

Distilling Self-Supervised Vision Transformers for Weakly-Supervised Few-Shot Classification & Segmentation

Dahyun Kang^{1,2*} Piotr Koniusz^{3,4} Minsu Cho² Naila Murray¹
¹Meta AI ²POSTECH ³Data61♥CSIRO ⁴Australian National University

Abstract

We address the task of weakly-supervised few-shot image classification and segmentation, by leveraging a Vision Transformer (ViT) pretrained with self-supervision. Our proposed method takes token representations from the self-supervised ViT and leverages their correlations, via self-attention, to produce classification and segmentation predictions through separate task heads. Our model is able to effectively learn to perform classification and segmentation in the absence of pixel-level labels during training, using only image-level labels. To do this it uses attention maps, created from tokens generated by the self-supervised ViT backbone, as pixel-level pseudo-labels. We also explore a practical setup with “mixed” supervision, where a small number of training images contains ground-truth pixel-level labels and the remaining images have only image-level labels. For this mixed setup, we propose to improve the pseudo-labels using a pseudo-label enhancer that was trained using the available ground-truth pixel-level labels. Experiments on Pascal-5ⁱ and COCO-20ⁱ demonstrate significant performance gains in a variety of supervision settings, and in particular when little-to-no pixel-level labels are available.

1. Introduction

Few-shot learning [22, 23, 77] is the problem of learning to perform a prediction task using only a small number of examples (*i.e.* supports) of the completed task, typically in the range of 1-10 examples. This problem setting is appealing for applications where large amounts of annotated data are expensive or impractical to collect. In computer vision, few-shot learning has been actively studied for image classification [24, 36, 46, 66, 70, 85] and semantic segmentation [17, 58, 61, 88, 89]. In this work, we focus on the combined task of few-shot classification and segmentation (FS-CS) [33], which aims to jointly predict (i) the presence of each support class, *i.e.*, multi-label classification, and (ii)

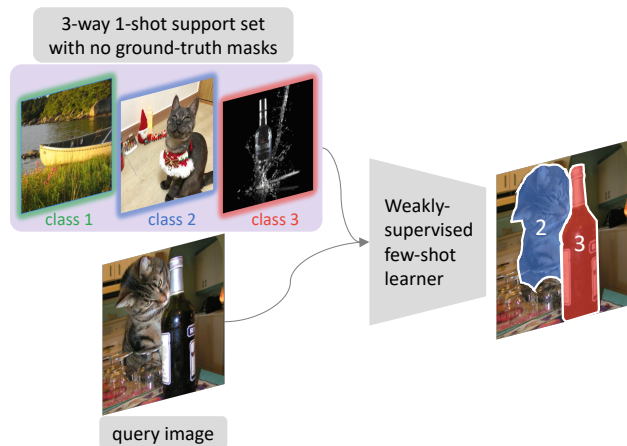


Figure 1. Weakly-supervised few-shot classification and segmentation. In this paper, we explore training few-shot models using little-to-no ground-truth segmentation masks.

its pixel-level semantic segmentation.

Few-shot learners are typically trained either by meta-learning [24, 56, 63, 71] or by transfer learning with fine-tuning [11, 12, 16]. Both these paradigms commonly assume the availability of a large-scale and fully-annotated training dataset. For FS-CS, we need Ground-Truth (GT) segmentation masks for query and support images during training. We also need these GT masks for support images during testing. A natural alternative to collecting such expensive segmentation masks would be to instead employ a weaker supervision, *e.g.* image-level [52, 59] or box-level labels [13, 30], for learning, *i.e.*, to adopt a weakly-supervised learning approach [5, 41, 43, 51]. Compared to conventional weakly-supervised learning however, weakly-supervised *few-shot* learning has rarely been explored in the literature and is significantly more challenging. This is because in few-shot learning, object classes are completely disjoint between training and testing, resulting in models which are susceptible to severe over-fitting to the classes seen during training.

In this work we tackle this challenging scenario, addressing weakly-supervised FS-CS where only image-level labels, *i.e.* class labels, are available (*cf.* Fig. 1). Inspired by

*Work done during an internship at FAIR.

recent work by Caron *et al.* [9], which showed that pixel-level semantic attention emerges from self-supervised vision transformers (ViTs) [18], we leverage attention maps from a frozen self-supervised ViT to generate pseudo-GT segmentation masks. We train an FS-CS learner on top of the frozen ViT using the generated pseudo-GT masks as pixel-level supervision. The FS-CS learner is a small transformer network that takes features from the ViT as input and is trained to predict the pseudo-GT masks. Our complete model thus learns to segment using its own intermediate byproduct, with no segmentation labeling cost, in a form of distillation [2, 28, 91, 94] between layers within a model.

We also explore a practical training setting where a limited number of training images have both image-level and pixel-level annotations, while the remaining images have only image-level labels. For this *mixed-supervised* setting, we propose to train an auxiliary mask enhancer that refines attention maps into pseudo-GT masks. The enhancer is supervised using the small number of available GT masks and is used to generate the final pseudo-GT masks for the training images that do not have GT masks.

Lastly, to effectively address this weakly-supervised FS-CS task, we propose a Classification-Segmentation Transformer (CST) architecture. CST takes as input ViT tokens for query and support images, computes correlations between them, and then predicts classification and segmentation outputs through separate task heads; the classification head is trained with class labels while the segmentation head is trained with either GT masks or pseudo-GT masks, depending on availability. Unlike prior work using ViTs that generate prediction for a single task using either the class token [18, 74] or image tokens [50, 65], CST uses each of them for each task prediction, which proves advantageous for both tasks. CST achieves moderate-to-large performance gains for FS-CS on all three supervision levels (image-level-, mixed-level, and pixel-level), when compared to prior state-of-the-art methods.

Our contributions are the following:

- i. We introduce a powerful baseline for few-shot classification and segmentation (FS-CS) using only image-level supervision, which leverages localized semantic information encoded in self-supervised ViTs [9].
- ii. We propose a learning setup for FS-CS with mixed supervision, and present an auxiliary mask enhancer to improve performance compared to image-level supervision only.
- iii. We design Classification-Segmentation Transformer, and a multi-task learning objective of classification and segmentation, which is beneficial for both tasks, and allows for flexibly tuning the degree of supervision.

2. Related work

Few-shot classification and segmentation. This paper tackles the recently proposed task of few-shot classification and segmentation (FS-CS) [33], which generalizes and combines the conventional few-shot classification (FS-C) [11, 24, 34, 63, 70] and few-shot segmentation (FS-S) [58, 71, 89] tasks. FS-C assumes that a query contains only *one* of the support classes, while FS-S assumes that the query always contains *all* of them. In contrast, FS-CS breaks the assumption that a query image contains one class (or all classes) from the support set. As a result, FS-CS models must predict i) the presence or absence of *each* support class in the query; *and* ii) pixel-level segmentation masks. Our work explores a more challenging weakly-supervised version of FS-CS with image-level annotations but few or no pixel-level annotations. Prior work on weakly-supervised FS-S used supervision such as language [40, 60] or bounding boxes [55]. Our work does not adopt any extra modalities and uses weaker (image-level) supervision than bounding boxes. Other applications of few-shot learning include few-shot object detection [92, 93], action recognition [72, 73], and metric learning [32].

Weakly-supervised segmentation. A common paradigm for weakly-supervised semantic segmentation uses class labels [1, 39] to generate pseudo-GT masks from class activation maps (CAM) [57, 95], followed by pseudo-supervised training [37, 96, 99]. CAM requires a classification backbone pre-trained from many examples on *seen* classes whereas, in our few-shot setting, we devise a pseudo-GT mask generation method that requires only a small number of examples, and can handle *unseen* classes. We also investigate the extent to which a pseudo-mask enhancement module, trained using a small amount of GT masks, can generate better pseudo-GT masks. While such “mixed-supervision” learning has been called semi-weakly-supervised segmentation [4, 38, 42, 75] and studied in the standard seen-class many-shot learning paradigm, our work is the first to propose mixed-supervision learning for few-shot semantic segmentation, to the best of our knowledge.

Self-supervised learning. Self-supervised learning [31] has proven highly effective for learning powerful representations for downstream tasks [7–10, 80, 97]. One popular approach [9, 25] trains a model to map an image and its augmented views to nearby coordinates in a representation space. This class-agnostic objective trains the model to capture view-invariant semantics that transcend designated image categories. In addition, DINO-trained ViTs [9] capture localized semantics in the learned token representations. In this work, we leverage these representations to create attention map-based pseudo-GT masks for weakly-supervised FS-CS, using a novel and efficient self-distillation approach. The representations have also ben-

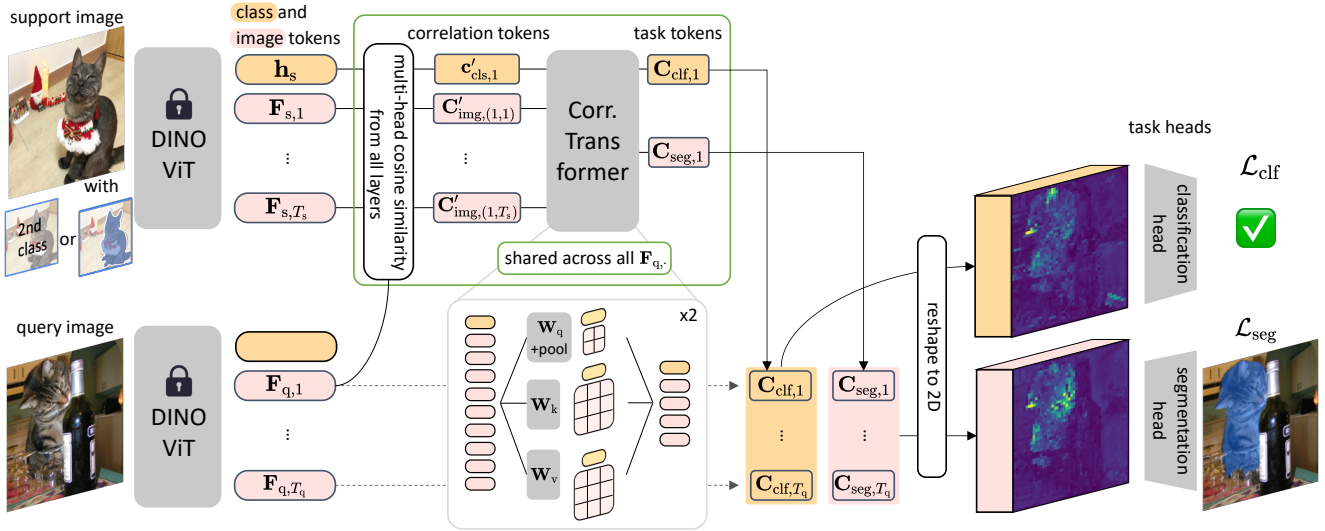


Figure 2. Classification-Segmentation Transformer. A query and a support image pass through a frozen DINO ViT to extract class and images tokens, which are cross-correlated to form correlation tokens. The correlation tokens are then fed to a correlation transformer and transform to task tokens. At the end, each classification and segmentation head takes each task token map and predicts each task.

edited several other dense prediction tasks, including unsupervised object discovery [62, 76] and unsupervised image/video segmentation [9, 86, 87, 98].

3. Weakly-supervised few-shot classification and segmentation

In this paper we tackle few-shot classification and segmentation (FS-CS) [33]. Given a small number K of labeled support images for each of N classes, an FS-CS model is trained to i) identify the presence of each class in a query image and ii) segment the associated pixels. This setting is conventionally called an N -way K -shot episode [53, 70]. We adopt a meta-learning approach [29, 63, 70] to train our model, where an annotated training dataset (with a class set disjoint from the test set classes) is used to generate a sequence of N -way K -shot episodes.¹ In our setting, annotations consist of class labels and segmentation masks. When ground-truth segmentation masks are available for all support and query images during training, we refer to this as pixel-level supervision. When only class labels are available we refer to this as image-level supervision. When ground-truth segmentation masks are available for only a small percentage of training images, we refer to this as mixed supervision.

In the remainder of this section, we first introduce our base model architecture, which we call Classification-Segmentation Transformer (CST), and then describe training paradigms for the different levels of supervision mentioned above.

¹Note that the terms, meta-training and meta-testing, are used interchangeably with training and testing in this paper.

3.1. Classification-Segmentation transformer

Our Classification-Segmentation Transformer (CST) architecture first computes the cross-correlation of DINO ViT [9, 18] tokens between support and query image pairs. It then generates two new sets of tokens which serve as input to the classification and segmentation prediction heads respectively. The objective function then combines the losses from each task head. The details of CST’s architecture are shown in Fig. 2 and described in the following paragraphs.

Correlation tokens. We use a frozen ViT architecture pre-trained with DINO as a backbone, which takes as input an image and produces $T = HW$ image (patch) tokens $\mathbf{F} \in \mathbb{R}^{C \times T}$ and one class token $\mathbf{h} \in \mathbb{R}^C$, where C denotes the token feature dimensionality. The terms “image” and “class” tokens are used in analogous fashion to previous work [9, 15, 18, 68]. The backbone first transforms a query image q and a support image s into two sets of tokens, $\{\mathbf{F}_q, \mathbf{h}_q\}$ and $\{\mathbf{F}_s, \mathbf{h}_s\}$. The query image tokens \mathbf{F}_q are cross-correlated, by cosine similarity,² with the support class token \mathbf{h}_s and the support image tokens \mathbf{F}_s , to generate a set of *correlation tokens* $\mathbf{c}_{\text{cls}} \in \mathbb{R}^{T_q}$ and $\mathbf{C}_{\text{img}} \in \mathbb{R}^{T_q \times T_s}$ respectively:

$$\mathbf{c}_{\text{cls}} = \bar{\mathbf{F}}_q^T \bar{\mathbf{h}}_s, \quad \mathbf{C}_{\text{img}} = \bar{\mathbf{F}}_q^T \bar{\mathbf{F}}_s, \quad (1)$$

where $\bar{\mathbf{F}}$ and $\bar{\mathbf{h}}$ denote l_2 -normalized tokens. We retain the semantic diversity of the M heads of the ViT backbone by computing M separate correlation tokens. In addition, prior work [26, 44] has demonstrated the value of incorporating multi-layer semantics. Similarly, we compute correlation

²The cosine similarity between two vectors is the uncentered Pearson correlation coefficient.

tokens for each of the L transformer layers of the ViT. In all, we compute ML different sets of correlation tokens and concatenate them along a new dimension to generate:

$$\mathbf{c}'_{\text{cls}} \in \mathbb{R}^{T_q \times ML}, \quad \mathbf{C}'_{\text{img}} \in \mathbb{R}^{T_q \times T_s \times ML}. \quad (2)$$

Correlation transformer. We first concatenate $\mathbf{c}'_{\text{cls},i} \in \mathbb{R}^{ML}$ and $\mathbf{C}'_{\text{img},i} \in \mathbb{R}^{T_s \times ML}$, where $i \in [1, \dots, T_q]$ is an index over the query image tokens, to produce $\mathbf{Z}_{0,i} \in \mathbb{R}^{(1+T_s) \times ML}$. Our correlation transformer is a 2-layer transformer [69] that takes $\mathbf{Z}_{0,i}$ and a segmentation mask label for the support image ($\mathbf{M}_{\text{gt}} \in \{0, 1\}^{T_s}$) as input, and returns two tokens - a classification task token and a segmentation task token - for query image token i . Each layer has the following form:

$$\mathbf{Z}'_{j,i} = \text{Norm}'_j(\text{MSA}_j(\mathbf{Z}_{j,i}, \mathbf{M}_{\text{gt}}) + \mathbf{Z}_{j,i}), \quad (3)$$

$$\mathbf{Z}_{j+1,i} = \text{Norm}_j(\text{MLP}_j(\mathbf{Z}'_{j,i}) + \mathbf{Z}'_{j,i}), \quad (4)$$

where MSA, Norm, and MLP correspond to a multi-head self-attention, a group normalization [78], and a linear layer, respectively, and j denotes the transformer layer index. Additionally, in each MSA layer, the generated query embeddings \mathbf{Q} are spatially pooled (using a 4×4 and 3×3 kernel in layers 1 and 2 respectively) by reshaping \mathbf{Q} to recover its H_s, W_s dimensions, applying spatial pooling, and flattening the output. This progressively reduces the number of image tokens, collapsing the input spatial size from $T_s = 12 \times 12$ to 1. This spatial pooling, illustrated in Fig. 2, reduces memory consumption, similarly to [14]. The output of the correlation transformer is then $\mathbf{Z}_{2,i} \in \mathbb{R}^{2 \times C'}$, where C' denotes the new token dimensionality after MLP projection. This process is broadcast over the correlation tokens $\mathbf{C}'_{\text{img},i}$ for each query image token i , resulting in $\mathbf{Z}_2 \in \mathbb{R}^{T_q \times 2 \times C'}$. We then split this tensor, along the second dimension, into classification and segmentation tokens $\mathbf{C}_{\text{clf}}, \mathbf{C}_{\text{seg}} \in \mathbb{R}^{T_q \times C'}$. Lastly, \mathbf{C}_{clf} and \mathbf{C}_{seg} are reshaped into spatial dimensions as $\mathbf{C}_{\text{clf}}, \mathbf{C}_{\text{seg}} \in \mathbb{R}^{H_q \times W_q \times C'}$ and then fed to their respective task heads, described next.

Classification and segmentation heads. The classification head consists of two 1×1 convolutions followed by global average pooling, and estimates the class presence probability \mathbf{y}_{clf} of the query image corresponding to the support image class. The segmentation head consists of two 3×3 convolutions followed by bilinear interpolation to the original query image size and predicts a segmentation map \mathbf{Y}_{seg} .

Learning objective. Given predictions from the task heads, the classification and segmentation losses are computed as the cross-entropy between the prediction and the ground-truth for each task:

$$\mathcal{L}_{\text{clf}} = -\mathbf{y}_{\text{gt}} \log \mathbf{y}_{\text{clf}}, \quad (5)$$

$$\mathcal{L}_{\text{seg}} = -\frac{1}{HW} \sum_{\mathbf{x} \in \{H\} \times \{W\}} \mathbf{Y}_{\text{gt}}(\mathbf{x}) \log \mathbf{Y}_{\text{seg}}(\mathbf{x}), \quad (6)$$




Figure 3. a) Image-level-supervision learners generate pseudo-GT masks by averaging multi-head cross-(self-) attention followed by thresholding (Sec. 3.2). **b) Mixed-supervision learners** first pre-train a mask enhancer with a few GT masks during training to enhance the pseudo-GT mask quality (Sec. 3.3).

where $\mathbf{y}_{\text{gt}} \in \{0, 1\}$ is the image-level ground-truth label and $\mathbf{Y}_{\text{gt}}(\mathbf{x}) \in \{0, 1\}$ is the ground-truth label for each pixel \mathbf{x} . The final learning loss combines the two losses so that the model learns both tasks jointly:

$$\mathcal{L} = \lambda \mathcal{L}_{\text{clf}} + \mathcal{L}_{\text{seg}}, \quad (7)$$

where λ is a balancing hyperparameter.

Inference. For N -way K -shot inference, K class logits and segmentation masks are generated using K support images, for each of N classes. These are then averaged, resulting in one estimate per class n , $\hat{\mathbf{y}}_{\text{clf}}^{(n)}$ and $\hat{\mathbf{Y}}_{\text{seg}}^{(n)}$. Multi-label classification $\tilde{\mathbf{y}}_{\text{clf}}^{(n)}$ is then obtained by thresholding each class response $\hat{\mathbf{y}}_{\text{clf}}^{(n)}$ with an arbitrarily chosen threshold of $\delta = 0.5$.

For segmentation, a pixel \mathbf{x} is labeled as the class with the highest response, or as background (*i.e.* the $N + 1$ -th class) if none of the N class responses exceeds δ . Formally:

$$\tilde{\mathbf{Y}}_{\text{seg}}(\mathbf{x}) = \begin{cases} N + 1 & \text{if } \forall \{\hat{\mathbf{Y}}_{\text{seg}}^{(n)}(\mathbf{x})\}_{n=1}^N \leq \delta, \\ \arg \max_{n \in \{1, \dots, N\}} \{\hat{\mathbf{Y}}_{\text{seg}}^{(n)}(\mathbf{x})\}_{n=1}^N. \end{cases} \quad (8)$$

3.2. Learning with image-level supervision

The previously-described learning objective assumes access to GT support and query segmentation mask labels, \mathbf{M}_{gt} and \mathbf{Y}_{gt} . To remove the dependence on such expensive annotations, in this section we propose a learning framework for FS-CS that relies only on image-level annotations, *i.e.*, ground-truth class labels. Using the same model architecture, we adopt a method akin to self-distillation [91, 94], that generates pseudo-GT masks, \mathbf{M}_{ps} and \mathbf{Y}_{ps} , and uses them in place of GT mask labels in equations 3 and 6.

Pseudo-GT mask generation. As empirically investigated in previous studies [3, 9], query-key attention maps can capture semantically salient foreground objects. Inspired by this, we generate pseudo-GT masks for query and support images on the fly by computing the cross- and self-attention from the last DINO ViT layer. We intercept the key and query embeddings $\mathbf{K}_s^m, \mathbf{Q}_s^m \in \mathbb{R}^{(1+T_s) \times C}$ and

$\mathbf{K}_q^m, \mathbf{Q}_q^m \in \mathbb{R}^{(1+T_q) \times C}$ for the support and query images respectively, for the m -th head from the last ViT layer. The support and query image pseudo-GT mask values \mathbf{M}_{ps,i_s} and \mathbf{Y}_{ps,i_q} (where i_s and i_q index the support and query image tokens respectively) are then computed as:

$$\begin{aligned} \mathbf{P}_s &\in \mathbb{R}^{M \times T_s}, & \mathbf{P}_{s,i_s}^m &= \overline{\mathbf{K}}_{s,i_s}^m \top \overline{\mathbf{Q}}_{s,c_s}^m, \\ \mathbf{P}_q &\in \mathbb{R}^{M \times T_q}, & \mathbf{P}_{q,i_q}^m &= \overline{\mathbf{K}}_{q,i_q}^m \top \overline{\mathbf{Q}}_{s,c_s}^m, \end{aligned} \quad (9)$$

$$\mathbf{M}'_{ps,i_s} = \frac{1}{M} \sum_{m=1}^M \mathbf{P}_{s,i_s}^m, \quad (10)$$

$$\mathbf{Y}'_{ps,i_q} = \begin{cases} \mathbf{0} & \text{if } \mathbf{y}_{\text{clf}} = \mathbf{0} \\ \frac{1}{M} \sum_{m=1}^M \mathbf{P}_{q,i_q}^m, & \end{cases}$$

$$\mathbf{M}_{ps} = \Pi(\xi(\mathbf{M}'_{ps}), \alpha); \quad \mathbf{Y}_{ps} = \Pi(\xi(\mathbf{Y}'_{ps}), \alpha), \quad (11)$$

with c_s denoting the index of the class token for the support image, $\xi(\cdot)$ performing spatial reshaping and bilinear interpolation to the input image size, and $\Pi(\cdot, \alpha)$ applying an element-wise binary threshold $\alpha = -0.1$ (set via hyperparameter search). The semantics captured in the support class token are matched with the localized semantics of the (query or support) image tokens, using (cross- or self-) attention. This localized semantic matching is then captured in the resulting pseudo-GT masks. Our pseudo-mask generation process brings negligible extra computation by leveraging internal byproducts from the ViT backbone.

This learning framework also relates to self-knowledge distillation [91, 94], in the sense that our model is supervised using its own (intermediate) features. We can think of the backbone as the teacher and the full CST model as the student. However, a key difference is that in our approach the student model learns additional parameters. In contrast, in strictly self-knowledge distillation approaches the student and teacher model architectures are identical.

3.3. Extension to mixed-supervision learning

As shown in Sec. 4, our pseudo-masks provide effective supervision for segmentation. However, there is unsurprisingly still a quality gap between these masks and the GT segmentation masks, which degrades segmentation performance. In this section, we propose a practical learning setting for FS-CS with *mixed supervision*, which reduces this degradation in segmentation performance, by using a small number of GT segmentation masks during training. Mixed-supervised FS-CS assumes that all training images contain image-level labels, and a small subset of them also contains pixel-level labels, which are relatively inexpensive compared to full supervision.

To obtain higher-quality pseudo-GT masks, we use the small number of images with GT masks to train a *mask enhancer*. We reshape \mathbf{P}_s and \mathbf{P}_q in Eq. (9) into spatial tensors



Figure 4. Examples of pseudo-GT & enhanced pseudo-GT masks.

and use them as inputs to the enhancer, which returns binary masks. During training the enhancer is supervised with the available GT masks. The mask enhancer is a network of three 3×3 convolutional layers followed by a Sigmoid. For mixed-supervision learning, we pre-train this mask enhancer and all pseudo-masks are obtained via the enhancer before being used as supervision during meta-training and meta-testing. Figure 3 compares the two proposed pseudo-GT mask generation processes under the image-level- and mixed-supervision learning settings.

4. Experiments

In this section we provide implementation details for our approach, and present quantitative and qualitative results to experimentally evaluate its performance. The attached supplementary material contains more analyses and details.

Architecture and training. To compare with prior work based on ResNet50 [33, 48, 67, 71], we use a ViT-small backbone [18] with 21M parameters that was pre-trained via self-supervision with DINO [9] on ImageNet 1K [54]. The size and training data are comparable with ResNet50 [27], which has 23M parameters and was also trained on ImageNet 1K, but using class labels as supervision. ViT-small consists of $L = 12$ transformer layers of $M = 6$ attention heads, with $C = 64$ -dimensional query-key-value embeddings. The backbone is frozen, as in previous work [67, 90], and the additional layers are trained using Adam [35] with a learning rate of 10^{-3} . Consistent with [33], CST is trained with 1-way 1-shot episodes and is used for any N -way K -shot inference. We set the task token dimensionality $C' = 128$, and the loss balancing hyperparameter $\lambda = 0.1$.

Datasets. We report results on two standard FS-CS benchmarks [33], Pascal-5ⁱ [20, 58] and COCO-20ⁱ [45, 49]. Both datasets use 4-fold validation, where the folds $i \in \{0, 1, 2, 3\}$ are created by splitting the 20 (resp. 80) classes into 4 mutually disjoint class sets, with 10K (resp. 100K) images per fold. For training and validation, support and query image pairs are sampled on the fly, as done in [33]. For each training episode, N support classes are chosen randomly from

method	ps-mask	1-way 1-shot					2-way 1-shot					learn. params.										
		classification 0/1 exact ratio (%)					segmentation mIoU (%)															
		5 ⁰	5 ¹	5 ²	5 ³	avg.	5 ⁰	5 ¹	5 ²	5 ³	avg.											
HSNet [48]	✗	84.5	84.8	60.8	85.3	78.9	20.0	23.5	16.2	16.6	19.1	70.8	67.0	36.5	70.1	61.1	11.0	23.4	15.4	17.0	17.7	2.6M
ASNet [33]	✗	80.2	84.0	66.2	82.7	78.3	11.7	21.1	13.4	16.2	15.6	67.6	70.2	44.5	69.3	62.9	10.3	20.7	12.6	15.9	14.9	1.3M
CST (image)	✗	83.9	80.7	65.9	82.7	78.3	11.9	20.7	15.0	16.3	16.0	72.6	65.3	43.2	69.5	62.6	11.9	20.5	15.0	16.6	15.7	0.4M
DINO [9]	◇	-	-	-	-	-	20.0	23.4	16.2	16.6	19.1	-	-	-	-	-	12.8	12.9	10.0	9.9	11.4	0
CST	✓	84.0	82.2	70.8	82.6	79.9	35.8	38.9	28.9	29.2	33.2	74.3	67.2	49.0	67.8	64.6	35.7	36.0	26.8	29.1	31.9	0.4M

Table 1. Comparing model performance on FS-CS when trained with image-level supervision on Pascal-5ⁱ [58]. “ps-mask” refers to using pseudo-GT masks during training. Our CST model compares favourably to prior works on classification, and outperforms them on segmentation. Results for DINO [9] were generated without any training on Pascal-5ⁱ (denoted by ◇). See section 4.1 for details.

method	ps-mask	1-way 1-shot		2-way 1-shot	
		cls. ER	seg. mIoU	clf. ER	seg. mIoU
CST	✗	74.5	10.3	62.3	9.8
DINO [9]	◇	-	12.1	-	7.4
CST	✓	78.2	19.6	62.4	18.3

Table 2. Comparing model performance on FS-CS when trained with image-level supervision on COCO-20ⁱ [49], averaged over 4 folds. See section 4.1 for details.

the split and K images are chosen randomly from each of the N target classes to construct an N -way K -shot support set. When a support image contains multiple classes, only one of them is set as its image class and the rest are set as background. While this introduces noise into support sets containing multi-class images, we find that it performs well in practice. A query image is also chosen randomly and may contain between 0 to N of the target classes. The image-level label for the query is then constructed as an N -dimensional binary vector containing 1 if the query contains the n -th class, and 0 otherwise. Similarly, its pixel-level label is defined as an $(N+1)$ -way segmentation mask, including background. For mixed-supervision learning, we sort the images in a training split by their image names and use pixel-level labels for the first *e.g.* 5% of images, and only image-level labels for the remainder. This effectively randomizes the subset of images with full-supervision.

Evaluation and metrics. We follow the established evaluation protocol for each FS-CS benchmark [33, 49, 58]. To evaluate multi-label classification, we use the 0/1 exact match ratio [19], where an image is considered to be correctly classified if the binary label for each of the N classes is correct. For semantic segmentation, we use mean IoU, $mIoU = \frac{1}{R} \sum_R IoU_r$ [58, 71], where IoU_r is the IoU of all images labeled with the r -th-class. The best model is chosen based on the highest validation mIoU averaged with 1000 random validation episodes.

4.1. Results with image-level supervision

Tables 1 and 2 compare the performance of our image-level-supervised CST with baseline methods. The only prior work that addresses this setting [33] generates a seg-

% GT training set masks →	0%	2.5%	5%	10%	100%
(a): pseudo-mask baseline	33.2	41.4	43.5	44.5	-
(b): (a) + mask enhancer	-	43.2	44.1	45.2	-
(c): (b) + GT test set \mathcal{S} masks	-	47.5	48.4	50.1	55.5

Table 3. Comparing segmentation mIoU of CST on FS-CS with mixed supervision on Pascal-5ⁱ, for different % of training set ground-truth masks, averaged over 4 folds. In (c), ground-truth masks were used for all support (\mathcal{S}) images during testing.

mentation mask that is then globally average-pooled. The cross-entropy between the average-pooled prediction and the ground-truth image label is used as the objective function. We adopt this approach for [48], which did not address this setting, to construct another baseline, and we also evaluate CST in this setting, which we denote with ✗. The result for DINO is computed by evaluating segmentation masks that were generated as in [9] (denoted by ◇), without using any training data from the target datasets. Our CST model, trained with pseudo-masks via self-distillation (denoted by ✓), performs comparably to prior work on classification and, because of the superior supervision provided by our pseudo-masks during training, outperforms these models on segmentation. In particular, CST achieves an over 100% improvement over ASNet [33] in both the 1-way 1-shot and 2-way 1-shot settings. We also note that training CST with pseudo-GT masks (✓) also improves classification performance by 1.6% (1-way 1-shot) and 2% (2-way 1-shot) on Pascal-5ⁱ when compared to training without (✗).

4.2. Results with mixed supervision

To validate the effectiveness of the mixed-supervision setting, we vary the percentage of pixel-level annotated images used in training and present the results in Table 3. We observe that using just 5% of GT masks is remarkably effective at improving performance, when compared to using 0% (43.5% vs. 33.2%). The mask enhancer brings much improvement with fewer ground-truth segmentation masks. The effect of the mask enhancer is visually demonstrated in Fig. 4. We observe that the mask enhancer reduces foreground noise and often filters out objects irrelevant to a support image. Note that mixed-supervision learning assumes access to pixel-level annotations only during training. Dur-

method	1-way 1-shot										2-way 1-shot										learn. params.
	classification 0/1 exact ratio (%)					segmentation mIoU (%)					classification 0/1 exact ratio (%)					segmentation mIoU (%)					
	5 ⁰	5 ¹	5 ²	5 ³	avg.	5 ⁰	5 ¹	5 ²	5 ³	avg.	5 ⁰	5 ¹	5 ²	5 ³	avg.	5 ⁰	5 ¹	5 ²	5 ³	avg.	
PANet [71]	69.9	67.7	68.8	69.4	69.0	32.8	45.8	31.0	35.1	36.2	56.2	47.5	44.6	55.4	50.9	33.3	46.0	31.2	38.4	37.2	23.6M
PFENet [67]	69.8	82.4	68.1	77.9	74.6	38.3	54.7	35.1	43.8	43.0	22.5	61.7	40.3	39.5	41.0	31.1	47.3	30.8	32.2	35.3	31.5M
HSNet [48]	86.6	84.8	76.9	86.3	83.7	49.1	59.7	41.0	49.0	49.7	68.0	73.2	57.0	70.9	67.3	42.4	53.7	34.0	43.9	43.5	2.6M
ASNet [33]	84.9	89.6	79.0	86.2	84.9	51.7	61.5	43.3	52.8	52.3	68.5	76.2	58.6	70.0	68.3	48.5	58.3	36.3	48.3	47.8	1.3M
CST	86.9	88.0	81.5	86.5	85.7	55.6	61.6	47.7	56.9	55.5	70.3	74.9	62.0	74.4	70.4	56.9	61.0	48.0	57.0	55.7	0.4M

Table 4. Comparing FS-CS models trained with pixel-level supervision on Pascal-5ⁱ.

method	1-way 1-shot		2-way 1-shot	
	clf. ER	seg. mIoU	clf. ER	seg. mIoU
PANet [71]	66.7	25.2	48.5	23.6
PFENet [67]	71.4	31.9	36.5	22.6
HSNet [48]	77.0	34.3	62.5	29.5
ASNet [33]	78.6	35.8	63.1	31.6
CST	80.8	38.3	64.0	36.2

Table 5. Comparing performance on FS-CS with pixel-level supervision on COCO-20ⁱ.




Figure 5. Comparing performance on FS-CS with pixel-level supervision on N -way 1-shot Pascal-5ⁱ.

ing testing, support images use pseudo-masks without (Table 3 (a)) or with (Table 3 (b)) mask enhancement. In (c), we use one-shot support GT mask during testing, but still use reduced amounts of GT masks in training. In this setting, the performance is only 5.4%-8.0% lower than the fully-supervised model, indicating that our approach allows for significant pixel-level annotation efficiency during training.

4.3. Results with full, *i.e.* pixel-level, supervision

Tables 4 and 5, as well as Fig. 5, compare our model’s performance on FS-CS with state-of-the-art methods whose results, included here, were initially reported in [33]. CST achieves significant improvements over previous work, which we attribute to (i) our correlation tokens that leverage the localized semantics encoded in our self-supervised ViT backbone model; and (ii) the two task heads of our transformer (see ablations in Sec. 4.4). Figure 6 visually demonstrates qualitative examples of final segmentation predictions with increasing levels of supervision, from image-level to pixel-level. Image-level supervised predictions roughly segment out objects that correspond with the support classes, and the predictions become less noisy as more pixel-level supervision is used.

labels →	class-sup.	self-sup.	self-sup.	self-sup.
archi. →	ViT [18]	CNN [27]	ViT [9]*	ViT [9]
seg. mIoU (%)	6.9	20.7	15.6	33.2

Table 6. FS-CS model comparison trained with image-level labels on 1-way 1-shot Pascal-5ⁱ by varying pseudo-mask generation models. Different columns denote model architectures for pseudo-labeling and their pre-training supervision levels. The method * uses the final layer tokens from [9] for generating pseudo-labels.

clf. & seg. prediction heads	superv. level	1-way 1-shot		2-way 1-shot	
		clf. ER	seg. mIoU	clf. ER	seg. mIoU
coupled	image	78.5	32.9	59.5	31.9
decoupled (CST)	image	79.9	33.2	64.5	31.9
coupled	pixel	84.3	54.2	67.0	53.0
decoupled (CST)	pixel	85.7	55.5	70.5	55.7

Table 7. Ablation study of coupling or decoupling classification and segmentation prediction heads for FS-CS on Pascal-5ⁱ.

4.4. Ablation study

Pseudo-labeling with different models. We explore other learning-free pseudo-labeling methods by using alternative pre-trained feature extractors to generate pseudo-GT masks. Rather than self-distillation, where we train CST with pseudo-GT masks generated from its own self-supervised DINO ViT backbone, in these experiments we train CST with pseudo-GT masks generated from class-supervised ViT [18], self-supervised DINO ResNet50 [27], or the output layer tokens from the self-supervised DINO ViT [9]. As ResNets do not have class tokens, we apply global-average pooling on the support feature map from the last ResNet layer and use the pooled feature as a class token for generating the pseudo-GT masks. As seen in Table 6, pseudo-GT masks generated from alternative feature extractors do not perform as well as the intermediate attention maps from DINO. We qualitatively observed that the alternatives’ less accurate pseudo-masks lead to lower segmentation accuracy when used in pseudo-supervised learning: the self-supervised ResNet50 produces pseudo-masks with indistinct object boundaries, while the class-supervised ViT produces noisy patches throughout the pseudo-mask (as was




Figure 6. 2-way 1-shot segmentation examples from CST trained with image-level, mixed, and pixel-level supervision. Note that image-level-supervised and mixed-supervised models do not have access to ground-truth support masks. In both cases, the model is still effective.




Figure 7. An example of channel-averaged C_{clf} , C_{seg} . Each task token tends to capture different aspects, *e.g.*, clearer edges of C_{seg} .

method	1-way 1-shot		1-way 5-shot		# learn. params.
	mIoU	FBIoU	mIoU	FBIoU	
RPM [83]	30.6	-	35.5	-	38.6M
RePRI [6]	34.0	-	42.1	-	-
SSP [21]	37.4	-	44.1	-	8.7M
MMNet [79]	37.5	-	38.2	-	10.4M
MLC [84]	33.9	-	40.6	-	8.7M
NTRENet [47]	39.3	68.5	40.3	69.2	-
CMN [81]	39.3	61.7	43.1	63.3	-
HSNet [48]	39.2	68.2	46.9	70.7	2.6M
DACM [82]	40.6	68.9	48.1	71.6	-
ASNet [33]	42.2	68.8	47.9	71.6	1.3M
CST	44.0	70.3	48.7	73.7	0.4M

Table 8. Comparing FS-S models trained with pixel-level supervision on COCO-20^f.

similarly observed in [9]).

Effect of decoupled task heads. Table 7 compares the impact of coupling and decoupling the classification and segmentation task heads. The coupled architecture has a single segmentation task head, and the image is classified with the foreground class if any pixels are predicted as foreground, and is otherwise classified as the background class. Having separate heads for each task benefits both tasks in both the image-level and pixel-level supervision settings. A similar observation was made for object detection [64] when decoupling localization from classification. Figure 7 visualizes an example of channel-averaged C_{clf} and C_{seg} tensors. We observe that these two intermediate representations tend to highlight different areas; C_{clf} has large responses in highly-localized class-specific image regions, while C_{seg} better distinguishes the overall object support and boundaries.

Comparison on FS-S. In Table 8, CST is also evaluated in the conventional FS-S setting [58, 71], *i.e.*, binary segmentation using mask supervision, and compared with previous works that use different ResNet50-based architectures. FS-S is equivalent to FS-CS with the precondition that a query and a support always share a class label; classification is not required. We train CST on this setting and find that it achieves state-of-the-art segmentation performance. It is also worth noting that it achieves this with fewer learned parameters than others, which is also the case for the comparisons in Table 1 and Table 4. Our small number of learned parameters may help to avoid overfitting.

5. Conclusion

In this work, we propose an approach to generating effective pseudo-GT masks for few-shot classification and segmentation (FS-CS) using the localized semantics encoded in the tokens of self-supervised vision transformers. We show that using these pseudo-GT masks, sometimes combined with a small number of GT masks, we can train effective FS-CS models. Our work is the first to explore this practical weakly-supervised FS-CS setting, which drastically reduces the amount of expensive pixel-level annotations needed for training and deployment. Extensive experiments show that our approach achieves significant performance gains when compared to state-of-the-art methods in a variety of supervision settings, on two FS-CS benchmarks.

Acknowledgements. This work was partly supported by the IITP grants (2022-0-00959: Few-Shot Learning of Causal Inference in Vision and Language, 2022-0-00113: Sustainable Collaborative Multi-modal Lifelong Learning) funded by Ministry of Science and ICT, Korea. PK was in part funded by CSIRO’s Machine Learning and Artificial Intelligence Future Science Platform (MLAI FSP) Spatiotemporal Activity.

References

- [1] Jiwoon Ahn, Sunghyun Cho, and Suha Kwak. Weakly supervised learning of instance segmentation with inter-pixel relations. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [2] Zeyuan Allen-Zhu and Yuanzhi Li. Towards understanding ensemble, knowledge distillation and self-distillation in deep learning. *arXiv preprint arXiv:2012.09816*, 2020. 2
- [3] Shir Amir, Yossi Gandelsman, Shai Bagon, and Tali Dekel. Deep vit features as dense visual descriptors. *arXiv preprint arXiv:2112.05814*, 2021. 4
- [4] Wonho Bae, Junhyug Noh, Milad Jalali Asadabadi, and Danica J Sutherland. One weird trick to improve your semi-weakly supervised semantic segmentation model. *Proc. International Joint Conference on Artificial Intelligence (IJCAI)*, 2022. 2
- [5] Hakan Bilen and Andrea Vedaldi. Weakly supervised deep detection networks. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 1
- [6] Malik Boudiaf, Hoel Kervadec, Ziko Imtiaz Masud, Pablo Piantanida, Ismail Ben Ayed, and Jose Dolz. Few-shot segmentation without meta-learning: A good transductive inference is all you need? In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 8
- [7] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In *Proc. European Conference on Computer Vision (ECCV)*, 2018. 2
- [8] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 2
- [9] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2021. 2, 3, 4, 5, 6, 7, 8
- [10] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *Proc. International Conference on Machine Learning (ICML)*, 2020. 2
- [11] Wei-Yu Chen, Yen-Cheng Liu, Zsolt Kira, Yu-Chiang Wang, and Jia-Bin Huang. A closer look at few-shot classification. In *International Conference on Learning Representations (ICLR)*, 2019. 1, 2
- [12] Yinbo Chen, Zhuang Liu, Huijuan Xu, Trevor Darrell, and Xiaolong Wang. Meta-baseline: Exploring simple meta-learning for few-shot learning. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2021. 1
- [13] Jifeng Dai, Kaiming He, and Jian Sun. Boxesup: Exploiting bounding boxes to supervise convolutional networks for semantic segmentation. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2015. 1
- [14] Zihang Dai, Guokun Lai, Yiming Yang, and Quoc Le. Funnel-transformer: Filtering out sequential redundancy for efficient language processing. *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 4
- [15] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 3
- [16] Guneet Singh Dhillon, Pratik Chaudhari, Avinash Ravichandran, and Stefano Soatto. A baseline for few-shot image classification. In *International Conference on Learning Representations*, 2019. 1
- [17] Nanqing Dong and Eric P Xing. Few-shot semantic segmentation with prototype learning. In *Proc. British Machine Vision Conference (BMVC)*, 2018. 1
- [18] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *Proc. International Conference on Learning Representations (ICLR)*, 2021. 2, 3, 5, 7
- [19] Thibaut Durand, Nazanin Mehrasa, and Greg Mori. Learning a deep convnet for multi-label classification with partial labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 647–657, 2019. 6
- [20] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International Journal of Computer Vision (IJCV)*, 2010. 5
- [21] Qi Fan, Wenjie Pei, Yu-Wing Tai, and Chi-Keung Tang. Self-support few-shot semantic segmentation. In *Proc. European Conference on Computer Vision (ECCV)*, 2022. 8
- [22] Li Fei-Fei, Rob Fergus, and Pietro Perona. One-shot learning of object categories. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2006. 1
- [23] Michael Fink. Object classification from a single example utilizing class relevance metrics. *Advances in Neural Information Processing Systems (NeurIPS)*, 2005. 1
- [24] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proc. International Conference on Machine Learning (ICML)*, 2017. 1, 2
- [25] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent—a new approach to self-supervised learning. *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 2
- [26] Bharath Hariharan, Pablo Arbeláez, Ross Girshick, and Jitendra Malik. Hypercolumns for object segmentation and fine-grained localization. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 3
- [27] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 5, 7

- [28] Geoffrey Hinton, Oriol Vinyals, Jeff Dean, et al. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 2
- [29] Sepp Hochreiter, A Steven Younger, and Peter R Conwell. Learning to learn using gradient descent. In *Proc. International Conference on Artificial Neural Networks (ICANN)*, 2001. 3
- [30] Cheng-Chun Hsu, Kuang-Jui Hsu, Chung-Chi Tsai, Yen-Yu Lin, and Yung-Yu Chuang. Weakly supervised instance segmentation using the bounding box tightness prior. *Advances in Neural Information Processing Systems (NeurIPS)*, 2019. 1
- [31] Longlong Jing and Yingli Tian. Self-supervised visual feature learning with deep neural networks: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2020. 2
- [32] Deunsol Jung, Dahyun Kang, Suha Kwak, and Minsu Cho. Few-shot metric learning: Online adaptation of embedding for retrieval. 2022. 2
- [33] Dahyun Kang and Minsu Cho. Integrative few-shot learning for classification and segmentation. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1, 2, 3, 5, 6, 7, 8
- [34] Dahyun Kang, Heeseung Kwon, Juhong Min, and Minsu Cho. Relational embedding for few-shot classification. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2021. 2
- [35] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *Proc. International Conference on Learning Representations (ICLR)*, 2015. 5
- [36] Gregory Koch, Richard Zemel, and Ruslan Salakhutdinov. Siamese neural networks for one-shot image recognition. In *International Conference on Machine Learning (ICML) Workshop on Deep Learning*, 2015. 1
- [37] Dong-Hyun Lee et al. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, volume 3, page 896, 2013. 2
- [38] Jungbeom Lee, Eunji Kim, Sungmin Lee, Jangho Lee, and Sungroh Yoon. Ficklenet: Weakly and semi-supervised semantic image segmentation using stochastic inference. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [39] Jungbeom Lee, Eunji Kim, and Sungroh Yoon. Anti-adversarially manipulated attributions for weakly and semi-supervised semantic segmentation. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [40] Yuan-Hao Lee, Fu-En Yang, and Yu-Chiang Frank Wang. A pixel-level meta-learner for weakly supervised few-shot semantic segmentation. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2022. 2
- [41] Dong Li, Jia-Bin Huang, Yali Li, Shengjin Wang, and Ming-Hsuan Yang. Weakly supervised object localization with progressive domain adaptation. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 1
- [42] Kunpeng Li, Ziyang Wu, Kuan-Chuan Peng, Jan Ernst, and Yun Fu. Tell me where to look: Guided attention inference network. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [43] Di Lin, Jifeng Dai, Jiaya Jia, Kaiming He, and Jian Sun. Scribblesup: Scribble-supervised convolutional networks for semantic segmentation. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3159–3167, 2016. 1
- [44] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017. 3
- [45] 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 *Proc. European Conference on Computer Vision (ECCV)*, 2014. 5
- [46] Bin Liu, Yue Cao, Yutong Lin, Qi Li, Zheng Zhang, Mingsheng Long, and Han Hu. Negative margin matters: Understanding margin in few-shot classification. In *Proc. European Conference on Computer Vision (ECCV)*, 2020. 1
- [47] Yuanwei Liu, Nian Liu, Qinglong Cao, Xiwen Yao, Junwei Han, and Ling Shao. Learning non-target knowledge for few-shot semantic segmentation. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 8
- [48] Juhong Min, Dahyun Kang, and Minsu Cho. Hypercorrelation squeeze for few-shot segmentation. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2021. 5, 6, 7, 8
- [49] Khoi Nguyen and Sinisa Todorovic. Feature weighting and boosting for few-shot segmentation. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2019. 5, 6
- [50] Zizheng Pan, Bohan Zhuang, Jing Liu, Haoyu He, and Jianfei Cai. Scalable vision transformers with hierarchical pooling. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2021. 2
- [51] Deepak Pathak, Evan Shelhamer, Jonathan Long, and Trevor Darrell. Fully convolutional multi-class multiple instance learning. *Proc. International Conference on Learning Representations (ICLR)*, 2015. 1
- [52] Xiaojuan Qi, Zhengzhe Liu, Jianping Shi, Hengshuang Zhao, and Jiaya Jia. Augmented feedback in semantic segmentation under image level supervision. In *Proc. European Conference on Computer Vision (ECCV)*, 2016. 1
- [53] Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. In *Proc. International Conference on Learning Representations (ICLR)*, 2017. 3
- [54] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. 5
- [55] Oindrila Saha, Zezhou Cheng, and Subhansu Maji. Improving few-shot part segmentation using coarse supervision. In *Proc. European Conference on Computer Vision (ECCV)*, 2022. 2

- [56] Jürgen Schmidhuber. *Evolutionary principles in self-referential learning, or on learning how to learn: the meta-meta-... hook*. PhD thesis, Technische Universität München, 1987. 1
- [57] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2017. 2
- [58] Amirreza Shaban, Shray Bansal, Zhen Liu, Irfan Essa, and Byron Boots. One-shot learning for semantic segmentation. In *Proc. British Machine Vision Conference (BMVC)*, 2017. 1, 2, 5, 6, 8
- [59] Wataru Shimoda and Keiji Yanai. Distinct class-specific saliency maps for weakly supervised semantic segmentation. In *Proc. European Conference on Computer Vision (ECCV)*, 2016. 1
- [60] Mennatullah Siam, Naren Doraiswamy, Boris N Oreshkin, Hengshuai Yao, and Martin Jagersand. Weakly supervised few-shot object segmentation using co-attention with visual and semantic embeddings. In *Proc. International Joint Conference on Artificial Intelligence (IJCAI)*, 2020. 2
- [61] Mennatullah Siam, Boris N. Oreshkin, and Martin Jagersand. Amp: Adaptive masked proxies for few-shot segmentation. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2019. 1
- [62] Oriane Siméoni, Gilles Puy, Huy V Vo, Simon Roburin, Spyros Gidaris, Andrei Bursuc, Patrick Pérez, Renaud Marlet, and Jean Ponce. Localizing objects with self-supervised transformers and no labels. In *Proc. British Machine Vision Conference (BMVC)*, 2021. 3
- [63] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017. 1, 2, 3
- [64] Guanglu Song, Yu Liu, and Xiaogang Wang. Revisiting the sibling head in object detector. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 8
- [65] Robin Strudel, Ricardo Garcia, Ivan Laptev, and Cordelia Schmid. Segmenter: Transformer for semantic segmentation. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2021. 2
- [66] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip HS Torr, and Timothy M Hospedales. Learning to compare: Relation network for few-shot learning. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1
- [67] Zhuotao Tian, Hengshuang Zhao, Michelle Shu, Zhicheng Yang, Ruiyu Li, and Jiaya Jia. Prior guided feature enrichment network for few-shot segmentation. In *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2020. 5, 7
- [68] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Herve Jégou. Training data-efficient image transformers & distillation through attention. In *Proc. International Conference on Machine Learning (ICML)*. PMLR, 2021. 3
- [69] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017. 4
- [70] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, and Daan Wierstra. Matching networks for one shot learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2016. 1, 2, 3
- [71] Kaixin Wang, Jun Hao Liew, Yingtian Zou, Daquan Zhou, and Jiashi Feng. Panet: Few-shot image semantic segmentation with prototype alignment. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2019. 1, 2, 5, 6, 7, 8
- [72] Lei Wang and Piotr Koniusz. Temporal-viewpoint transportation plan for skeletal few-shot action recognition. 2022. 2
- [73] Lei Wang and Piotr Koniusz. Uncertainty-dtw for time series and sequences. In *Proc. European Conference on Computer Vision (ECCV)*, 2022. 2
- [74] 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 *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2021. 2
- [75] Yuhang Wang, Jing Liu, Yong Li, and Hanqing Lu. Semi-and weakly-supervised semantic segmentation with deep convolutional neural networks. In *Proc. ACM Multimedia Conference (ACMMM)*, 2015. 2
- [76] Yangtao Wang, Xi Shen, Shell Xu Hu, Yuan Yuan, James L Crowley, and Dominique Vaufreydaz. Self-supervised transformers for unsupervised object discovery using normalized cut. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 3
- [77] Yaqing Wang, Quanming Yao, James T Kwok, and Lionel M Ni. Generalizing from a few examples: A survey on few-shot learning. *ACM computing surveys (csur)*, 2020. 1
- [78] Yuxin Wu and Kaiming He. Group normalization. In *Proc. European Conference on Computer Vision (ECCV)*, 2018. 4
- [79] Zhonghua Wu, Xiangxi Shi, Guosheng Lin, and Jianfei Cai. Learning meta-class memory for few-shot semantic segmentation. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2021. 8
- [80] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [81] Guo-Sen Xie, Huan Xiong, Jie Liu, Yazhou Yao, and Ling Shao. Few-shot semantic segmentation with cyclic memory network. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2021. 8
- [82] Zhitong Xiong, Haopeng Li, and Xiao Xiang Zhu. Doubly deformable aggregation of covariance matrices for few-shot segmentation. In *Proc. European Conference on Computer Vision (ECCV)*, 2022. 8
- [83] Boyu Yang, Chang Liu, Bohao Li, Jianbin Jiao, and Ye Qixiang. Prototype mixture models for few-shot semantic segmentation. In *Proc. European Conference on Computer Vision (ECCV)*, 2020. 8

- [84] Lihe Yang, Wei Zhuo, Lei Qi, Yinghuan Shi, and Yang Gao. Mining latent classes for few-shot segmentation. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2021. 8
- [85] Han-Jia Ye, Hexiang Hu, De-Chuan Zhan, and Fei Sha. Few-shot learning via embedding adaptation with set-to-set functions. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 1
- [86] Zhaoyuan Yin, Pichao Wang, Fan Wang, Xianzhe Xu, Hanling Zhang, Hao Li, and Rong Jin. Transfgu: a top-down approach to fine-grained unsupervised semantic segmentation. In *Proc. European Conference on Computer Vision (ECCV)*, 2022. 3
- [87] Andrii Zadaianchuk, Matthaeus Kleindessner, Yi Zhu, Francesco Locatello, and Thomas Brox. Unsupervised semantic segmentation with self-supervised object-centric representations. *arXiv preprint arXiv:2207.05027*, 2022. 3
- [88] Chi Zhang, Guosheng Lin, Fayao Liu, Jiushuang Guo, Qingyao Wu, and Rui Yao. Pyramid graph networks with connection attentions for region-based one-shot semantic segmentation. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2019. 1
- [89] Chi Zhang, Guosheng Lin, Fayao Liu, Rui Yao, and Chunhua Shen. Canet: Class-agnostic segmentation networks with iterative refinement and attentive few-shot learning. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1, 2
- [90] Gengwei Zhang, Guoliang Kang, Yi Yang, and Yunchao Wei. Few-shot segmentation via cycle-consistent transformer. *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. 5
- [91] Linfeng Zhang, Chenglong Bao, and Kaisheng Ma. Self-distillation: Towards efficient and compact neural networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2021. 2, 4, 5
- [92] Shan Zhang, Naila Murray, Lei Wang, and Piotr Koniusz. Time-reversed diffusion tensor transformer: A new tenet of few-shot object detection. In *Proc. European Conference on Computer Vision (ECCV)*, 2022. 2
- [93] Shan Zhang, Lei Wang, Naila Murray, and Piotr Koniusz. Kernelized few-shot object detection with efficient integral aggregation. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [94] Zhilu Zhang and Mert R Sabuncu. Self-distillation as instance-specific label smoothing. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 2, 4, 5
- [95] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2
- [96] Yi Zhu, Zhongyue Zhang, Chongruo Wu, Zhi Zhang, Tong He, Hang Zhang, R Manmatha, Mu Li, and Alexander J Smola. Improving semantic segmentation via efficient self-training. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2021. 2
- [97] Chengxu Zhuang, Alex Lin Zhai, and Daniel Yamins. Local aggregation for unsupervised learning of visual embeddings. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2019. 2
- [98] Adrian Ziegler and Yuki M Asano. Self-supervised learning of object parts for semantic segmentation. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 3
- [99] Yuliang Zou, Zizhao Zhang, Han Zhang, Chun-Liang Li, Xiao Bian, Jia-Bin Huang, and Tomas Pfister. Pseudoseg: Designing pseudo labels for semantic segmentation. *Proc. International Conference on Learning Representations (ICLR)*, 2021. 2