Supplementary Information for
Recombinator Networks: Learning Coarse-to-Fine Feature Aggregation

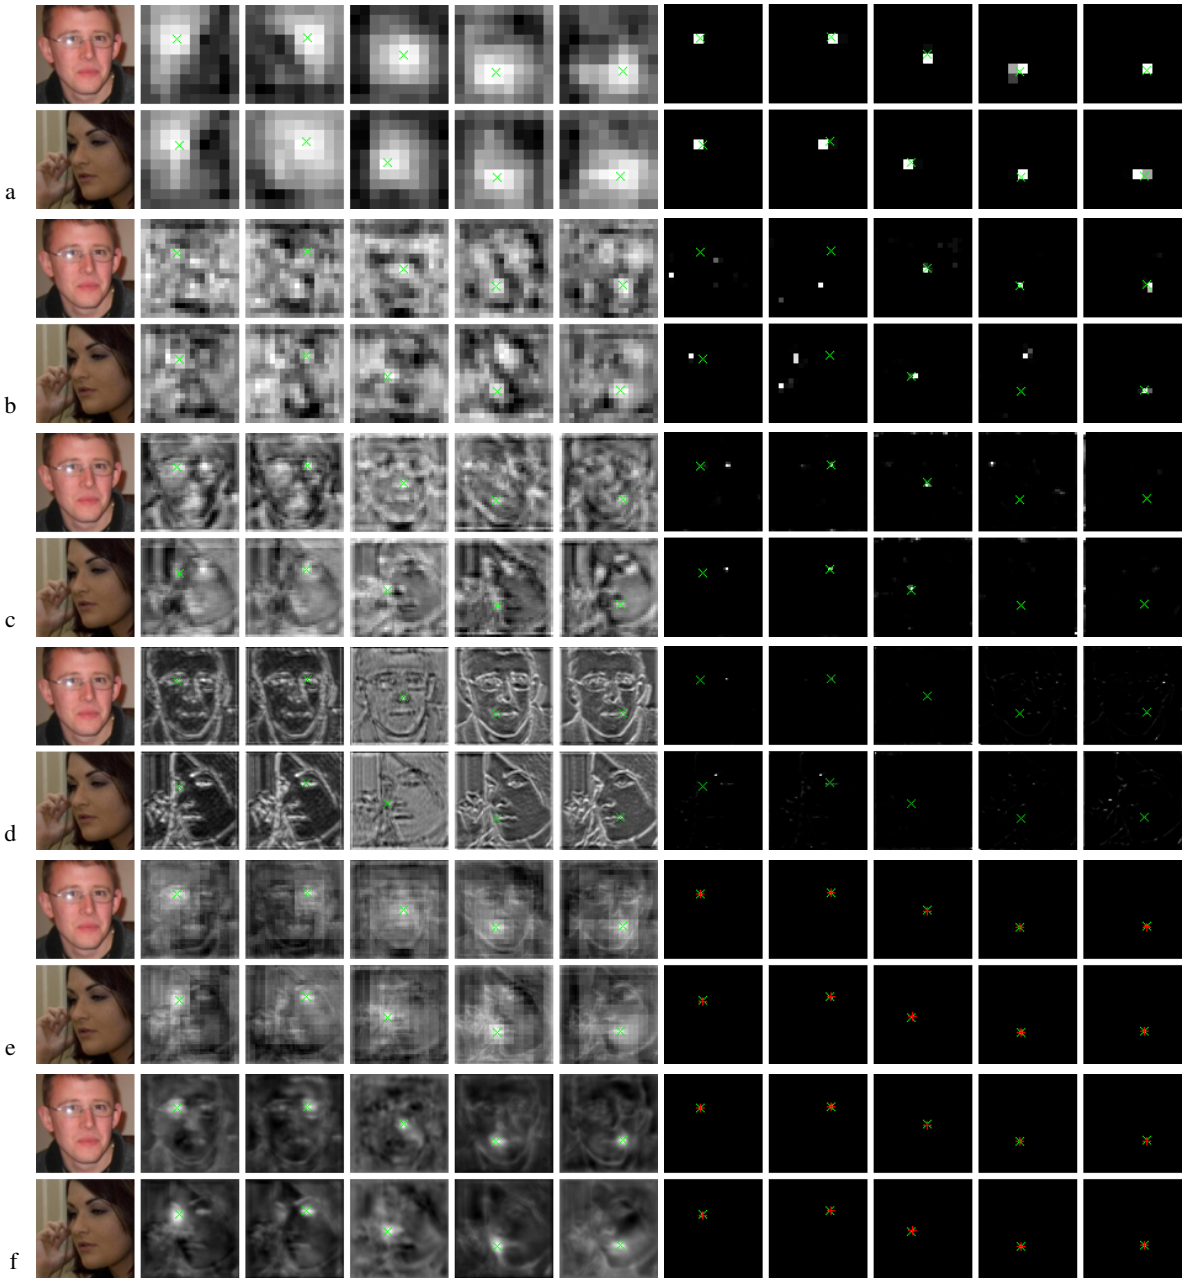

Figure S1. Sub-figures a, b, c, d show pre-sum (left) and softmax (right) of the coarsest to finest branches in a 4-branch SumNet model. The softmax used in these branches are only for illustration purposes and is not part of the trained model. Sub-figure e (left) shows the sum of branches in the SumNet model and Sub-figure f (left) depicts the pre-softmax values in RCN. The true keypoint locations are shown by green cross in all figures to show their relative correspondence with the branch activations. SumNet and RCN's predictions are shown by red plus on the post-softmax maps in Sub-figures e (right) and f (right), respectively. In each row the images correspond to the keypoints in this order from left to right: left-eye, right-eye, nose, left-mouth, right-mouth. Best viewed electronically with zoom.

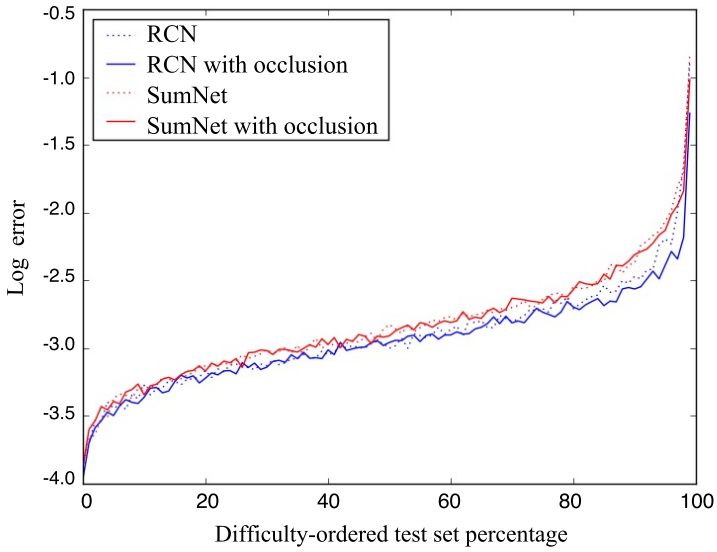

Figure S2. The performance of SumNet and RCN models with and without occlusion pre-processing on the merged AFW and AFLW test sets as the difficulty of the examples increase (lower is better). To get this plot, we note for each test set image (including both AFLW and AFW) the average error over the four models and use this as a notion of that image's difficulty. We then sort all images by difficulty and get each model's log error (using Eq. 4) on each test example. Finally, we plot each model's performance on the sorted test set examples from the easiest (0% difficulty) to the most difficult (100% difficulty) percentage of the test set examples. The plot shows RCN performs better than SumNet, especially on the harder examples. The occlusion pre-processing helps RCN on most difficult examples (difficulty $>$ 65%), while it slightly helps SumNet.