Supplementary Material for submission 2147: Traditional Saliency Reloaded: A Good Old Model in New Shape

Simone Frintrop, Thomas Werner, and Germán M. García
Institute of Computer Science III
Rheinische Friedrich-Wilhelms-Universität Bonn, Germany
frintrop@iai.uni-bonn.de

In the supplementary material, we show the complete set of evaluation plots. Sec. 1 shows some evaluations of the most important parameters of our system and Sec. 2 shows the evaluation plots on four different datasets.

1 Influence of Parameters

Fig. 1 shows an evaluation of the most important parameters of our system, namely the choice of the layers of the pyramids (left/middle) and the choice of the center-surround ratio for creating the surround pyramid from the center pyramid (right). For each plot, we keep the other parameters which are not evaluated fixed to the optimal value.

It can be seen that the system is not very sensitive to both parameters: similar values also perform quite well. Even a single layer obtains reasonable results if the layer is chosen well (layer 2 or 3), which is interesting for high-speed performance on real-time systems. However, this diminishes also the ability of the system to generalize, thus we recommend to use at least 3 layers for good performance.

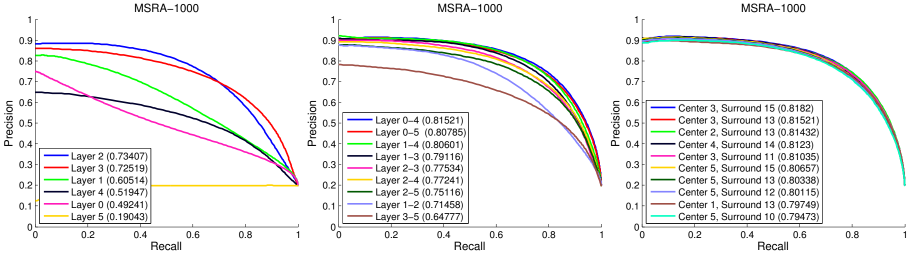

Fig. 1: Evaluation of parameters for our saliency system on the MSRA-1000 dataset. Left/middle: effect of layer choice of the pyramids. Right: Effect of Center-Surround Ratio parameter.

2 Comparative Results on Benchmarks

We show the results of our VOCUS2 saliency system (V2-Basic = pixel-precise; V2-LP = pixel-precise with location prior; V2-Prop = segment-based generated with object proposals) obtained in four different datasets for salient object segmentation: MSRA 10k [5][1] in Fig. 2, ECSSD [13][2] in Fig. 3, PASCAL-S [10][3] in Fig. 4 and SED [3][4] in Fig. 5. We compare our method with the following saliency methods: Itti's iNVT [7], the SaliencyToolbox (STB) [12], HZ08 [6], AIM [4], AC09 [1], AC10 [2], CoDi [9], HSaliency [13], Yang 2013 [14], and DRFI [8].

In terms of weighted F-measure [11] the results show that our method that integrates segmentation (V2-Prop) is the best in the MSRA 10k (Fig. 2, left), PASCAL-S (Fig. 4, left) and SED2 (Fig. 5, bottom left) datasets; it is third on SED1 (Fig. 5, top left) and second on ECSSD (Fig. 3, left). The weighted F-measure plots are separated into the methods that incorporate segmentation (left hand side) and those that do not (right hand side).

In the precision-recall plots, our V2-LP method outperforms all other pixel-precise methods, except in SED2, in which the V2-Basic approach is better. This can be easily explained since the SED2 dataset incorporates less center bias since each frame contains 2 objects. Among the segment-based approaches, the V2-Prop method is usually ranked 2nd or 3rd. The gap between our method and the top methods results at least partly from the limitations of the precision-recall plots as announced by [11]. Please note also that we intentionally optimized our code only on the MSRA-1000 dataset and not, as some other methods, on all test datasets.

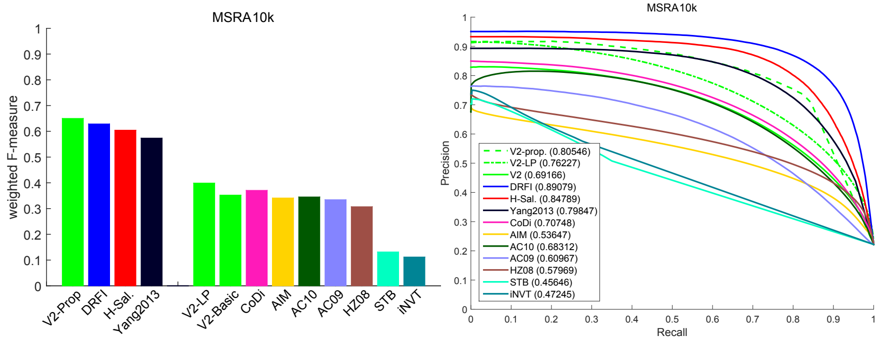

Fig. 2: MSRA10k Dataset results. Left: weighted F-measure from [11]. Right: Precision-recall curves (AUC-values in parentheses). V2-Basic: pixel-precise saliency map. V2-LP: pixel-precise saliency map with location prior. V2-Prop: proposal maps obtained from generating object proposals and fusing them to a saliency map.

[1] http://mmcheng.net/msra10k/

[2] http://www.cse.cuhk.edu.hk/leojia/projects/hsaliency/dataset.html

[3] http://cbi.gatech.edu/salobj/

[4] http://www.wisdom.weizmann.ac.il/~vision/Seg_Evaluation_DB/dl.html

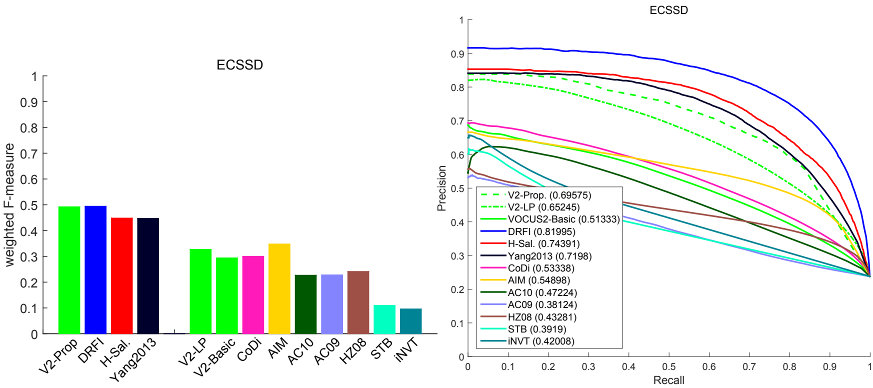

Fig. 3: ECSSD Dataset results. Left: weighted F-measure from [11]. Right: Precision-recall curves (AUC-values in parentheses). V2-Basic: pixel-precise saliency map. V2-LP: pixel-precise saliency map with location prior. V2-Prop: proposal maps obtained from generating object proposals and fusing them to a saliency map.

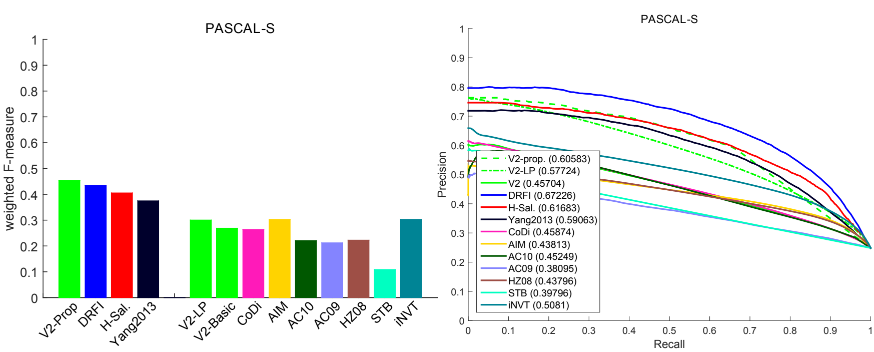

Fig. 4: PASCAL-S Dataset results. Left: weighted F-measure from [11]. Right: Precision-recall curves (AUC-values in parentheses). V2-Basic: pixel-precise saliency map. V2-LP: pixel-precise saliency map with location prior. V2-Prop: proposal maps obtained from generating object proposals and fusing them to a saliency map.

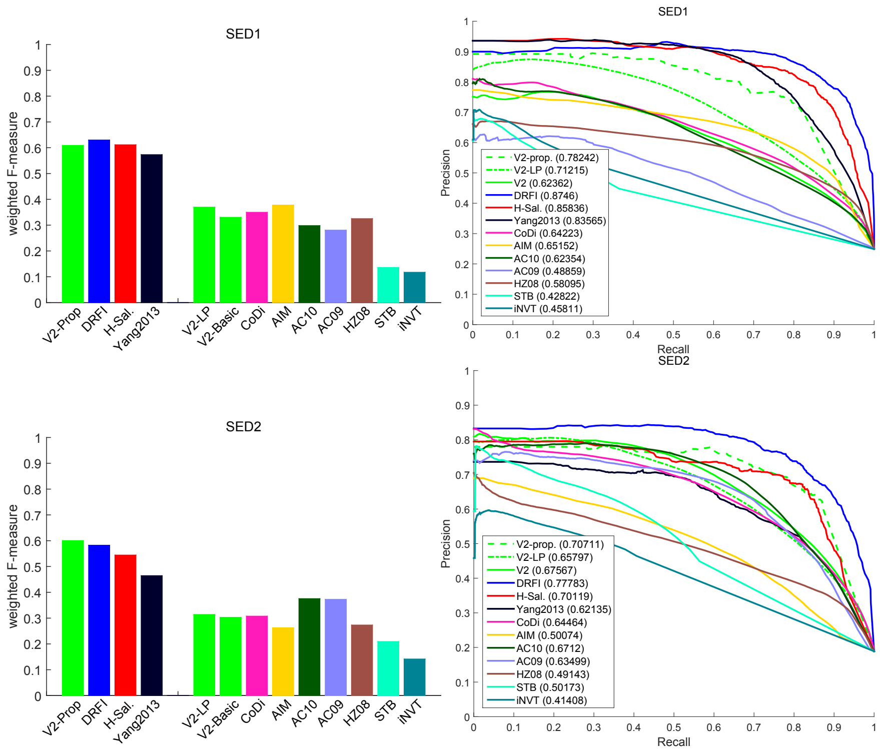

Fig. 5: SED1 and SED2 Datasets results. Left: weighted F-measure from [11]. Right: Precision-recall curves (AUC-values in parentheses). V2-Basic: pixel-precise saliency map. V2-LP: pixel-precise saliency map with location prior. V2-Prop: proposal maps obtained from generating object proposals and fusing them to a saliency map.

References

1. R. Achanta, S. Hemami, F. Estrada, and S. Süsstrunk. Frequency-tuned salient region detection. In *CVPR*, 2009.
2. R. Achanta and S. Süsstrunk. Saliency Detection using Maximum Symmetric Surround. In *ICIP*, 2010.
3. S. Alpert, M. Galun, R. Basri, and A. Brandt. Image segmentation by probabilistic bottom-up aggregation and cue integration. In *CVPR*, 2007.
4. N. D. B. Bruce and J. K. Tsotsos. Saliency, attention, and visual search: An information theoretic approach. *J. of Vision*, 9(3), 2009.
5. M.-M. Cheng, N. J. Mitra, X. Huang, P. H. S. Torr, and S.-M. Hu. Global contrast based salient region detection. *TPAMI*, 37(3):569–582, 2015.
6. X. Hou and L. Zhang. Dynamic visual attention: Searching for coding length increments. In *Advances in Neural Information Processing Systems*, 2008.
7. L. Itti, C. Koch, and E. Niebur. A model of saliency-based visual attention for rapid scene analysis. *TPAMI*, 20(11), 1998.
8. H. Jiang, J. Wang, Z. Yuan, Y. Wu, N. Zheng, and S. Li. Salient object detection: A discriminative regional feature integration approach. In *CVPR*, 2013.
9. D. A. Klein and S. Frintrop. Salient Pattern Detection using W_2 on Multivariate Normal Distributions. In *DAGM-OAGM*, 2012.
10. Y. Li, X. Hou, C. Koch, J. M. Rehg, and A. L. Yuille. The secrets of salient object segmentation. In *CVPR*, 2014.
11. R. Margolin, L. Zelnik-Manor, and A. Tal. How to evaluate foreground maps? 2014.
12. D. Walther and C. Koch. Modeling attention to salient proto-objects. *Neural Networks*, 2006.
13. Q. Yan, L. Xu, J. Shi, and J. Jia. Hierarchical Saliency Detection. In *CVPR*, 2013.
14. C. Yang, L. Zhang, H. Lu, X. Ruan, and M.-H. Yang. Saliency Detection via Graph-based Manifold Ranking. In *CVPR*, 2013.