

Investigating Deep Learning for Face Recognition

Ganesan, Diwakar

This project significantly increased the accuracy of face recognition for biometrics and forensics by testing normalization techniques, distance metrics, and training set sizes for the popular Siamese deep learning algorithm. After creating a baseline network to serve as a frame of reference, I performed a marginal analysis of the hyperparameters of my system. Histogram equalization, a classic method of increasing contrast, increased performance across multiple network configurations, from 44% to 62% on the Good partition of the Good, Bad and Ugly dataset at a false accept rate of 0.1%. The dataset consisted of 1085 images of 437 subjects. Other normalization techniques such as I-chrominance and logarithmic normalization increased accuracy only marginally over the raw RGB images. This verification rate is above the level of baseline algorithms but not at the state of the art, which requires many more images than I had access to in my research. I also found that choice of distance metric (L1, L2, Chi-squared, etc.) is highly dependent on image preprocessing. The Chi-squared method can compensate for a lack of preprocessing, and preprocessing can compensate for a poorer distance metric. Finally, I found the size of the training dataset must be tuned to the depth of the network. Networks with only two layers lost over 4% in accuracy when additional layers were added, whereas by adding just one additional convolutional layer the losses were recouped. Together, the marginal and dependency analyses provide critical insights into deep learning systems to allow for maximum accuracy and efficiency.