results** (§5.2) on COCO-Stuff [10] test.

test. It outperforms all the baselines. Notably, with MiT-B4 [118] as the network backbone, our approach earns an mIoU score of **43.3%**, establishing a new state-of-the-art.

**Qualitative Results.** Fig. 4 provides qualitative comparison of **Ours** against Segformer [118] on representative examples in the three datasets. We observe that our approach is able to handle diverse challenging scenarios and produce more accurate results (as highlighted in red dashed boxes).

### 5.3. Scalability to Large-Vocabulary Semantic Segmentation

Today, rigorous evaluation of semantic segmentation models is mostly performed in a few category regime (*e.g.*, 19/150/172 classes for Cityscapes/ADE20K/COCO-Stuff), while the generalization to more natural large-vocabulary setting is ignored. In this section, we demonstrate the remarkable superiority of our method in large-vocabulary setting. We start with the default setting in ADE20K [140] which includes 150 semantic concepts. Then, we gradually increase the number of concepts based on their visibility frequency, and train/test models on the selected number of classes. In this experiment, we use MiT-V2 [118] as the backbone and train models for 40k iterations.

The results are summarized in Table 4, from which we find that: *i)* For the parametric scheme, the amount of prototype parameters increases with vocabulary size. For the extreme case of 10 prototypes and 847 classes, the number of prototype parameters is 6.5 M, accounting for  $\sim 20\%$  of total parameters (*i.e.*, 33.96 M). In sharp contrast, our scheme requires no any learnable prototype parameters. *ii)* Our

Method	# Proto	150 classes		300 classes		500 classes		700 classes		847 classes	
		mIoU (%)	# Param (M)	mIoU (%)	# Param (M)	mIoU (%)	# Param (M)	mIoU (%)	# Param (M)	mIoU (%)	# Param (M)
parametric nonparametric (Ours)	1	45.1	27.48 (0.12)	36.5	27.62 (0.23)	25.7	27.80 (0.39)	19.8	27.98 (0.54)	16.5	28.11 (0.65)
	1	<b>45.5</b> $\uparrow$ <b>0.4</b>	27.37 (0)	<b>37.2</b> $\uparrow$ <b>0.7</b>	27.37 (0)	<b>26.8</b> $\uparrow$ <b>1.1</b>	27.37 (0)	<b>21.2</b> $\uparrow$ <b>1.4</b>	27.37 (0)	<b>18.1</b> $\uparrow$ <b>1.6</b>	27.37 (0)
parametric nonparametric (Ours)	10	45.7	28.56 (1.2)	37.0	29.66 (2.3)	26.6	31.26 (3.9)	20.8	32.86 (5.4)	17.7	33.96 (6.5)
	10	<b>46.4</b> $\uparrow$ <b>0.7</b>	27.37 (0)	<b>37.8</b> $\uparrow$ <b>0.8</b>	27.37 (0)	<b>27.9</b> $\uparrow$ <b>1.3</b>	27.37 (0)	<b>22.1</b> $\uparrow$ <b>1.3</b>	27.37 (0)	<b>19.4</b> $\uparrow$ <b>1.7</b>	27.37 (0)

Table 4. **Scalability study** (§5.3) of our nonparametric model against the parametric baseline (*i.e.*, SegFormer [118]) on ADE20K [140]. For each model variant, we report its segmentation mIoU, parameter numbers of the entire model as well as the prototypes (in the bracket).

$\mathcal{L}_{CE}$ (Eq. 7)	$\mathcal{L}_{PPC}$ (Eq. 11)	$\mathcal{L}_{PPD}$ (Eq. 12)	mIoU (%)	# Prototype	mIoU (%)	Coefficient $\mu$	mIoU (%)	Distance Measure	mIoU (%)
✓			45.0	$K = 1$	45.5	$\mu = 0$	44.9	Standard	45.7
✓	✓		45.9	$K = 5$	46.0	$\mu = 0.9$	45.9	Huberized	45.2
✓		✓	45.4	$K = 10$	46.4	$\mu = 0.99$	46.0	Cosine	46.4
✓	✓	✓	46.4	$K = 20$	46.5	$\mu = 0.999$	46.4		
				$K = 50$	46.4	$\mu = 0.9999$	46.3		

(a) Training Objective  $\mathcal{L}$

(b) Prototype Number  $K$

(c) Momentum Coefficient  $\mu$

(d) Distance Measure

Table 5. A set of **ablative studies** (§5.4) on ADE20K [140] val. All model variants use MiT-B2 [118] as the backbone.

method achieves consistent performance elevations against the parametric counterpart under all settings. These results well demonstrate the utility of our nonparametric scheme for unrestricted open-vocabulary semantic segmentation.

## 5.4. Diagnostic Experiment

To investigate the effect of our core designs, we conduct ablative studies on ADE20K [140] val. We use MiT-B2 [118] as the backbone and train models for 40K iterations.

**Training Objective.** We first investigate our overall training objective (*cf.* Eq. 13). As shown in Table 5a, the model with  $\mathcal{L}_{CE}$  alone achieves an mIoU score of 45.0%. Adding  $\mathcal{L}_{PPC}$  or  $\mathcal{L}_{PPD}$  individually brings gains (*i.e.*, **0.9%/0.4%**), revealing the value to explicitly learn pixel-prototype relations. Combining all the losses together leads to the best performance, yielding an mIoU score of 46.4%.

**Prototype Number Per Class  $K$ .** Table 5b reports the performance of our approach with regard to the number of prototype per class. For  $K = 1$ , we directly represent each class as the mean embedding of its pixel samples. The pixel assignment is based simply on ground-truth labels, without using online clustering (Eqs. 8-9). This baseline obtains a score of 45.5%. Further, when using more prototypes (*i.e.*,  $K = 3$ ), we see a clear performance boost (*i.e.*, 45.5%  $\rightarrow$  46.0%). The score further improves when allowing 5 or 10 prototypes; however, increasing  $K$  beyond 10 gives marginal returns in performance. As a result, we set  $K = 10$  for a better trade-off between accuracy and computation cost. This study confirms our motivation to use multiple prototypes for capturing intra-class variations.

**Coefficient  $\mu$ .** Table 5c quantifies the effect of momentum coefficient ( $\mu$  in Eq. 14) which controls the speed of prototype updating. The model performs reasonably well using a relatively large coefficient (*i.e.*,  $\mu \in [0.999, 0.9999]$ ), showing that a slow updating is beneficial. When  $\mu$  is 0.9 or 0.99, the performance decreases, and drops considerably at

the extreme case of  $\mu = 0$ .

**Distance Measure.** By default, we use cosine distance (refer to as ‘Cosine’) to measure pixel-prototype similarity as denoted in Eq. 6, Eq. 11 and Eq. 12. However, other choices are also applicable. Here we study two alternatives. The first is the standard Euclidean distance (*i.e.*, ‘Standard’), *i.e.*,  $\langle \mathbf{x}, \mathbf{y} \rangle = \|\mathbf{x} - \mathbf{y}\|_2$ . In contrast to ‘Cosine’, here  $\mathbf{x}$  and  $\mathbf{y}$  are un-normalized real-valued vectors. To handle the non-differentiability in ‘Standard’, we further study an approximated Huber-like function [51] (‘Huberized’), *i.e.*,  $\langle \mathbf{x}, \mathbf{y} \rangle = \delta(\sqrt{\|\mathbf{x} - \mathbf{y}\|^2/\delta^2 + 1} - 1)$ . The hyper-parameter  $\delta$  is empirically set to 0.1. As we find from Table 5d that ‘Cosine’ performs much better than other un-normalized Euclidean measurements. The Huberized norm does not show any advantage over ‘Standard’.

## 6. Conclusion and Discussion

The vast majority of recent effort in this field seek to learn parametric class representations for pixel-wise recognition. In contrast, this paper explores an exemplar-based regime. This leads to a nonparametric segmentation framework, where several typical points in the embedding space are selected as class prototypical representation, and distance to the prototypes determines how a pixel sample is classified. It enjoys several advantages: **i)** explicit prototypical representation for class-level statistics modeling; **ii)** better generalization with nonparametric pixel-category prediction; and **iii)** direct optimization of the feature embedding space. Our framework is elegant, general, and yields outstanding performance. It also comes with some intriguing questions. For example, to pursue better interpretability, one can optimize the prototypes to directly resemble pixel- or region-level observations [52, 65]. Overall, we feel the results in this paper warrant further exploration in this direction.

**Acknowledgements** This work was supported by CCF-Baidu Open Fund and ARC DECRA DE220101390.



## References

- [1] Agnar Aamodt and Enric Plaza. Case-based reasoning: Foundational issues, methodological variations, and system approaches. *AI Communications*, 7(1):39–59, 1994. 6
- [2] Yuki Markus Asano, Christian Rupprecht, and Andrea Vedaldi. Self-labelling via simultaneous clustering and representation learning. In *ICLR*, 2020. 4, 6
- [3] Andreas Backhaus and Udo Seiffert. Classification in high-dimensional spectral data: Accuracy vs. interpretability vs. model size. *Neurocomputing*, 131:15–22, 2014. 2
- [4] Vijay Badrinarayanan, Alex Kendall, and Roberto Cipolla. Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE TPAMI*, 39(12):2481–2495, 2017. 5
- [5] Miguel A Bautista, Artsiom Sanakoyeu, Ekaterina Tikhoncheva, and Björn Ommer. Cliqecnn: Deep unsupervised exemplar learning. In *NeurIPS*, 2016. 6
- [6] Gedas Bertasius, Jianbo Shi, and Lorenzo Torresani. Semantic segmentation with boundary neural fields. In *CVPR*, 2016. 5
- [7] Michael Biehl, Barbara Hammer, Petra Schneider, and Thomas Villmann. Metric learning for prototype-based classification. In *Innovations in Neural Information Paradigms and Applications*, pages 183–199. 2009. 2
- [8] Michael Biehl, Barbara Hammer, and Thomas Villmann. Distance measures for prototype based classification. In *International Workshop on Brain-Inspired Computing*, 2013. 3, 6
- [9] Christopher M Bishop. Pattern recognition. *Machine learning*, 128(9), 2006. 6
- [10] Holger Caesar, Jasper Uijlings, and Vittorio Ferrari. Coco-stuff: Thing and stuff classes in context. In *CVPR*, 2018. 2, 6, 7
- [11] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, 2020. 1
- [12] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In *ECCV*, 2018. 6
- [13] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. In *NeurIPS*, 2020. 4, 6
- [14] Liang-Chieh Chen, Jonathan T Barron, George Papandreou, Kevin Murphy, and Alan L Yuille. Semantic image segmentation with task-specific edge detection using cnns and a discriminatively trained domain transform. In *CVPR*, 2016. 5
- [15] 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 TPAMI*, 40(4):834–848, 2017. 1, 5, 6
- [16] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *ECCV*, 2018. 5, 7
- [17] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, 2020. 6
- [18] Wuyang Chen, Xinyu Gong, Xianming Liu, Qian Zhang, Yuan Li, and Zhangyang Wang. FASTERseg: Searching for faster real-time semantic segmentation. In *ICLR*, 2019. 5
- [19] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020. 6
- [20] Bowen Cheng, Alexander G. Schwing, and Alexander Kirillov. Per-pixel classification is not all you need for semantic segmentation. In *NeurIPS*, 2021. 1, 5, 7
- [21] MMSegmentation Contributors. MMSegmentation: Openmmlab semantic segmentation toolbox and benchmark. <https://github.com/open-mmlab/mms Segmentation>, 2020. 6
- [22] 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*, 2016. 2, 6, 7
- [23] Thomas Cover and Peter Hart. Nearest neighbor pattern classification. *IEEE TIT*, 13(1):21–27, 1967. 1, 3, 6
- [24] Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *NeurIPS*, 2013. 4
- [25] Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable convolutional networks. In *CVPR*, 2017. 5
- [26] Bert De Brabandere, Davy Neven, and Luc Van Gool. Semantic instance segmentation for autonomous driving. In *CVPR Workshop*, 2017. 6
- [27] Henghui Ding, Xudong Jiang, Ai Qun Liu, Nadia Magnat Thalmann, and Gang Wang. Boundary-aware feature propagation for scene segmentation. In *CVPR*, 2019. 5
- [28] Henghui Ding, Xudong Jiang, Bing Shuai, Ai Qun Liu, and Gang Wang. Semantic correlation promoted shape-variant context for segmentation. In *CVPR*, 2019. 7
- [29] Nanqing Dong and Eric P Xing. Few-shot semantic segmentation with prototype learning. In *BMVC*, 2018. 6
- [30] 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. In *ICLR*, 2020. 5, 7
- [31] Richard O Duda, Peter E Hart, et al. *Pattern classification and scene analysis*, volume 3. Wiley New York, 1973. 1, 3, 6
- [32] David Eigen and Rob Fergus. Nonparametric image parsing using adaptive neighbor sets. In *CVPR*, 2012. 5
- [33] Jerome H Friedman. *The elements of statistical learning: Data mining, inference, and prediction*. Springer, 2009. 3, 6
- [34] Jun Fu, Jing Liu, Haijie Tian, Yong Li, Yongjun Bao, Zhiwei Fang, and Hanqing Lu. Dual attention network for scene segmentation. In *CVPR*, 2019. 1, 5, 7
- [35] Jun Fu, Jing Liu, Yuhang Wang, Yong Li, Yongjun Bao, Jinhui Tang, and Hanqing Lu. Adaptive context network for scene parsing. In *ICCV*, 2019. 7
- [36] Salvador Garcia, Joaquin Derrac, Jose Cano, and Francisco Herrera. Prototype selection for nearest neighbor classi-

- fication: Taxonomy and empirical study. *IEEE TPAMI*, 34(3):417–435, 2012. 3, 6
- [37] Jacob Goldberger, Geoffrey E Hinton, Sam Roweis, and Russ R Salakhutdinov. Neighbourhood components analysis. In *NeurIPS*, 2004. 6
- [38] Samantha Guerriero, Barbara Caputo, and Thomas Mensink. Deepncm: Deep nearest class mean classifiers. In *ICLR workshop*, 2018. 6
- [39] Michael Gutmann and Aapo Hyvärinen. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *AISTATS*, 2010. 6
- [40] Raia Hadsell, Sumit Chopra, and Yann LeCun. Dimensionality reduction by learning an invariant mapping. In *CVPR*, 2006. 2, 6
- [41] Adam W Harley, Konstantinos G Derpanis, and Iasonas Kokkinos. Segmentation-aware convolutional networks using local attention masks. In *ICCV*, 2017. 5
- [42] Junjun He, Zhongying Deng, and Yu Qiao. Dynamic multi-scale filters for semantic segmentation. In *ICCV*, 2019. 5
- [43] Junjun He, Zhongying Deng, Lei Zhou, Yali Wang, and Yu Qiao. Adaptive pyramid context network for semantic segmentation. In *CVPR*, 2019. 5
- [44] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020. 6
- [45] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 2, 5, 7
- [46] R Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Phil Bachman, Adam Trischler, and Yoshua Bengio. Learning deep representations by mutual information estimation and maximization. In *ICLR*, 2019. 6
- [47] Chi-Wei Hsiao, Cheng Sun, Hwann-Tzong Chen, and Min Sun. Specialize and fuse: Pyramidal output representation for semantic segmentation. In *ICCV*, 2021. 5
- [48] Hanzhe Hu, Deyi Ji, Weihao Gan, Shuai Bai, Wei Wu, and Junjie Yan. Class-wise dynamic graph convolution for semantic segmentation. In *ECCV*, 2020. 5
- [49] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *CVPR*, 2018. 1, 5
- [50] Zilong Huang, Xinggang Wang, Lichao Huang, Chang Huang, Yunchao Wei, and Wenyu Liu. Ccnet: Criss-cross attention for semantic segmentation. In *ICCV*, 2019. 5
- [51] Peter J Huber. Robust regression: asymptotics, conjectures and monte carlo. *The annals of statistics*, pages 799–821, 1973. 8
- [52] Jyh-Jing Hwang, Stella X Yu, Jianbo Shi, Maxwell D Collins, Tien-Ju Yang, Xiao Zhang, and Liang-Chieh Chen. Segsort: Segmentation by discriminative sorting of segments. In *ICCV*, 2019. 5, 8
- [53] Shipra Jain, Danda Pani Paudel, Martin Danelljan, and Luc Van Gool. Scaling semantic segmentation beyond 1k classes on a single gpu. In *ICCV*, 2021. 5
- [54] Saumya Jetley, Bernardino Romera-Paredes, Sadeep Jayasumana, and Philip Torr. Prototypical priors: From improving classification to zero-shot learning. *arXiv preprint arXiv:1512.01192*, 2015. 6
- [55] Xu Ji, Joao F Henriques, and Andrea Vedaldi. Invariant information clustering for unsupervised image classification and segmentation. In *ICCV*, 2019. 4
- [56] Zhenchao Jin, Tao Gong, Dongdong Yu, Qi Chu, Jian Wang, Changhu Wang, and Jie Shao. Mining contextual information beyond image for semantic segmentation. In *ICCV*, 2021. 5
- [57] Zhenchao Jin, Bin Liu, Qi Chu, and Nenghai Yu. Isnet: Integrate image-level and semantic-level context for semantic segmentation. In *ICCV*, 2021. 5
- [58] Mahmut Kaya and Hasan Şakir Bilge. Deep metric learning: A survey. *Symmetry*, 11(9):1066, 2019. 6
- [59] Tsung-Wei Ke, Jyh-Jing Hwang, Ziwei Liu, and Stella X Yu. Adaptive affinity fields for semantic segmentation. In *ECCV*, 2018. 7
- [60] Barbara J Knowlton and Larry R Squire. The learning of categories: Parallel brain systems for item memory and category knowledge. *Science*, 262(5140):1747–1749, 1993. 2, 3
- [61] Teuvo Kohonen. The self-organizing map. *Neurocomputing*, 21(1-3):1–6, 1998. 6
- [62] Shu Kong and Charless C Fowlkes. Recurrent pixel embedding for instance grouping. In *CVPR*, 2018. 5
- [63] Hanchao Li, Pengfei Xiong, Jie An, and Lingxue Wang. Pyramid attention network for semantic segmentation. *arXiv preprint arXiv:1805.10180*, 2018. 5
- [64] Junnan Li, Pan Zhou, Caiming Xiong, and Steven CH Hoi. Prototypical contrastive learning of unsupervised representations. In *ICLR*, 2020. 2, 6
- [65] Oscar Li, Hao Liu, Chaofan Chen, and Cynthia Rudin. Deep learning for case-based reasoning through prototypes: A neural network that explains its predictions. In *AAAI*, 2018. 6, 8
- [66] Xiangtai Li, Xia Li, Li Zhang, Guangliang Cheng, Jianping Shi, Zhouchen Lin, Shaohua Tan, and Yunhai Tong. Improving semantic segmentation via decoupled body and edge supervision. In *ECCV*, 2020. 5
- [67] Xia Li, Yibo Yang, Qijie Zhao, Tiancheng Shen, Zhouchen Lin, and Hong Liu. Spatial pyramid based graph reasoning for semantic segmentation. In *CVPR*, 2020. 7
- [68] Xia Li, Zhisheng Zhong, Jianlong Wu, Yibo Yang, Zhouchen Lin, and Hong Liu. Expectation-maximization attention networks for semantic segmentation. In *ICCV*, 2019. 5
- [69] Yanwei Li, Lin Song, Yukang Chen, Zeming Li, Xiangyu Zhang, Xinggang Wang, and Jian Sun. Learning dynamic routing for semantic segmentation. In *CVPR*, 2020. 5
- [70] Guosheng Lin, Anton Milan, Chunhua Shen, and Ian Reid. Refinenet: Multi-path refinement networks for high-resolution semantic segmentation. In *CVPR*, 2017. 5
- [71] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. 6
- [72] Chenxi Liu, Liang-Chieh Chen, Florian Schroff, Hartwig Adam, Wei Hua, Alan L Yuille, and Li Fei-Fei. Auto-deeplab: Hierarchical neural architecture search for semantic image segmentation. In *CVPR*, 2019. 5
- [73] Ce Liu, Jenny Yuen, and Antonio Torralba. Sift flow: Dense correspondence across scenes and its applications. *IEEE TPAMI*, 33(5):978–994, 2010. 5
- [74] Ce Liu, Jenny Yuen, and Antonio Torralba. Nonparametric scene parsing via label transfer. *IEEE TPAMI*,

- 33(12):2368–2382, 2011. 5
- [75] Mingyuan Liu, Dan Schonfeld, and Wei Tang. Exploit visual dependency relations for semantic segmentation. In *CVPR*, 2021. 5
- [76] Weyang Liu, Yandong Wen, Zhiding Yu, and Meng Yang. Large-margin softmax loss for convolutional neural networks. In *ICML*, 2016. 3
- [77] Ziwei Liu, Xiao Xiao Li, Ping Luo, Chen Change Loy, and Xiaoou Tang. Deep learning markov random field for semantic segmentation. *IEEE TPAMI*, 40(8):1814–1828, 2017. 5
- [78] 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*, 2021. 2, 5, 6, 7
- [79] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015. 1, 5, 6, 7
- [80] Tomasz Malisiewicz and Alexei A Efros. Recognition by association via learning per-exemplar distances. In *CVPR*, 2008. 5
- [81] Sachin Mehta, Mohammad Rastegari, Anat Caspi, Linda Shapiro, and Hannaneh Hajishirzi. Espnet: Efficient spatial pyramid of dilated convolutions for semantic segmentation. In *ECCV*, 2018. 5
- [82] Tim Meinhardt, Alexander Kirillov, Laura Leal-Taixe, and Christoph Feichtenhofer. Trackformer: Multi-object tracking with transformers. *arXiv preprint arXiv:2101.02702*, 2021. 1
- [83] Pascal Mettes, Elise van der Pol, and Cees Snoek. Hyper-spherical prototype networks. In *NeurIPS*, 2019. 3, 6
- [84] Tom M Mitchell. *The need for biases in learning generalizations*. Department of Computer Science, Laboratory for Computer Science Research, 1980. 3
- [85] Vladimir Nekrasov, Hao Chen, Chunhua Shen, and Ian Reid. Fast neural architecture search of compact semantic segmentation models via auxiliary cells. In *CVPR*, 2019. 5
- [86] Davy Neven, Bert De Brabandere, Marc Proesmans, and Luc Van Gool. Instance segmentation by jointly optimizing spatial embeddings and clustering bandwidth. In *CVPR*, 2019. 6
- [87] Allen Newell, Herbert Alexander Simon, et al. *Human problem solving*, volume 104. 1972. 6
- [88] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 6
- [89] Tianyu Pang, Kun Xu, Yinpeng Dong, Chao Du, Ning Chen, and Jun Zhu. Rethinking softmax cross-entropy loss for adversarial robustness. In *ICLR*, 2020. 2, 4
- [90] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 2015. 5
- [91] Eleanor H Rosch. Natural categories. *Cognitive psychology*, 4(3):328–350, 1973. 2, 3
- [92] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael S. Bernstein, Alexander C. Berg, and Fei-Fei Li. Imagenet large scale visual recognition challenge. *IJCV*, 115(3):211–252, 2015. 6
- [93] Ruslan Salakhutdinov and Geoff Hinton. Learning a non-linear embedding by preserving class neighbourhood structure. In *Artificial Intelligence and Statistics*, pages 412–419, 2007. 6
- [94] Atsushi Sato and Keiji Yamada. Generalized learning vector quantization. In *NeurIPS*, 1995. 6
- [95] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *CVPR*, 2015. 3, 6
- [96] John Shawe-Taylor, Nello Cristianini, et al. *Kernel methods for pattern analysis*. Cambridge university press, 2004. 6
- [97] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015. 5
- [98] Jake Snell, Kevin Swersky, and Richard S Zemel. Prototypical networks for few-shot learning. In *NeurIPS*, 2017. 6
- [99] Kihyuk Sohn. Improved deep metric learning with multi-class n-pair loss objective. In *NeurIPS*, 2016. 6
- [100] Robin Strudel, Ricardo Garcia, Ivan Laptev, and Cordelia Schmid. Segmenter: Transformer for semantic segmentation. In *ICCV*, 2021. 1, 5, 7
- [101] Guolei Sun, Wenguan Wang, Jifeng Dai, and Luc Van Gool. Mining cross-image semantics for weakly supervised semantic segmentation. In *ECCV*, 2020. 5
- [102] Joseph Tighe and Svetlana Lazebnik. Superparsing: scalable nonparametric image parsing with superpixels. In *ECCV*, 2010. 5
- [103] Antonio Torralba, Rob Fergus, and William T Freeman. 80 million tiny images: A large data set for nonparametric object and scene recognition. *IEEE TPAMI*, 30(11):1958–1970, 2008. 5
- [104] Wouter Van Gansbeke, Simon Vandenhende, Stamatios Georgoulis, Marc Proesmans, and Luc Van Gool. Scan: Learning to classify images without labels. In *ECCV*, 2020. 6
- [105] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. 1, 5
- [106] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. In *NeurIPS*, 2016. 6
- [107] Jiang Wang, Yang Song, Thomas Leung, Chuck Rosenberg, Jingbin Wang, James Philbin, Bo Chen, and Ying Wu. Learning fine-grained image similarity with deep ranking. In *CVPR*, 2014. 6
- [108] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Mingkui Tan, Xinggang Wang, et al. Deep high-resolution representation learning for visual recognition. *IEEE TPAMI*, 2020. 2, 6, 7
- [109] Kaixin Wang, Jun Hao Liew, Yingtian Zou, Daquan Zhou, and Jiashi Feng. Panet: Few-shot image semantic segmentation with prototype alignment. In *ICCV*, 2019. 6
- [110] Wenguan Wang, Tianfei Zhou, Siyuan Qi, Jianbing Shen, and Song-Chun Zhu. Hierarchical human semantic parsing with comprehensive part-relation modeling. *IEEE TPAMI*, 2021. 5
- [111] Wenguan Wang, Tianfei Zhou, Fisher Yu, Jifeng Dai, Ender Konukoglu, and Luc Van Gool. Exploring cross-image

- pixel contrast for semantic segmentation. In *ICCV*, 2021. [2](#), [5](#), [7](#)
- [112] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *CVPR*, 2018. [5](#)
- [113] Yuqing Wang, Zhaoliang Xu, Xinlong Wang, Chunhua Shen, Baoshan Cheng, Hao Shen, and Huaxia Xia. End-to-end video instance segmentation with transformers. In *CVPR*, 2021. [1](#)
- [114] Yandong Wen, Kaipeng Zhang, Zhifeng Li, and Yu Qiao. A discriminative feature learning approach for deep face recognition. In *ECCV*, 2016. [3](#)
- [115] Zhirong Wu, Alexei A Efros, and Stella X Yu. Improving generalization via scalable neighborhood component analysis. In *ECCV*, 2018. [2](#), [6](#)
- [116] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *CVPR*, 2018. [2](#), [6](#)
- [117] Tete Xiao, Yingcheng Liu, Bolei Zhou, Yuning Jiang, and Jian Sun. Unified perceptual parsing for scene understanding. In *ECCV*, 2018. [7](#)
- [118] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. In *NeurIPS*, 2021. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#), [8](#)
- [119] Junyuan Xie, Ross Girshick, and Ali Farhadi. Unsupervised deep embedding for clustering analysis. In *ICML*, 2016. [6](#)
- [120] Wenjia Xu, Yongqin Xian, Jiuniu Wang, Bernt Schiele, and Zeynep Akata. Attribute prototype network for zero-shot learning. In *NeurIPS*, 2020. [6](#)
- [121] Kouta Nakata Yaling Tao, Kentaro Takagi. Clustering-friendly representation learning via instance discrimination and feature decorrelation. In *ICLR*, 2021. [6](#)
- [122] Xueting Yan, Ishan Misra, Abhinav Gupta, Deepti Ghadiyaram, and Dhruv Mahajan. Clusterfit: Improving generalization of visual representations. In *CVPR*, 2020. [6](#)
- [123] Hong-Ming Yang, Xu-Yao Zhang, Fei Yin, and Cheng-Lin Liu. Robust classification with convolutional prototype learning. In *CVPR*, 2018. [3](#), [6](#)
- [124] Maoke Yang, Kun Yu, Chi Zhang, Zhiwei Li, and Kuiyuan Yang. Denseaspp for semantic segmentation in street scenes. In *CVPR*, 2018. [5](#)
- [125] Yi Yang, Yueting Zhuang, and Yunhe Pan. Multiple knowledge representation for big data artificial intelligence: framework, applications, and case studies. *Frontiers of Information Technology & Electronic Engineering*, 22(12):1551–1558, 2021. [6](#)
- [126] Changqian Yu, Jingbo Wang, Changxin Gao, Gang Yu, Chunhua Shen, and Nong Sang. Context prior for scene segmentation. In *CVPR*, 2020. [5](#)
- [127] Changqian Yu, Jingbo Wang, Chao Peng, Changxin Gao, Gang Yu, and Nong Sang. Learning a discriminative feature network for semantic segmentation. In *CVPR*, 2018. [5](#)
- [128] Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. In *ICLR*, 2016. [5](#)
- [129] Yuhui Yuan, Xilin Chen, and Jingdong Wang. Object-contextual representations for semantic segmentation. In *ECCV*, 2020. [5](#), [7](#)
- [130] Yuhui Yuan, Rao Fu, Lang Huang, Weihong Lin, Chao Zhang, Xilin Chen, and Jingdong Wang. Hrformer: High-resolution transformer for dense prediction. In *NeurIPS*, 2021. [7](#)
- [131] Yuhui Yuan, Jingyi Xie, Xilin Chen, and Jingdong Wang. Segfix: Model-agnostic boundary refinement for segmentation. In *ECCV*, 2020. [5](#)
- [132] Hang Zhang, Kristin Dana, Jianping Shi, Zhongyue Zhang, Xiaogang Wang, Amrith Tyagi, and Amit Agrawal. Context encoding for semantic segmentation. In *CVPR*, 2018. [5](#)
- [133] Kai Zhang, James T Kwok, and Bahram Parvin. Prototype vector machine for large scale semi-supervised learning. In *ICML*, 2009. [2](#)
- [134] Xiao Zhang, Rui Zhao, Yu Qiao, and Hongsheng Li. Rbf-softmax: Learning deep representative prototypes with radial basis function softmax. In *ECCV*, 2020. [2](#), [4](#)
- [135] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *CVPR*, 2017. [1](#), [5](#), [7](#)
- [136] Hengshuang Zhao, Yi Zhang, Shu Liu, Jianping Shi, Chen Change Loy, Dahua Lin, and Jiaya Jia. Pscanet: Point-wise spatial attention network for scene parsing. In *ECCV*, 2018. [5](#), [7](#)
- [137] Mingmin Zhen, Jinglu Wang, Lei Zhou, Shiwei Li, Tianwei Shen, Jiaxiang Shang, Tian Fang, and Long Quan. Joint semantic segmentation and boundary detection using iterative pyramid contexts. In *CVPR*, 2020. [5](#)
- [138] Shuai Zheng, Sadeep Jayasumana, Bernardino Romera-Paredes, Vibhav Vineet, Zhizhong Su, Dalong Du, Chang Huang, and Philip HS Torr. Conditional random fields as recurrent neural networks. In *ICCV*, 2015. [5](#)
- [139] Sixiao Zheng, Jiachen Lu, Hengshuang Zhao, Xiatian Zhu, Zekun Luo, Yabiao Wang, Yanwei Fu, Jianfeng Feng, Tao Xiang, Philip HS Torr, et al. Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In *CVPR*, 2021. [1](#), [2](#), [3](#), [5](#), [7](#)
- [140] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *CVPR*, 2017. [2](#), [6](#), [7](#), [8](#)
- [141] Tianfei Zhou, Liulei Li, Xueyi Li, Chun-Mei Feng, Jianwu Li, and Ling Shao. Group-wise learning for weakly supervised semantic segmentation. *IEEE TIP*, 31:799–811, 2021. [5](#)