Supplementary Material for Superpixel Sampling Networks

Varun Jampani[1], Deqing Sun[1], Ming-Yu Liu[1], Ming-Hsuan Yang[1,2], Jan Kautz[1]

[1]NVIDIA [2]UC Merced
{vjampani,deqings,mingyul,jkautz}@nvidia.com, mhyang@ucmerced.edu

In Section 1, we formally define the Acheivable Segmentation Accuracy (ASA) used for evaluating superpixels. Then, in Section 2, we report F-measure and Compactness scores with more visual results on different datasets. We also include a supplementary video[1] that gives an overview of Superpixel Sampling Networks (SSN) with a glimpse of experimental results.

1 Evaluation Metrics

Here, we formally define the Achievable Segmentation Accuracy (ASA) metric that is used in the main paper. Given an image I with n pixels, let $H \in \{0, 1, \cdots, m\}^{n \times 1}$ denotes the superpixel segmentation with m superpixels. H is composed of m disjoint segments, $H = \bigcup_{j=1}^{m} H^j$, where j^{th} segment is represented as H^j. Similarly, let $G \in \{0, 1, \cdots, w\}^{n \times 1}$ denotes ground-truth (GT) segmentation with w segments. $G = \bigcup_{l=1}^{w} G^l$, where G^l denotes l^{th} GT segment.

ASA Score. The ASA score between a given superpixel segmentation H and the GT segmentation G is defined as

$$ASA(H, G) = \frac{1}{n} \sum_{H^j \in S} \max_{G^l} |H^j \cap G^l|, \tag{1}$$

where $|H^j \cap G^l|$ denotes the number of overlapping pixels between S^j and G^l. To compute ASA, we first find the GT segment that overlaps the most with each of the superpixel segments and then sum the number of overlapping pixels. As a normalization, we divide the number of overlapping pixels with the number of image pixels n. In other words, ASA represents an upper bound on the accuracy achievable by any segmentation step performed on the superpixels.

Boundary Precision-Recall. Boundary Recall (BR) measures how well the boundaries of superpixel segmentation aligns with the GT boundaries. Higher BR score need not correspond to higher quality of superpixels. Superpixels with high BR score can be irregular and may not be useful in practice. Following reviewers' suggestions, we report Boundary Precision-Recall curves instead of just Boundary Recall scores.

[1] https://www.youtube.com/watch?v=q37MxZolDck

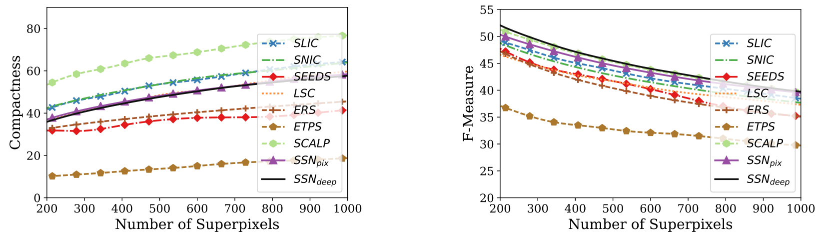

(a) Compactness. (b) F-measure.

Fig. 1: **BSDS500 test results.** (a) Compactness and (b) F-measure plots. SSN$_{deep}$ outperforms other techniques also in terms of F-measure while maintaing the same compactness as that of SSN$_{pix}$.

We also report F-measure and Compactness in the next section (Section 2.1). We use the evaluation scripts from [8] with default parameters to compute Boundary Precision-Recall, F-measure and Compactness.

2 Additional Experimental Results

2.1 Compactness and F-measure

We compute compactness (CO) of different superpixels on the BSDS dataset (Fig. 1(a)). SSN superpixels have only slightly lower CO compared to widely-used SLIC showing the practical utility of SSN. SSN$_{deep}$ has similar CO as SSN$_{pix}$ showing that training SSN, while improving ASA and boundary adherence, does not destroy compactness. More importantly, we find SSN to be flexible and responsive to task-specific loss functions and one could use more weight (λ) for the compactness loss (Eq. 6 in the main paper) if more compact superpixels are desired. In addition, we also plot F-measure scores in Fig. 1(b). In summary, SSN$_{deep}$ also outperforms other techniques in terms of F-measure while maintaining the compactness as that of SSN$_{pix}$. This shows the robustness of SSN with respect to different superpixel aspects.

2.2 Additional Visual Results

In this section, we present additional visual results of different techniques and on different datasets. Figs. 2, 3 and 4 show superpixel visual results on three segmentation benchmarks of BSDS500 [2], Cityscapes [4] and PascalVOC [5] respectively. For comparisons, we show the superpixels obtained with 3 existing superpixel techniques of SLIC [1], LSC [6] and ERS [7]. Fig. 5 shows additional visual results on MPI-Sintel [3] where we present sample segmented flows obtained using different types of superpixels.

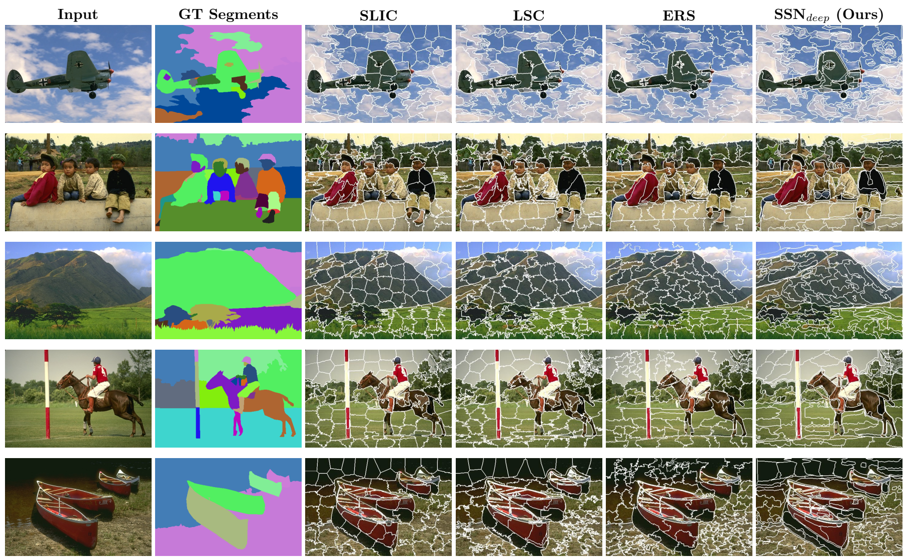

Fig. 2: **Additional visual results on BSDS500 test images.** SSN$_{deep}$ tends to produce smoother object contours and more superipxels near object boundaries in comparison to other superpixel techniques.

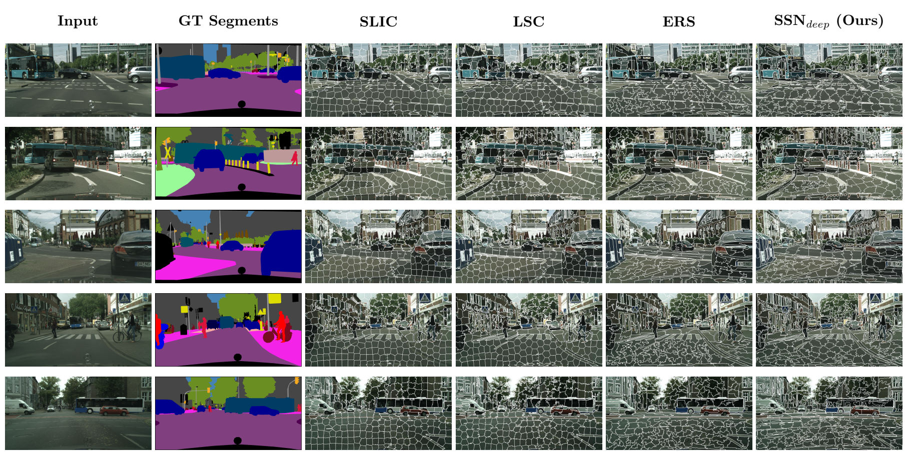

Fig. 3: **Additional visual results on Cityscapes validation images.** SSN$_{deep}$ tend to generate bigger superpixels on uniform regions (such as road) and more superpixels on smaller objects.

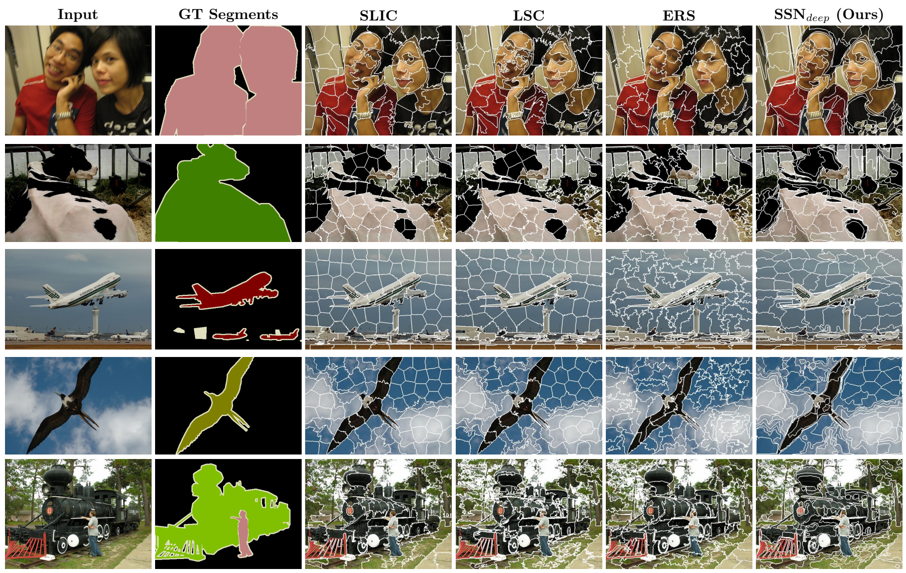

Fig. 4: **Additional visual results on PascalVOC validation images.** SSN$_{deep}$ tends to produce smoother object contours and more superipxels near object boundaries in comparison to other superpixel techniques.

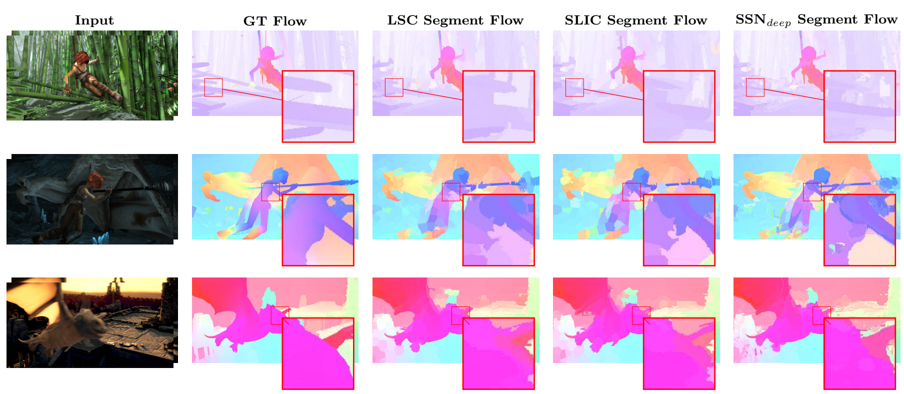

Fig. 5: **Additional visual results on Sintel images.** Segmented flow visuals obtained with different types of superpixels indicate that SSN$_{deep}$ superpixels can better represent GT optical flow compared to other techniques.

References

1. Achanta, R., Shaji, A., Smith, K., Lucchi, A., Fua, P., Süsstrunk, S.: SLIC superpixels compared to state-of-the-art superpixel methods. IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI) **34**(11), 2274–2282 (2012)
2. Arbelaez, P., Maire, M., Fowlkes, C., Malik, J.: Contour detection and hierarchical image segmentation. IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI) **33**(5), 898–916 (2011)
3. Butler, D.J., Wulff, J., Stanley, G.B., Black, M.J.: A naturalistic open source movie for optical flow evaluation. In: European Conference on Computer Vision (ECCV). pp. 611–625. Springer (2012)
4. Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset for semantic urban scene understanding. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2016)
5. Everingham, M., Eslami, S.A., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The Pascal visual object classes challenge: A retrospective. International Journal of Computer Vision (IJCV) **111**(1), 98–136 (2015)
6. Li, Z., Chen, J.: Superpixel segmentation using linear spectral clustering. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2015)
7. Liu, M.Y., Tuzel, O., Ramalingam, S., Chellappa, R.: Entropy rate superpixel segmentation. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2011)
8. Stutz, D., Hermans, A., Leibe, B.: Superpixels: An evaluation of the state-of-the-art. Computer Vision and Image Understanding **166**(C), 1–27 (2018)