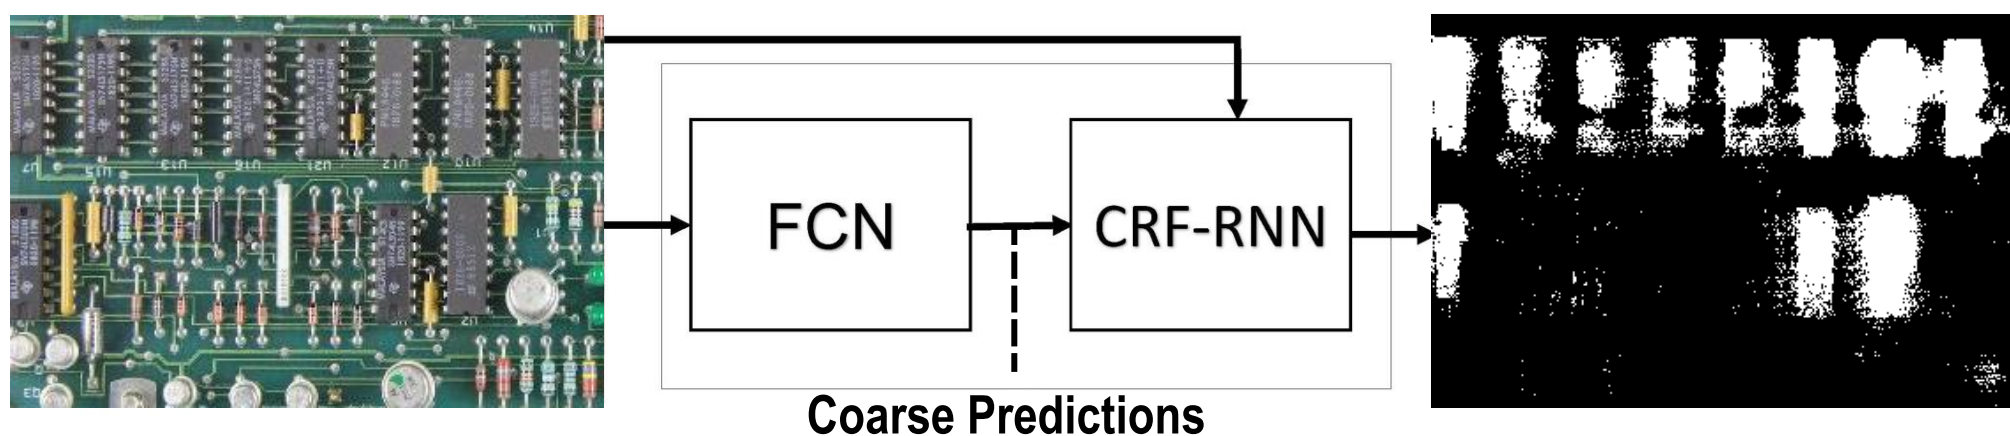


Improving Semantic Segmentation Results using a CRF-RNN

By Brendan Crabb and Chris Ward

The conditional random field as a recurrent neural network (CRF-RNN) adds a conditional random field layer to a fully convolutional neural network (FCN-8) architecture ¹

Schematic visualization of the CRF-RNN network:



The CRF layer provides an image-dependent smoothing term, which encourages assigning similar labels to pixels with similar properties

The CRF-RNN network was trained end-to-end for 100,000 iterations using an NVIDIA GeForce GTX 980 GPU

FCN-8 Results

Pixel Accuracy:	0.8913
Frequency-weighted IU:	0.8265

CRF-RNN Results

Pixel Accuracy:	0.9031
Frequency-weighted IU:	0.8464




Figure 1. Example Input Image



Figure 2. FCN-8 semantic segmentation results




Figure 3. CRF-RNN semantic segmentation results

1. Zheng, Shuai, et al. "Conditional random fields as recurrent neural networks." *Proceedings of the IEEE International Conference on Computer Vision*. 2015.