Semantic Segmentation with Active Semi-Supervised Representation Learning

Aneesh Rangnekar¹ aneesh.rangnekar@mail.rit.edu Christopher Kanan² ckanan@cs.rochester.edu

Matthew Hoffman¹ mjhsma@rit.edu

- ¹ Rochester Institute of Technology Rochester, NY, USA
- ² University of Rochester Rochester, NY, USA

Abstract

Obtaining human per-pixel labels for semantic segmentation is incredibly laborious, often making labeled dataset construction prohibitively expensive. Here, we endeavor to overcome this problem with a novel algorithm that combines semi-supervised and active learning, resulting in the ability to train an effective semantic segmentation algorithm with significantly lesser labeled data. To do this, we extend the prior state-of-the-art S4AL algorithm by replacing its mean teacher approach for semi-supervised learning with a self-training approach that improves learning with noisy labels. We further boost the neural network's ability to query useful data by adding a contrastive learning head, which leads to better understanding of the objects in the scene, and hence, better queries for active learning. We evaluate our method on CamVid and CityScapes datasets, the de-facto standards for active learning for semantic segmentation. We achieve more than 95% of the network's performance on CamVid and CityScapes datasets, utilizing only 12.1% and 15.1% of the labeled data, respectively. We also benchmark our method across existing stand-alone semi-supervised learning methods on the CityScapes dataset and achieve superior performance without any bells or whistles.

1 Introduction

Getting labels for supervised learning problems is challenging, especially for semantic segmentation where these labels are needed on a per-pixel level. The most widely used methods for reducing the need for labels are semi-supervised learning (SSL) and pool-based active learning (AL). Semi-supervised learning proposes strategies to use the unlabeled dataset alongside labeled samples, typically by maintaining an exponential moving average of the network to predict pseudo labels [13, 24, 54], and pool-based active learning queries the most informative samples within the unlabeled data pool, in terms of pixels, regions, or entire images, to add to the labeling pool via a predetermined scoring mechanism [14, 154], [25], [26]]. Recently, both EquAL [155] and S4AL [255] have pioneered combining active and semisupervised learning for semantic segmentation, resulting in significant gains compared to using either in isolation. However, these methods are complicated - the former relies on enforcing consistency in the predictions without accounting for their correctness, and the latter



Figure 1: An example from the CityScapes dataset for Active learning: (a) Shows the image, (b) The ground truth labels, (c) The predictions from a network trained only with crossentropy loss, and (d) The predictions from a network trained with contrastive representation learning. The red box highlights the improved predictions in for the train and the nearby vegetation over standard training with cross-entropy loss.

relies on complex data augmentation schemes for convergence. Here we propose S4AL+, which greatly simplifies the S4AL method by replacing its mean-teacher framework with self-training, thereby eliminating the use of complex data augmentations [53, 50].

Most methods for active learning rely on cross-entropy loss for training their networks [13, 29, 59]. While cross-entropy is a widely used loss function in semantic segmentation, it operates on a individual pixels of the semantic map and fails to take into consideration possible cues of similarity and differences in areas throughout the entire image or set of images to strengthen its learning further (Fig. 1). The computer vision community has developed several algorithms to address this problem, notably Conditional Random Fields (CRF) [5], Affinity Nets [23], Region Mutual Information Loss (RMI) [13] and Contrastive Learning [22, 12]. We hypothesize that boosting cross-entropy performance with any of these methods can improve the quality of pseudo labels, thereby improving the quality of queried labeling instances.

In this work, we show how active learning for semantic segmentation can be improved with a straightforward technique: at each active learning cycle, we produce pseudo labels for the dataset with the self-training framework [13] and leverage contrastive representation learning to improve the boundaries between different classes [12]. Our combination of self-training with contrastive representation learning in S4AL+ enables querying superior samples for active learning, which results in more efficient learning. While S4AL+ was designed for active learning, we also demonstrate state-of-the-art performance on semi-supervised learning.

2 Background

Pool-based Active Learning for Semantic Segmentation is a technique for ranking unlabeled data points on their importance based on machine learning methods including, but not limited to, consistency [13, 13, 21, 13], diversity [11, 13] and feature level learning [1, 29, 11, 19]. These frameworks can be further classified based on their approach for querying the data (image [11, 15, 15, 19], region [1, 15, 21, 29, 12] or pixel [11]). Notably, EquAL [13], Minimax [11], and S4AL [29] proposed including the unlabeled dataset as part of the training protocol to achieve better image-level (Minimax) or region-level (EquAL, S4AL) results. Unlike these approaches, we show that adding a simple contrastive embed-

ding head and self-training to boost the networks' quality of pseudo labels is sufficient to attain state-of-the-art results.

Semi-Supervised Learning for Semantic Segmentation commonly learns representations using a teacher-student framework [[12], [14], [13], [25], [26], [53], [46]], and recently, combining teacher-student with contrastive embeddings [[1], [22], [52], [45]]. Most teacher-student frameworks use the mean teacher framework as their foundation [[55]], wherein a slowly updated version (teacher) of the continuously updated model (student) predicts reliable pseudo labels on the unlabeled data points for joint training. The exceptions to this line of approach are ST++ [[10] and USRN [[12] which use self-training based learning [[53]]. These algorithms complement our work, which differs in two aspects of learning. Unlike ST++, we use pixellevel (local) metrics instead of image-level (global) metrics for determining the quality of pseudo labels, and unlike USRN, we use contrastive representation learning to improve the quality of pseudo labels instead of clustering to improve the data distribution.

Supervised Contrastive Learning stems from adapting the one-to-many contrastive loss to the many-to-many pipeline for mapping similarities and differences between data points in the representation space [22]. Multiple frameworks aim to boost the pixel-wise predictions in semantic segmentation by incorporating contrastive learning with memory banks (prototypes) for supervised learning [53, 52, 53] and semi-supervised learning [9, 53, 53]. Our work employs a much lighter technique, ReCo [22], which continuously updates the contrastive embeddings per data iteration and achieves a comparative performance to its memory banks counterparts.

3 Our Active Learning Framework & Algorithm: S4AL+

Active learning methods work with an initial labeled pool (D_L, x_l) and a relatively larger unlabeled pool (D_U, x_u) , intending to query instances from D_U for labeling. This task is performed at every active learning cycle (AL_T) , whose stopping criteria (AL_C) is either preset or determined as the number of cycles required to achieve comparable performance to supervised learning $(D_L, D_U \rightarrow D_L)$.

We build S4AL+ on the foundations of S4AL [29], self-training [53] and contrastive representation learning [22]. Following S4AL, we query the instances x_u at every AL_T at region level using the average entropy of the network prediction's in the region. At the end of every active learning cycle, we query for labeled data from D_U by ranking the regions according to network's entropy and we account for pseudo labels by considering every pixel whose prediction confidence is above a fixed threshold. In this manner, the network has richer labels within the unlabeled pool at the beginning of the next active learning cycle and does not need to iteratively predict pseudo labels every data iteration [29, 53].

To ensure high quality within the queried regions, we investigate allowing the network to 'peep' into the data within the unlabeled pool by using semi-supervised learning. We achieve this with two simple adjustments: 1) We train the network with the self-training framework [53, 53], and 2) We add a contrastive embedding head to strengthen class-wise representations [22] (Fig. 2). We briefly discuss the motivations and functionality of both our modifications, before discussing our findings.



Figure 2: Schematic showing how our modified DeepLabv3+ framework employs contrastive representation learning to improve the performance of pseudo labels per learning cycle, which improves the performance of our active learning framework. Note how road keys group together (red) and ensure distance from sidewalk (orange), and similarly, bus keys (yellow) separate from train (green). The confidence in the prediction map highlights the keys for smart selection during contrastive learning (black arrow).

3.1 Self-training

Motivation. Mean-teacher based approaches for semi-supervised learning, wherein an exponentially moving average of the student network is used to predict pseudo labels on D_U , requires complex data augmentations (variants of CutMix [1]] or ClassMix [23]) and regularizations to function successfully on semantic segmentation. Multiple research studies in the areas of semi-supervised learning and active learning ([13, 29, 62]) use modified data augmentation pipelines to ensure robust pseudo labels. On the contrary, self-training uses a relatively simple data augmentation scheme, while managing to achieve comparable performance.

The self-training pipeline for semi-supervised semantic segmentation can be summarized in the following steps:

- 1. Train the network N_T on D_L^{wa} with CE,
- 2. Obtain pseudo labels for D_U from N_T with a determined threshold,
- 3. Jointly train D_L^{wa} and D_U^{sa} on N_S with CE,
- 4. $N_S \longrightarrow N_T$,
- 5. Repeat 2 to 4 till convergence,

where N_S , N_T , wa, sa, *CE* indicate student network, teacher network, weak data augmentation, strong data augmentation, and cross-entropy loss respectively. Self-training also relies on injecting noise into the network for augmentations, which we achieve with Dropout for MobileNetV2 and Stochastic Depth for the ResNet variants in our framework [III, III, EII]. This pipeline can be retrofitted for active learning with two minor adjustments. First, we modify Step 5 to repeat itself every AL_T , until we reach the stopping criteria cycle AL_C .

We also query x_u within D_U for labeling during every iteration of Step 2. As a result, we obtain the relatively easier set of labels in D_U via pseudo labeling and the harder set of labels in D_U via manual annotations at each active learning cycle (AL_T) . This results in a larger portion of information available for the network during the next training phase of the active learning cycle (AL_T) .

3.2 Contrastive Representation Learning

Motivation. We face two challenges when directly applying the framework for semantic segmentation: 1) for classification, the entire feature representation for the network corresponds to a single image category, and 2) cross-entropy loss works on a per-pixel basis and ignores the possibility of learning from other pixels in the image (and the batch) to improve its understanding. Hence, we explore the field of supervised contrastive learning and augment our network with a contrastive embedding head for learning meaningful representations (2). Specifically, we adapt the ReCo loss [22] into our training pipeline due to its simplicity and lower memory overhead.

We modify our network for semantic segmentation by adding another decoder head that accounts for learning representations and is trained via the ReCo loss (Fig. 2). Assuming that r indicates the representations obtained in parallel to the network's class-wise predictions, and C is the total number of classes present in the dataset during the mini-batch under training, Eqn. 1 summarizes the fundamentals of the ReCo loss [22] as follows:

$$L_{\text{ReCo}} = \sum_{c \in \mathcal{C}} \sum_{r_q \sim \mathcal{R}_q^c} -\log \frac{\exp(r_q \cdot r_k^{c,+}/\tau)}{\exp(r_q \cdot r_k^{c,+}/\tau) + \sum_{r_k^- \sim \mathcal{R}_k^c} \exp(r_q \cdot r_k^-/\tau)},$$
(1)

where

- \mathcal{R}_q^c represents the positive set containing all representations whose ground truth labels belong to class c,
- $r_k^{c,+}$ represents the positive anchor, which is the mean of r_q per class c,
- r_k^- are all representations within R_k^c , which is the negative key set containing all representations whose ground truth label is not class c, and
- τ is the scalar temperature control parameter.

Eqn. 1 accounts for all positive and negative keys present within a mini-batch whose size can grow exponentially based on the number of images and their resolutions. ReCO alleviates this problem by constructing a pair-wise similarity graph among the mean representations of every class $(r_k^{c,+}, c \in C)$ and actively sampling for a meaningful set of negative keys $r_k^$ per \mathcal{R}_q^c , whose quantity is a tunable hyper-parameter (RC_K). This ensures that semantically different classes, for example, Bus and Vegetation, are rarely sampled as a pair for learning, while semantically similar classes like Bus and Train are seen more often.

Similarly, r_q are chosen based on the threshold RC_{δ_s} for every corresponding prediction confidence from the network and a tunable hyper-parameter (RC_Q) which determines the quantity, similar to RC_K . RC_{δ_s} ensures that positive keys belong to areas within the prediction map which can benefit from representation learning. In this manner, ReCo ensures that only the most informative sets (both positive and negative) are trained while maintaining a minimal computation overhead. We replace the CE loss function in Steps 1 and 3 with a combination of CE and ReCo loss to further boost the network's learning potential.

3.3 Summary: S4AL+ vs. Other Methods

Our approach differs from S4AL, and other approaches in the semi-supervised and active learning literature in three ways: 1) We do not use the mean teacher pipeline to predict pseudo labels at every mini-batch iteration, in this manner, we save on the training time and also potential bias from the labels imbalance being carried forward every step, 2) We do not rely on heavy augmentations like CutMix [11] or ClassMix [12], and 3) We use contrastive learning with self-training to utilize the unlabeled pool to its maximum potential (Fig. 1).

4 Experiments and Results

Before presenting our results, we first introduce the metrics and datasets used in experiments.

4.1 Datasets, metrics and training protocol

Datasets. We conduct experiments on the CamVid, and CityScapes datasets [\square , \blacksquare]. CamVid has a default resolution of 720 × 960, with 11 classes of interest, and CityScapes has a default resolution of 1024 × 2048, with 11 classes of interest. For active learning, we down-sample CamVid and CityScapes to 360 × 480 and 688 × 688, respectively, for training and evaluation, following [\square , \blacksquare]. For semi-supervised learning, we crop images within the CityScapes dataset at 768 × 768 for training, and maintain the original resolution for evaluation, following [\square , \blacksquare].

Metrics. We evaluate our algorithm's performance for active learning by measuring the proportion of additional labeled data required to reach more than 95% performance in terms of mean IoU (mIoU) on fully supervised learning. We evaluate our algorithm's performance for semi-supervised learning by measuring the mIoU achieved using a specific % of fully supervised data as the labeled data.

Network Architectures. We use MobileNetV2 [50] with a modified stride of 16 as the backbone for all experiments on active learning (following [29, 59]), except in comparison to EquAL [13], wherein we use ResNet-50 with a modified stride of 8. We also use ResNet-50 [16] and ResNet-101-DeepStem [13], with modified stride of 8, for experiments on semi-supervised learning as per the protocol in the comparative state-of-the-art. We use the DeepLabv3+ framework for semantic segmentation throughout all our experiments [5].

Network Optimization. We begin all our experiments with an initial learning rate of 1×10^{-2} and use the "poly" learning schedule to gradually decrease the learning rate, similar to [**B**, **2**, **5**]. We use SGD for optimization and a weight decay of 0.0001 in all our experiments. We train the networks for 100 epochs on the CamVid and CityScapes datasets for active learning, and 240 epochs on the CityScapes dataset for semi-supervised learning. For all data augmentation schemes, we use random application of image resizing and horizontal flipping as the weak augmentations set (*wa*), and Gaussian blur, Color Jitter, and CutOut [**D**] as the strong augmentations set (*sa*). We do not create our own splits, but instead use the data splits provided in individual benchmarks to minimize disparity in our findings. Tables 1 and **S1** contain the details of our training procedures on active learning and semi-supervised learning respectively.



Figure 3: We demonstrate that our framework, S4AL+, works the best for active learning on the CamVid (a) and CityScapes (b) datasets. Both graphs show the mIoU improvements relative to the amount of labeled data utilized in comparison with the previous state-of-theart frameworks, DEAL [53] and S4AL [23], and the supervised learning performance.

4.2 Results: Active Learning

We conduct experiments for the full active learning system with the CamVid and CityScapes datasets with MobileNetV2 encoder and DeepLabv3+ framework. In both cases, we begin with 10% of the data as our labeled dataset and gradually query for additional labeled data per active learning cycle ([[25], [55]]). We query four regions per image of 30 × 30 and 43 × 43 for each active learning cycle on CamVid and CityScapes, respectively, following [[25]]. We also assign pseudo labels to the pixels that meet our prediction confidence threshold of 0.7 in every image in the unlabeled data pool.

Fig. 3 shows our results on the CamVid and CityScapes dataset. Our approach achieves over 95% of the full dataset performance with only 12.1% of the labeled pixel data for the CamVid dataset, and 15% labeled pixel data on the CityScapes dataset. We observe a significantly larger boost in the performance during the final active learning cycles, as we allow the network to train for a longer time, motivated by [**D**, **C**]. Fig. 4 shows the difference in outputs from S4AL and S4AL+, highlighting the advantage of using contrastive learning to learn better object features. We refer the reader to Tables **S3** and **S4** in the supplemental for a discussion on the class-wise performance on both datasets.

We further compare our system to EquAL [\Box], another region-based selection method, that uses a ResNet-50 encoder with DeepLab-v3+ [\Box]. Starting with 8% labeled data and a budget of 12% labeled data on the CamVid dataset, our approach achieved an mIOU of 66.4 on CamVid, compared to 63.4 from EquAL and 65.3 from S4AL. When starting with 3.5% data on the CityScapes dataset, our approach achieves a mIoU of 67.5 with only 10% labeled data, compared to EquAL's 67.4 with 12% data and S4AL's 66.7 with 10% data, thus demonstrating superior performance of our method on multiple training processes.



Figure 4: Comparison of our framework to the previous state-of-the-art S4AL. The areas of interest are highlighted and refer to regions wherein our approach supersedes S4AL in terms of correct pixel prediction.

4.3 Results: Semi-supervised Semantic Segmentation

We compare our approach without the active learning pipeline to semi-supervised learning methods. Specifically, we do not query for any additional data from the unlabeled pool of images and only use ground truth for those images that fall in the training set. We train the network on the labeled training data, assign the pseudo labels based on the prediction confidence, and retrain the network jointly with labeled and (pseudo-labeled) unlabeled data twice, before reporting our results.

For semi-supervised Learning, we vary the percentage (%) of labeled images within training data. Fig. 5(a) shows the performance of our framework in terms of previous stateof-the-art approaches. We compare S4AL+ with Mean Teacher [53], Cross Consistency Training (CCT) [73], Cross Pseudo Supervision (CPS) [73], Adaptive Equalized Learning (AEL) [73] and Unreliable Pseudo Labels (U²PL) [53] which use the ResNet-101-DeepStem encoder. We observe a slight drop in the performance on the 1/16 labeled data scenario, but perform at-par or better than other approaches in all other scenarios. Our method is also superior in terms of efficiency as 1) it does not use multiple sets of encoders and decoders (like CCT and CPS), 2) does not complex memory bank mechanisms to actively sample for long tail distribution classes (MT, AEL and U²PL), and 3) does not actively obtain pseudo labels at every data iteration (MT, AEL and U²PL). We refer the readers to Table S2 in the supplemental for in-depth results on the CityScapes dataset.

RANGNEKAR, KANAN, HOFFMAN: S4AL+

Table 1: Summary of all Active learning experiments on CamVid and CityScapes datasets. LD and ULD indicate the percentage of labeled data and unlabeled data respectively. We observe that using 1:2 labeled to unlabeled image ratio, and training for longer epochs, on the final active learning cycle, result in superior gains with respect to mIoU on both datasets.

Dataset	Experiment	Epochs	Batch Size (Labeled)	Batch Size (Unlabeled)	Iterations Per Epoch	Contrastive Representation Learning?	mIoU
CamVid	10% LD	100	4	-	50		58.42
	10% LD	100	4	-	50	Yes	60.11
	10% LD, ULD	100	2	2	100	Yes	60.50
	10% LD, ULD	100	2	4	100	Yes	60.92
	10% LD, 2.1% LD, ULD	100	2	4	100	Yes	61.80
	10% LD, 2.1% LD, ULD	150	2	4	100	Yes	62.65
CityScapes	10% LD	100	4	-	100		54.84
	10% LD	100	4	-	100	Yes	55.45
	10% LD, ULD	100	2	2	300	Yes	56.31
	10% LD, ULD	100	2	4	300	Yes	57.12
	10% LD, 1.3% LD, ULD	100	2	4	300	Yes	58.14
	10% LD, 2.6% LD, ULD	100	2	4	300	Yes	59.85
	10% LD, 3.8% LD, ULD	100	2	4	300	Yes	60.52
	10% LD, 5.1% LD, ULD	100	2	4	300	Yes	61.47
	10% LD, 5.1% LD, ULD	200	2	4	300	Yes	63.60

5 Discussion

Our results enable us to answer two questions:

- 1. *Does self-training help the case of active learning?* Yes. As a stand-alone adjustment, self-training boosts the initial mIoU by 1.8 points on the CamVid dataset and 1.2 points on the CityScapes dataset when using the MobileNetV2 encoder.
- Does representation learning help the case of active learning? Yes again. We observe a boost of 2 points in the initial mIoU for the CamVid dataset and 0.8 points on the CityScapes dataset with the MobileNetV2 encoder. We also verify this for semi-supervised learning in Fig. 5(b) and Table S2, wherein we use variants of a stronger ResNet encoder.

6 Conclusion

We propose a solution for improving active learning for semantic segmentation motivated by the desire to eliminate dependency on data augmentation schemes that involve randomness to ensure robust pseudo labels. We achieve this by formulating active learning as a stepwise semi-supervised learning problem, using self-training compared to the popular teacherstudent-based framework. We enrich this framework with a representation learning head that ensures the network can maximize its learning potential. While both our components seem reasonably straightforward, we emphasize that our goal was not to create something entirely new but to develop a solution from the existing literature. S4AL+ is a strong, but elegant, method for active and semi-supervised learning, with potential for improvement (Fig. 6), as it achieves state-of-the-art results in both domains.

RANGNEKAR, KANAN, HOFFMAN: S4AL+



Figure 5: (a) Comparison of our framework in terms of mIoU to the previous state-of-theart approaches for semi-supervised learning on CityScapes dataset. (b) Comparison of improvements gained by using contrastive learning (ReCo [22]) for semi-supervised learning on CityScapes dataset.



Figure 6: Knowledge distillation? S4AL+ still makes incorrect predictions under cases of extra-similar object appearances. However, we show that this is a limitation of the encoder, MobileNetV2, by comparing its outputs post active learning to the stand-alone output of the same approach trained under the semi-supervised learning framework (with a similar amount of initial labeled data) and the significantly more powerful ResNet-101-DeepStem encoder.

Acknowledgements

This work was supported by the Dynamic Data Driven Applications Systems Program, Air Force Office of Scientific Research under Grant FA9550-19-1-0021. We gratefully acknowledge the support of NVIDIA Corporation with the donations of the Titan X and Titan Xp Pascal GPUs used for this research, and DataCrunch.io for the usage of Nvidia V100 GPUs.

References

- Inigo Alonso, Alberto Sabater, David Ferstl, Luis Montesano, and Ana C Murillo. Semi-supervised semantic segmentation with pixel-level contrastive learning from a class-wise memory bank. In *Proceedings of the IEEE/CVF International Conference* on Computer Vision, pages 8219–8228, 2021.
- [2] Lucas Beyer, Xiaohua Zhai, Amélie Royer, Larisa Markeeva, Rohan Anil, and Alexander Kolesnikov. Knowledge distillation: A good teacher is patient and consistent. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10925–10934, 2022.

- [3] Gabriel J Brostow, Julien Fauqueur, and Roberto Cipolla. Semantic object classes in video: A high-definition ground truth database. *Pattern Recognition Letters*, 30(2): 88–97, 2009.
- [4] Arantxa Casanova, Pedro O. Pinheiro, Negar Rostamzadeh, and Christopher J. Pal. Reinforced active learning for image segmentation. In *International Conference on Learning Representations*, 2020. URL https://openreview.net/forum? id=SkgC6TNFvr.
- [5] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin P. Murphy, and Alan Loddon Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40:834–848, 2018.
- [6] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018.
- [7] Xiaokang Chen, Yuhui Yuan, Gang Zeng, and Jingdong Wang. Semi-supervised semantic segmentation with cross pseudo supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2613–2622, June 2021.
- [8] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference* on computer vision and pattern recognition, pages 3213–3223, 2016.
- [9] Terrance DeVries and Graham W Taylor. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017.
- [10] Sayna Ebrahimi, William Gan, Dian Chen, Giscard Biamby, Kamyar Salahi, Michael Laielli, Shizhan Zhu, and Trevor Darrell. Minimax active learning. *arXiv preprint arXiv:2012.10467*, 2020.
- [11] William Falcon et al. Pytorch lightning. *GitHub. Note: https://github. com/PyTorchLightning/pytorch-lightning*, 3(6), 2019.
- [12] Geoff French, Samuli Laine, Timo Aila, Michal Mackiewicz, and Graham Finlayson. Semi-supervised semantic segmentation needs strong, varied perturbations, 2020.
- [13] S. Alireza Golestaneh and Kris Kitani. Importance of self-consistency in active learning for semantic segmentation. *BMVC*, 2020.
- [14] Dayan Guan, Jiaxing Huang, Aoran Xiao, and Shijian Lu. Unbiased subclass regularization for semi-supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9968–9978, June 2022.

- [15] Elmar Haussmann, Michele Fenzi, Kashyap Chitta, Jan Ivanecky, Hanson Xu, Donna Roy, Akshita Mittel, Nicolas Koumchatzky, Clement Farabet, and Jose M Alvarez. Scalable active learning for object detection. In 2020 IEEE Intelligent Vehicles Symposium (IV), pages 1430–1435. IEEE, 2020.
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [17] Tong He, Zhi Zhang, Hang Zhang, Zhongyue Zhang, Junyuan Xie, and Mu Li. Bag of tricks for image classification with convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 558– 567, 2019.
- [18] Hanzhe Hu, Fangyun Wei, Han Hu, Qiwei Ye, Jinshi Cui, and Liwei Wang. Semisupervised semantic segmentation via adaptive equalization learning. *Advances in Neural Information Processing Systems*, 34, 2021.
- [19] Gao Huang, Yu Sun, Zhuang Liu, Daniel Sedra, and Kilian Q Weinberger. Deep networks with stochastic depth. In *European conference on computer vision*, pages 646– 661. Springer, 2016.
- [20] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *Advances in Neural Information Processing Systems*, 33:18661–18673, 2020.
- [21] Bo Li and Tommy Sonne Alstrøm. On uncertainty estimation in active learning for image segmentation. arXiv preprint arXiv:2007.06364, 2020.
- [22] Shikun Liu, Shuaifeng Zhi, Edward Johns, and Andrew J Davison. Bootstrapping semantic segmentation with regional contrast. *arXiv preprint arXiv:2104.04465*, 2021.
- [23] Sifei Liu, Shalini De Mello, Jinwei Gu, Guangyu Zhong, Ming-Hsuan Yang, and Jan Kautz. Learning affinity via spatial propagation networks. *Advances in Neural Information Processing Systems*, 30, 2017.
- [24] Yuyuan Liu, Yu Tian, Yuanhong Chen, Fengbei Liu, Vasileios Belagiannis, and Gustavo Carneiro. Perturbed and strict mean teachers for semi-supervised semantic segmentation. 2022.
- [25] Sudhanshu Mittal, Maxim Tatarchenko, and Thomas Brox. Semi-supervised semantic segmentation with high-and low-level consistency. *IEEE transactions on pattern analysis and machine intelligence*, 2019.
- [26] Viktor Olsson, Wilhelm Tranheden, Juliano Pinto, and Lennart Svensson. Classmix: Segmentation-based data augmentation for semi-supervised learning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 1369–1378, January 2021.
- [27] Yassine Ouali, Celine Hudelot, and Myriam Tami. Semi-supervised semantic segmentation with cross-consistency training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.

- [28] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. URL http://papers.neurips.cc/paper/ 9015-pytorch-an-imperative-style-high-performance-deep-learn pdf.
- [29] Aneesh Rangnekar, Christopher Kanan, and Matthew Hoffman. Semantic segmentation with active semi-supervised learning. *arXiv preprint arXiv:2203.10730*, 2022.
- [30] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [31] Gyungin Shin, Weidi Xie, and Samuel Albanie. All you need are a few pixels: Semantic segmentation with pixelpick. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, pages 1687–1697, October 2021.
- [32] Yawar Siddiqui, Julien Valentin, and Matthias Nießner. Viewal: Active learning with viewpoint entropy for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9433–9443, 2020.
- [33] Samarth Sinha, Sayna Ebrahimi, and Trevor Darrell. Variational adversarial active learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vi*sion, pages 5972–5981, 2019.
- [34] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(56):1929–1958, 2014. URL http: //jmlr.org/papers/v15/srivastaval4a.html.
- [35] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weightaveraged consistency targets improve semi-supervised deep learning results. In Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17, page 1195–1204, Red Hook, NY, USA, 2017. Curran Associates Inc. ISBN 9781510860964.
- [36] Wenguan Wang, Tianfei Zhou, Fisher Yu, Jifeng Dai, Ender Konukoglu, and Luc Van Gool. Exploring cross-image pixel contrast for semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 7303–7313, 2021.
- [37] Yuchao Wang, Haochen Wang, Yujun Shen, Jingjing Fei, Wei Li, Guoqiang Jin, Liwei Wu, Rui Zhao, and Xinyi Le. Semi-supervised semantic segmentation using unreliable pseudo labels. In *Proceedings of the IEEE/CVF International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.

- [38] Qizhe Xie, Minh-Thang Luong, Eduard Hovy, and Quoc V Le. Self-training with noisy student improves imagenet classification. In *Proceedings of the IEEE/CVF conference* on computer vision and pattern recognition, pages 10687–10698, 2020.
- [39] Shuai Xie, Zunlei Feng, Ying Chen, Songtao Sun, Chao Ma, and Mingli Song. Deal: Difficulty-aware active learning for semantic segmentation. In *Proceedings of the Asian Conference on Computer Vision*, 2020.
- [40] Lihe Yang, Wei Zhuo, Lei Qi, Yinghuan Shi, and Yang Gao. St++: Make selftraining work better for semi-supervised semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision and Pattern Recognition* (*CVPR*), 2022.
- [41] Sangdoo Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6023–6032, 2019.
- [42] Feihu Zhang, Philip Torr, René Ranftl, and Stephan Richter. Looking beyond single images for contrastive semantic segmentation learning. Advances in neural information processing systems, 34:3285–3297, 2021.
- [43] Shuai Zhao, Yang Wang, Zheng Yang, and Deng Cai. Region mutual information loss for semantic segmentation. Advances in Neural Information Processing Systems, 32, 2019.
- [44] Tianfei Zhou, Wenguan Wang, Ender Konukoglu, and Luc Van Gool. Rethinking semantic segmentation: A prototype view. In *Proceedings of the IEEE/CVF Conference* on Computer Vision and Pattern Recognition, pages 2582–2593, 2022.
- [45] Yanning Zhou, Hang Xu, Wei Zhang, Bin Gao, and Pheng-Ann Heng. C3-semiseg: Contrastive semi-supervised segmentation via cross-set learning and dynamic classbalancing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7036–7045, 2021.
- [46] Yi Zhu, Zhongyue Zhang, Chongruo Wu, Zhi Zhang, Tong He, Hang Zhang, R Manmatha, Mu Li, and Alexander Smola. Improving semantic segmentation via selftraining. arXiv preprint arXiv:2004.14960, 2020.