

Combining Cues for Boundary Detection Using the “Mixture of Specialists” Model

Chunhong Zhou¹ and Bartlett W. Mel^{1,2}

¹Department of Biomedical Engineering

²Neuroscience Graduate Program

University of Southern California

Biological visual systems are highly adept at combining cues to improve the sensitivity and selectivity of target detection. For example, depth, motion, intensity, color, and texture cues can together increase the reliability of boundary detection in natural scenes where camouflage abounds. But given a set of single-cue boundary detectors, how should their outputs be combined to arrive at overall boundary probability? Previous work has often assumed that cues are gaussian and class-conditionally (CC) independent; this leads to a linear cue combination rule with weights depending on individual cue reliabilities (Jacobs 1995, Yuille and Bulthoff 1996). However, the physical processes underlying edge formation in natural images may lead to exponential rather than gaussian distributions (Balboa and Grzywacz, 2003), and higher-order dependencies between edge detectors abound (Schwartz & Simoncelli, 2001). We studied cue combination using co-localized oriented edge detectors operating on red-green and blue-yellow opponent channels. Using the Martin-Fowlkes-Malik (2004) human-segmented database as ground truth, we found that the CC joint distributions of co-localized color edge detectors showed strong statistical dependencies, as expected. The dependencies were well explained by a generative model in which two independent, exponentially distributed raw edge values were multiplied by a third exponentially distributed common factor due to local lighting conditions. The modulated edge values were then passed through a saturating nonlinearity reflecting the range compression seen at all levels of the visual system. We derived a softmax-like divisive normalization scheme from the generative model, leading to a cue-combination rule that ranged from linear to MAX-like depending on joint cue values. We discuss the extensibility of the “mixture of specialists” framework to high-dimensional cue-combination problems, and the possible neural substrate for the scheme’s underlying computations.

This work was supported by ONR, NSF, ARO, and NIH.