

Semantic Video Segmentation by Gated Recurrent Flow Propagation

Supplementary Material

David Nilsson¹ and Cristian Sminchisescu^{1,2}

¹Department of Mathematics, Faculty of Engineering, Lund University

²Institute of Mathematics of the Romanian Academy

{david.nilsson, cristian.sminchisescu}@math.lth.se

1. Video Files

We provide several videos illustrating the qualitative benefits of our semantic video segmentation method (ability to temporally segment elongated structures e.g. light poles, considerably less flickering) compared to the single frame semantic video segmentation baseline it is built upon and refines. For the 3 example videos the files we include the video on the left is the input video we want to segment, the middle segmentation is the baseline we used, the LRR-4x network, and to the right we can see the segmentation our method GRFP produced, GRFP(LRR-4x, FlowNet2).

See table 7 in the main paper where we show quantitatively that the temporal consistency is higher for GRFP(LRR-4x, FlowNet2) than for the LRR-4x baseline.