
Semantic Segmentation with Scarce Data

Isay Katsman^{*1} Rohun Tripathi^{*1} Andreas Veit¹ Serge Belongie¹

Abstract

Semantic segmentation is a challenging vision problem that usually necessitates the collection of large amounts of finely annotated data, which is often quite expensive to obtain. Coarsely annotated data provides an interesting alternative as it is usually substantially more cheap. In this work, we present a method to leverage coarsely annotated data along with fine supervision to produce better segmentation results than would be obtained when training using only the fine data. We validate our approach by simulating a scarce data setting with less than 200 low resolution images from the Cityscapes dataset and show that our method substantially outperforms solely training on the fine annotation data by an average of 15.52% mIoU and outperforms the coarse mask by an average of 5.28% mIoU.

1. Introduction

Semantic segmentation is the task of performing per-pixel annotation with a discrete set of semantic class labels. Current state of the art methods in segmentation use deep neural networks (Zhao et al., 2017; Shelhamer et al., 2015) which rely on massive per-pixel labeled datasets (Zhou et al., 2017; Everingham et al., 2009; Cordts et al., 2016). The labels for these datasets are finely annotated and are generally very expensive to collect. In this paper, we consider the scarce data paradigm, a setting where annotating images with fine details is extremely expensive. On the other hand, coarse or noisy segmentation annotations are substantially cheaper to collect, for example, through use of automated heuristics.

Recently, Veit et al. (Veit et al., 2017) introduced a semi-supervised framework that jointly learns from clean and noisy labels to more accurately classify images. They show

^{*}Equal contribution ¹Department of Computer Science, Cornell Tech and Cornell University, NY, USA. Correspondence to: Isay Katsman <isk22@cornell.edu>, Rohun Tripathi <rt443@cornell.edu>.

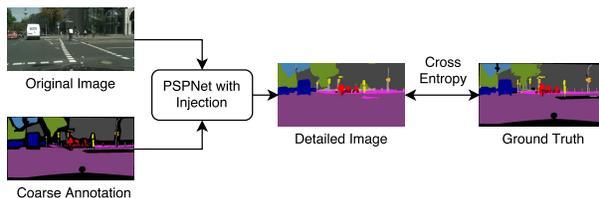


Figure 1. Overview. The proposed method is trained to improve the detail of coarse annotations by using them together with regular images in a PSPNet pipeline supervised with fine annotations.

that the paradigm of first training a network on noisy labels and then fine-tuning on the fine labels is non-optimal. Specifically, this naïve approach does not fully leverage the information in the clean annotations. In this work, we propose a related approach for semantic segmentation, in which the goal is to learn jointly from both coarse and fine segmentation masks to provide better image segmentations. The proposed model is able to take the input image together with a coarse mask to produce a detailed version, Figure 1. In particular, we model the segmentation network as conditionally dependent on the coarse masks, based on the intuition that the coarse masks provide a good starting point for the segmentation task.

We use the Cityscapes dataset to evaluate our method. This segmentation dataset has the important characteristic of providing both fine and coarse annotations. Fine annotations are limited to only 5000 images, whereas coarse annotations are available for these 5000 images and for an additional 20000 images. To simulate a setting of scarce data, we limit our training set to a small number of finely annotated images. We primarily focus on the scarce data setting, which we define as having less than or equal to 200 finely annotated images (together with their corresponding coarse masks).

This paper’s main contributions are as follows. We present a method for segmentation in the scarce data setting. We utilize the coarse and fine masks jointly and produce better segmentation results than would be obtained when training solely with fine data or using the coarse masks alone. Our approach is validated in a scarce data setting on the Cityscapes dataset, with average gains of 15.52% mIoU¹

¹Averages are taken over our tested scarce finely annotated

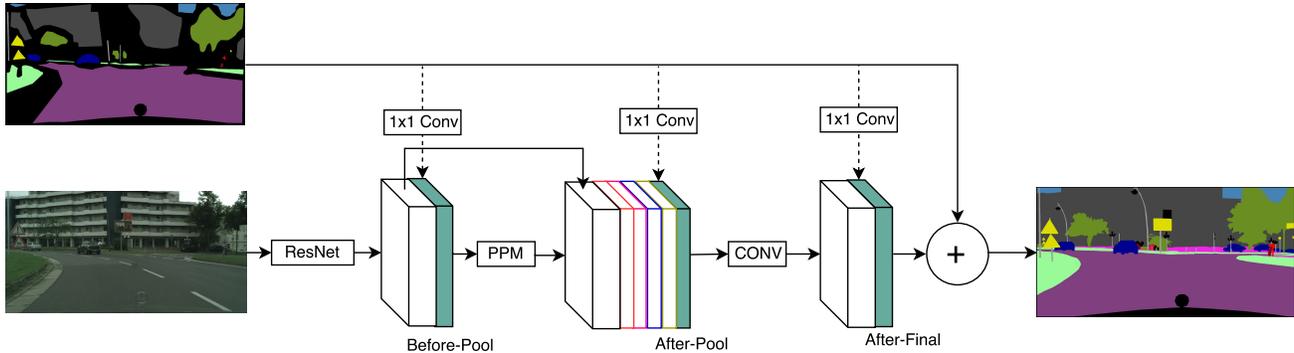


Figure 2. Overview of the proposed detailer network. The solid green boxes represent the possible locations for the injection of the coarse mask embedding. Only one of the injection locations is used for each variant. The architecture is based on the PSPNet.

over the corresponding baseline PSPNet and average gains of 5.28% over using the coarse mask directly as predictions.

2. Related Work

By adapting Convolutional Neural Networks for semantic segmentation, substantial progress has been made in many segmentation benchmarks (Badrinarayanan et al., 2017; Ronneberger et al., 2015; Shelhamer et al., 2015; Chen et al., 2018; Zhao et al., 2017). Semi-supervised semantic segmentation is an active field of research given that collection of large datasets of finely annotated data is expensive (Pinheiro et al., 2016; Garcia-Garcia et al., 2017; Papandreou et al., 2015). In (Pinheiro et al., 2016), the authors propose a refinement module to merge rich low level features with high-level object features learned in the upper layers of a CNN. Another method for semi-supervised segmentation is (Papandreou et al., 2015), which uses Expectation-Maximization (EM) methods for training DCNN models.

Providing low level information (such as the coarse masks in our method) in the form of an embedding layer is explored in previous works (Veit et al., 2017; Perez et al., 2017). In (Veit et al., 2017), the noisy labels are jointly embedded with the visual features extracted from an Inception-V3 (Szegedy et al., 2016) ConvNet. We explore an analogous approach, which concatenates the coarse masks with the convolution blocks at different network locations in the scarce data semantic segmentation domain.

3. Approach

Our goal is to leverage the small set of fine annotations and their coarse counterparts to learn a mapping from the coarse to fine annotation space. Both the coarse and fine segmentation labels consist of 20 class labels (one being the void

dataset sizes of 10, 25, 50, 100, 200.

label). Formally, we have a small dataset V of triplets, each of which consists of fine annotation v , coarse annotation y and image m . We thus have $V = \{(v_i, y_i, m_i), \dots\}$. Each annotation is a mask, with pixels that can take 20 values (each value representing a different class). Our goal is to design an effective approach to leverage the quality of the labels in v to map from the coarse labels in y to improved segmentation masks. The mapping is conditioned on the input image. We coin the name “detailer” for an architecture that achieves this goal.

3.1. Architecture of Detailer

We utilize a PSPNet for our classifier baseline with training details as described in the original paper. We propose a detailer architecture that is a modified PSPNet, depicted in Figure 2. For each triplet in V , we forward m_i through the network and insert an embedding of the coarse annotation, y_i , into the network. The embedding is inserted either before the PPM module (before-pool), after the PPM module (after-pool) or after the final convolution block (after-final).

To embed the coarse mask, we expand the 2D annotation image of size $W \times H$ to a $19 \times W \times H$ tensor, where W and H are the height and width of the mask and each channel is a binary $W \times H$ mask. The channel-wise binary vector for each pixel is a one-hot encoding for each pixel, i.e., it takes 1 at the k^{th} index, where k is the class of the pixel. A 1×1 convolution is applied on this mask tensor to produce its embedding. The number of filters used for this function is a hyper-parameter, with the default value being 800 filters. This embedding is concatenated with the visual feature tensor at the configured network location.

Another key detail of the detailer network is an identity-skip connection that adds the coarse annotation labels from the training set to the output of the cleaning module. The skip connection is inspired by the approach from (He et al., 2016).

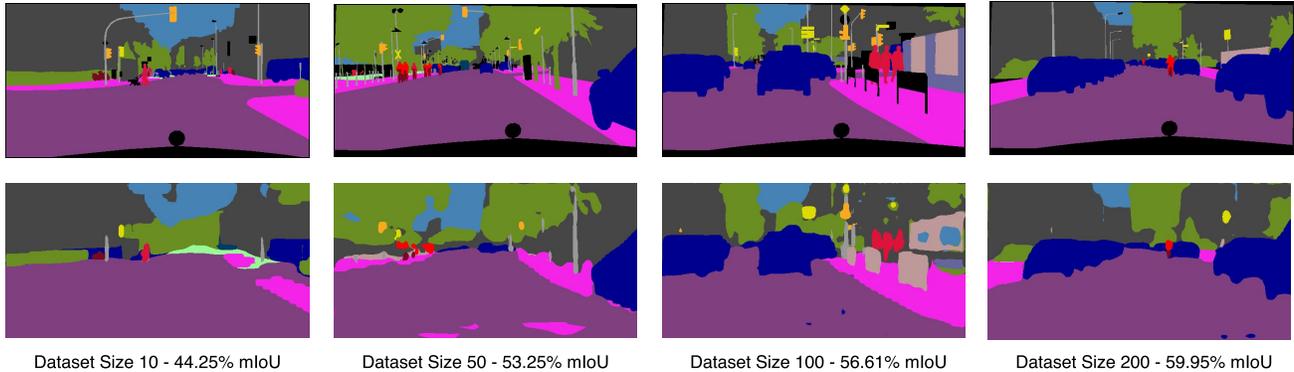


Figure 3. **Qualitative results of the proposed method** on Cityscapes (Cordts et al., 2016). Top images show the ground truth and the bottom images show the result of our model. The models used for predicted images are trained on 10, 50, 100, and 200 images with a low resolution of 340×340 . Pixels that correspond to black pixels in the ground truth are ignored during validation, as 'ignore labels'.

Due to this connection, the network only needs to learn the difference between the fine and the coarse annotations labels, thus simplifying the learning task. For the i^{th} image, the coarse mask, y_i and the predicted image annotation tensor, p_i , both take the shape $19 \times W \times H$. The first denotes a one-hot encoding tensor described earlier and the second represents the corrections for the predicted class values. These two tensors are added together for the final detailed prediction.

4. Experimental Setting

Our experiments are done with simulated scarce subsets of the Cityscapes dataset, using mIoU as our evaluation metric.

4.1. Dataset

We evaluate our proposed architecture on the Urban Landscape Cityscapes Dataset (Cordts et al., 2016). The dataset is suited for our task as it contains a coarse annotation along with a fine annotation for 5000 images. The coarse mask is made with an objective to annotate as many pixels as possible within 7 minutes. This was achieved by labeling coarse polygons under the sole constraint that each polygon must only include pixels belonging to a single class. These coarse masks have very low false positives, 97% of all labeled pixels in the coarse annotations were assigned the same class as the fine annotations. In comparison, the annotation process of each of the fine masks takes more than 1.5 hours on average.

The mIoU of the coarse masks compared against the fine masks for the validation set is 47.81%. Because our setting uses the coarse masks at test time, an mIoU of 47.81% can be achieved by directly outputting the coarse mask. Thus,

our detailer result should exceed this value at test time.

4.2. Training Details

We use data augmentation as follows - random sizing (between 0.5 and 2.0), random cropping, random rotation (between -10 and +10 degrees) and random horizontal flip. Resnet-101 pretrained on ImageNet was used to initialize the initial layers of the PSPNet. Accordingly, pixel values are normalized using the ImageNet mean and standard deviation. To train the networks we use cross-entropy loss and optimized the network using SGD with an initial learning rate 0.01, polynomial learning rate decay of 0.9 every iteration and momentum set to 0.99. We stopped training after 10K mini-batches for datasets sizes of 200 and under, and after 90k mini-batches for dataset sizes above 200.

5. Results

Our approach shows that for dataset sizes that are very small, it is still possible to learn a segmentation network with relatively high performance. We show that our detailer outperforms both the coarse mask and the regular classifier for small dataset sizes. Surprisingly, we also notice that although higher resolution images typically work much better than low resolution images when a lot of segmentation data is available, the same does not hold true in the scarce data setting for the detailer (elaboration in "Higher Resolution" section).

When utilizing a very small dataset of fine and coarse annotations to train, our detailer network performs reasonably well on the Cityscapes validation set, as shown in Table 1. For every size, the detailer outperforms the classifier. The average gain is 15.52% mIoU over the classifier perfor-

Dataset size	Detailer mIoU	Classifier mIoU
10	44.25%	23.54%
25	51.34%	35.49%
50	53.26%	38.54%
100	56.61%	42.75%
200	59.95%	47.50%
500	60.45%	45.78%
2975 (complete)	64.94%	64.11%

Table 1. **Segmentation performance** for the proposed detailer and PSPNet with varying dataset size. Models are trained with an image resolution of 340×340 . The proposed detailer clearly outperforms the standard PSPNet for all dataset sizes with the largest improvements in the regime with very few training examples.

Model	mIoU
Coarse Mask	47.81%
Classifier	42.75%
Detailer	56.61%
Classifier Composite	58.08%
Detailer Composite	60.60%

Table 2. **mIoU for different models.** We compare models trained on images with resolution 340×340 with a dataset size of 100. We observe that the detailer outperforms the classifier in the standard and the composite settings.

mance. Starting from only a dataset of 10 images up to 200 images, we also outperform the coarse mask by an average of 5.28% (the coarse mask mIoU is 47.81%). Qualitative results are shown in Figure 3.

Composite Prediction. A reasonable concern is that our improvement over the baseline classifier model shown in Table 1 is solely due to the fact that we incorporate the coarse mask (which is already at 47.81%). To confirm that our coarse injection model (Figure 2) improves results, we define the notion of a composite model. A composite model takes the coarse mask and adds the predictions of the trained segmentation model to it for all pixels that were not assigned a label by the coarse mask (pixels which had ‘ignore’). From Table 3 we see that if we take the trained detailer and classifier and make them composite, the detailer outperforms the classifier, indicating that the injection model approach does in fact improve performance.

Higher Resolution. As opposed to using a scarce data low-resolution setting, it is also presumable that a slightly more expensive setting of scarce data with high-resolution could occur in the real-world. To simulate this, we use the same approach as before but instead of re-sizing the Cityscapes data to 340×340 , we re-size to 850×850 . If we compare the detailer results at both resolutions (Table 3), we see that

Dataset Size	Detailer (higher res)	Detailer (lower res)
10	40.91%	44.20%
25	46.19%	52.33%
50	48.67%	51.51%
100	52.87%	53.96%
200	57.88%	56.27%
500	59.83%	58.10%

Table 3. **mIoU for varying dataset size for detailer networks.** We compare models trained on images with resolutions 340×340 and 850×850 while varying the number of training examples from 10 to 500. We observe that the former outperforms the latter when trained with lower number of images and the latter performs better as the number of data points exceeds 200.

for most dataset sizes, at low resolution the detailer does better than at high resolution. This is an important result as it shows that in the scarce data setting while using a coarse mask, using low resolution annotations is preferable to high resolution annotations. Our results indicate that one can save resources by providing coarse and fine annotations for a low resolution 340×340 instead of providing annotations at a high resolution of 850×850 .

Distillation Our detailer predicts detailed masks using coarse masks as input. These detailed masks can be used to distill the semantic information to a student network, thereby removing our need for coarse masks at test time. We observe that training a PSPNet network using just 100 detailed masks² gives 37.85% mIoU, exceeding the 34.61% mIoU obtained when training a PSPNet with the corresponding 100 coarse masks. This result shows that the detailing method captures important semantic information which can be transferred to a student network.

6. Conclusion

We proposed a method of leveraging coarse annotations to improve performance of deep neural networks for semantic segmentation in the scarce data setting. Our results show that in the case of scarce data, results can be significantly improved with the addition of coarse annotations. We observe and improvement in segmentation performance of the proposed detailer of in average 15.52% mIoU over the performance of a vanilla PSPNet. Additionally, we observe that although models trained on higher resolution images tend to perform better given sufficient training data, in the setting of scarce training data low resolution images lead to better performance.

²These detailed masks are obtained from our detailer trained on a scarce dataset of 100 images.

References

- Badrinarayanan, Vijay, Kendall, Alex, and Cipolla, Roberto. Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39:2481–2495, 2017.
- Chen, Liang-Chieh, Papandreou, George, Kokkinos, Iasonas, Murphy, Kevin, and Yuille, Alan L. 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.
- Cordts, Marius, Omran, Mohamed, Ramos, Sebastian, Rehfeld, Timo, Enzweiler, Markus, Benenson, Rodrigo, Franke, Uwe, Roth, Stefan, and Schiele, Bernt. The cityscapes dataset for semantic urban scene understanding. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3213–3223, 2016.
- Everingham, Mark, Gool, Luc Van, Williams, Christopher K. I., Winn, John M., and Zisserman, Andrew. The pascal visual object classes (voc) challenge. *International Journal of Computer Vision*, 88:303–338, 2009.
- Garcia-Garcia, Alberto, Orts, Sergio, Oprea, Sergiu, Villena-Martinez, Victor, and Rodríguez, José García. A review on deep learning techniques applied to semantic segmentation. *CoRR*, abs/1704.06857, 2017.
- He, Kaiming, Zhang, Xiangyu, Ren, Shaoqing, and Sun, Jian. Deep residual learning for image recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, 2016.
- Papandreou, George, Chen, Liang-Chieh, Murphy, Kevin P., and Yuille, Alan L. Weakly-and semi-supervised learning of a deep convolutional network for semantic image segmentation. *2015 IEEE International Conference on Computer Vision (ICCV)*, pp. 1742–1750, 2015.
- Perez, Ethan, de Vries, Harm, Strub, Florian, Dumoulin, Vincent, and Courville, Aaron C. Learning visual reasoning without strong priors. *CoRR*, abs/1707.03017, 2017.
- Pinheiro, Pedro H. O., Lin, Tsung-Yi, Collobert, Ronan, and Dollár, Piotr. Learning to refine object segments. In *ECCV*, 2016.
- Ronneberger, Olaf, Fischer, Philipp, and Brox, Thomas. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 2015.
- Shelhamer, Evan, Long, Jonathan, and Darrell, Trevor. Fully convolutional networks for semantic segmentation. *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3431–3440, 2015.
- Szegedy, Christian, Vanhoucke, Vincent, Ioffe, Sergey, Shlens, Jonathon, and Wojna, Zbigniew. Rethinking the inception architecture for computer vision. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2818–2826, 2016.
- Veit, Andreas, Alldrin, Neil, Chechik, Gal, Krasin, Ivan, Gupta, Abhinav, and Belongie, Serge J. Learning from noisy large-scale datasets with minimal supervision. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 6575–6583, 2017.
- Zhao, Hengshuang, Shi, Jianping, Qi, Xiaojuan, Wang, Xiaogang, and Jia, Jiaya. Pyramid scene parsing network. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 6230–6239, 2017.
- Zhou, Bolei, Zhao, Hang, Puig, Xavier, Fidler, Sanja, Barriuso, Adela, and Torralla, Antonio. Scene parsing through ade20k dataset. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017.