# Universal Semi-Supervised Semantic Segmentation Supplementary Material

Tarun Kalluri<sup>1</sup>

Girish Varma<sup>1</sup>

Manmohan Chandraker<sup>2</sup>

C V Jawahar<sup>1</sup>

<sup>1</sup>Center for Visual Information Technology, IIIT Hyderabad <sup>2</sup>University of California, San Diego

tarun.05.kalluri@gmail.com

## **A. Training Details**

We give details of the parameters used for training the universal segmentation models in various settings. We use the openly available PyTorch implementation of dilated residual network [4], with encoders designed using ResNet-18 (drn-d-22), ResNet-50 (drn-d-54) as well as ResNet-101 (drn-d-105) architectures. The decoder consists of a 1x1 convolution layer followed by a bilinear upsampling layer. We train every model on 2 Nvidia GeForce GTX 1080 GPUs for 200 epochs. During training, we use a crop size of 512x512 for Cityscapes and IDD datasets, 360x480 for the Camvid dataset and 480x640 for the SUN-RGB dataset. Validation mIoUs are reported on the standard resolutions from the dataset. We employ SGD learning algorithm with an initial learning rate of 0.001 and a momentum of 0.9, along with a poly learning rate schedule with a power of 0.9 [1]. We use a batch size of 10, and take the embedding dimension to be 128. The default values for  $\alpha$  and  $\beta$  are taken to be 1.

## **B.** Label Embeddings

## **B.1.** Calculating the label embeddings

In this section, we describe the method used to obtain the vector representations for the labels. For each dataset separately, we train an end-to-end segmentation network from scratch using only the limited training data available in that dataset. We use this trained segmentation network to calculate the encoder outputs of the training data at each pixel. Typically, the size of the output dimension of the encoder at each pixel (512 for a ResNet encoder) is not equal to the dimension of the label embeddings (d=128, in our case). So we first apply a dimensionality reduction technique like PCA to reduce the dimension of the outputs to match the dimension of the label embeddings d, and then calculate the class wise centroids to obtain the label embeddings.

| Method               |       | N=50  |       | N=100 |       |       |  |  |  |  |
|----------------------|-------|-------|-------|-------|-------|-------|--|--|--|--|
| lineunou             | CS    | CVD   | Avg.  | CS    | Avg.  |       |  |  |  |  |
| $Ours[\theta = 0.5]$ | 33.28 | 48.7  | 40.99 | 33.51 | 49.49 | 41.50 |  |  |  |  |
| Ours[Word2Vec]       | 33.48 | 53.19 | 43.34 | 36.18 | 52.72 | 44.45 |  |  |  |  |
| Ours[K=1]            | 34.01 | 53.23 | 43.62 | 41.03 | 54.62 | 47.83 |  |  |  |  |
| Ours[K=3]            | 35.23 | 52.38 | 43.81 | 41.82 | 54.96 | 48.39 |  |  |  |  |
| Ours[K=5]            | 33.76 | 52.77 | 43.27 | 40.08 | 55.02 | 47.55 |  |  |  |  |
| Direct SER           | 21.36 | 35.24 | 28.30 | 23.7  | 30.67 | 27.19 |  |  |  |  |

Table 1: Extension of Table 3 from the original paper.  $\theta$  is the update factor during training, and the default value is 1. K is the number of embeddings per label. *Ours[Word2Vec]* uses word vectors as label embeddings. The model gives best performance for K=3,  $\theta = 1$  while using prototype embeddings.

#### **B.2. Updating the label embeddings**

In our original experiments, we fixed the pretrained label embeddings over the phase of training the universal model. Here, we present a method to jointly train the segmentation model as well as the label embeddings. We initialize the embeddings with the values computed from the pretrained networks, and make use of the following exponentially weighted average rule to update the centroids at the  $t^{th}$  time step.

$$\mathbf{c}_{t}^{(k)} = \theta \cdot \mathbf{c}_{t-1}^{(k)} + (1-\theta) \cdot \mu_{L}(\mathcal{F}_{t}(x_{u}^{(k)})).$$
(1)

In Eq (1),  $c_{t-1}^{(k)}$  denotes the centroids at the  $(t-1)^{\text{th}}$  time step,  $\mathcal{F}_t$  is the state of the encoder module at the  $t^{th}$  time step and  $\mu_L$  calculates the class wise centroids.  $\theta$  is the update factor, where a value of  $\theta = 1$  implies that the centroids are not updated from their initial state, and a value of  $\theta = 0$  means that the centroids are calculated afresh at each update. We make an update to the centroids after every mini-batch of the original training data.

From Table 1, experiments with  $\theta = 0.5$  suggests that jointly training the network as well as updating the label embeddings reduces the performance compared to having fixed label embeddings. We believe that this is primarily due to having insufficient training data for jointly updating embeddings as well as the network weights, although this merits a deeper investigation.

### **B.3.** Multiple Label Embeddings

Many labels in a segmentation dataset often appear in more than one visual form or modalities. For example, road class can appear as dry road, wet road, shady road etc., or a class labeled as building can come in different structures and sizes. To better capture the multiple modalities involved in the visual information of the label, we propose using multiple embeddings for each label instead of a single mean centroid. This is analogous to polysemy in vocabulary, where many words can have multiple meanings and can occur in different contexts, and context specific word vector representations are used to capture this behavior. To calculate the multiple label embeddings, we perform K-means clustering of the pixel level encoder feature representations calculated from networks pretrained on the limited supervised data, and calculate similarity scores with all the multiple label embeddings.

Table 1 shows that using K=3 embeddings per label gives an advantage over using 1 embedding per label, so apparently some amount of over segmentation helps. However, further increasing K to 5 hurts the performance, as not all the labels benefit from having multiple modalities per label. So, an interesting future direction can be to examine optimum number of embeddings per label.

Particularly, from Table 2, it is evident that classes like *Road*, *Building*, *Person* etc. benefit largely from having multiple embeddings per label.

### **B.4.** Choice of label embeddings

In our work, we chose the pixel level class prototypes to be the label embeddings. We believe that this helps in better capturing visual information from the images compared to other approaches like Word2Vec [2]. To this end, we provide results of our approach replacing the prototype label embeddings with word vectors of the labels, by using the publicly available 128 dimensional word vectors for the labels from the Cityscapes and CamVid datasets.

From Table 1, having class prototypes as label embeddings, which are computed from the labeled data, performs better than using Word2vec based embeddings, which capture semantics of the word meaning rather than the visual appearance of the label. The performance improvement is more evident in case of N=100, which demonstrates that in presence of sufficient labeled data, class prototypes are better suited as label embeddings than word vector representations. Similar observations have been made in [3] as well.

### C. Direct Softmax Entropy Regularization

Entropy regularization is used to enhance the confidence of predictions made on unlabeled samples. In the case of deep neural networks, applying this directly to the softmax scores will make the predictions confident by simply increasing the weights of the last layer. So, we follow an approach where we calculate similarity between normalized label prototypes and encoder embeddings through our entropy module. *Direct SER* result from Table 1 further demonstrates the fact that applying SER (softmax entropy regularization) directly to our network shows inferior performance compared to the proposed entropy module based approach.

### **D.** Qualitative examples

From the qualitative examples presented in Figure 1 for Cityscapes and Figure 2 for IDD, we can observe that the entropy module is particularly useful in better segmentation of finer objects and foreground classes.

### References

- L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2018.
- [2] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119, 2013. 2
- [3] J. Snell, K. Swersky, and R. Zemel. Prototypical networks for few-shot learning. In *Advances in Neural Information Processing Systems*, pages 4077–4087, 2017.
- [4] F. Yu, V. Koltun, and T. A. Funkhouser. Dilated residual networks. In *CVPR*, volume 2, page 3, 2017. 1

| Method            | Road                           | SideWalk                       | Building                       | Wall                     | Fence                           | Pole                            | Traffic light                  | Traffic Sign                   | Vegetation                     | Train                          | Sky                            | Person                         | Rider                          | Car                            | Truck                          | Bus                      | Train                          | MotorCycle                    | Bicycle                        | mloU                           |
|-------------------|--------------------------------|--------------------------------|--------------------------------|--------------------------|---------------------------------|---------------------------------|--------------------------------|--------------------------------|--------------------------------|--------------------------------|--------------------------------|--------------------------------|--------------------------------|--------------------------------|--------------------------------|--------------------------|--------------------------------|-------------------------------|--------------------------------|--------------------------------|
| K=1<br>K=3<br>K=5 | 92.18<br><b>93.10</b><br>89.58 | 51.29<br><b>56.82</b><br>48.50 | 80.07<br>80.48<br><b>81.21</b> | 00.00<br>00.03<br>14.55  | <b>24.01</b><br>17.84<br>07.89  | <b>33.73</b><br>32.49<br>27.77  | <b>26.16</b><br>24.07<br>22.72 | <b>38.71</b><br>33.51<br>33.95 | 82.30<br><b>82.52</b><br>84.36 | 36.39<br><b>38.52</b><br>34.33 | <b>81.61</b><br>80.12<br>80.52 | 54.38<br>53.22<br><b>54.39</b> | 20.48<br>15.35<br><b>22.51</b> | <b>81.71</b><br>81.34<br>81.52 | 02.37<br><b>07.79</b><br>02.41 | <b>22.79</b> 20.79 09.43 | 03.85<br>04.18<br><b>07.28</b> | 01.31<br>22.57<br>10.40       | 46.23<br><b>49.90</b><br>48.18 | 41.03<br><b>41.82</b><br>40.08 |
| Metho             | od                             | Sky                            | Ę                              | .ima                     | Pole                            |                                 | Road                           | Davia                          | 1 avc.                         | Tree                           |                                | Sign                           | Fence                          |                                | Car                            | Ped                      | T CTI.                         | Bicy.                         |                                | mloU                           |
| K=1<br>K=3<br>K=5 |                                | 86.3<br><b>87.67</b><br>86.07  | 77<br><b>78</b><br>76          | .23<br><b>.51</b><br>.39 | <b>17.1</b> 3<br>16.37<br>15.25 | <b>3</b> 8<br>7 8<br>5 <b>8</b> | 4.99<br>4.84<br><b>57.87</b>   | 53.<br>53.<br><b>63</b>        | 35<br>18<br>. <b>6</b>         | 70.57<br><b>73.33</b><br>70.95 | 31<br><b>3</b> 4<br>3          | 1.99<br><b>4.02</b><br>2.6     | 32.4<br>27.7<br><b>34.</b> 0   | 5 7<br>1 7<br>6 7              | 72.94<br>74.42<br><b>76.63</b> | 36.<br><b>40.</b><br>30. | 61<br><b>36</b><br>42          | <b>37.22</b><br>34.24<br>30.9 |                                | 54.61<br>54.96<br><b>55.02</b> |

Table 2: Class-wise IoU values for the 19 classes in Cityscapes dataset and 11 classes in the CamVid dataset for different K, for N=100 on Resnet-18.



(a) Original Image

(b) Without Entropy Module

(c) With Entropy Module

(d) Ground Truth Segmentation

Figure 1: Qualitative examples from the cityscapes dataset with and without the proposed entropy regularization module, when trained on a universal model on CS+CamVid.



(a) Original Image

(b) Without Entropy Module

(c) With Entropy Module

(d) Ground Truth Segmentation

Figure 2: Qualitative examples from the IDD dataset with and without the proposed entropy regularization module, when trained on a universal model on IDD+CS.