

Modeling Color Constancy with Neural Networks

Vlad Cardei, Brian Funt and Kobus Barnard

vcardei@cs.sfu.ca, funt@cs.sfu.ca, kobus@cs.sfu.ca

School of Computing Science, Simon Fraser University, British Columbia, V5A 1S6, Canada

The many algorithms used for color correction make a series of assumptions that try to constrain the problem of finding the scene illuminant under which a given image was taken. In contrast, the neural network we have developed¹ has no explicit constraints. All rules are implicitly learned from the training set, which contains a large number of artificially generated scenes. The network estimates the chromaticity of the illuminant under which the given image was taken. This allows for a diagonal transformation² of the image to another illuminant.

The neural network uses a binarized input of all the chromaticity values found in the image. First, each (R,G,B) pixel in the image is transformed into rg-chromaticity space ($r=R/(R+G+B)$ and $g=G/(R+G+B)$). Second, the chromaticity space is uniformly tessellated into bins and binarized so that the bins have values of either 0 or 1 indicating whether or not the bin's chromaticity range is present in the image. Although this binning has the disadvantage that it discards part of the color resolution, it has a big advantage, which is that it provides a permutation-independent input to the neural net. All variations due to the image geometry are eliminated and only the image colors are used.

The network that we used is a *Perceptron* with two hidden layers. We experimented with many different architectures, but the one that yielded the best results had 3600-200-32-2 nodes. Because of the large size of the network, we added a new *adaptive technique* to the existing network, which shortens the training time by almost an order of magnitude. The performance of the network remained unaffected. With the adaptive reconfiguration technique, the first hidden layer is not fully connected to the input layer. Instead only 150 connections are made to the input layer from any node in the first hidden layer. Initially, these connections are distributed at random. Since the gamut of all possible colors does not occupy the whole input space, the input space initially is not used efficiently – there will be many input neurons that will never receive any activation because their inputs map to colors outside the gamut. The adaptive technique consists of deleting those links to the input layer that were never active during one training epoch and replacing them with new links created at random. This process stabilizes after only three or four epochs, at which time all input links point to active areas in the input space.

The network was trained using Back-propagation without momentum. Different learning rates were used for each layer, which improved the training speed and stability. The learning rates were 50 for the first hidden layer, 10 for the second one and 0.1 for the output layer. The Euclidean distance in the chromaticity space between the target output (r, g) and the estimated one (r_e, g_e) defined the error function. The network was trained with a large number of automatically generated scenes, each with a random number of colors ranging from 2 to 80. The database of illuminants consists of 89 different natural illuminants that were measured with a spectroradiometer. The database of reflectances (surfaces) is composed of 260 surface reflectance functions. For each illuminant, the number of scenes ranged from 100 to 1000. Noise and specular reflectance were also modeled, in order to improve the performance when applied to real images. The network performed much better than the other conventional color constancy algorithms.

1. Funt, B., Cardei, V. and Barnard, K., "Learning Color Constancy", *Proc. Fourth IS&T/SID Color Imaging Conf.*, pp. 58-60, Scottsdale, Nov. 19-22, 1996
2. Finlayson, G., Drew, M., and Funt., B., "Color Constancy: Generalized Diagonal Transforms Suffice", *J. Opt. Soc. Am. A*, 11(11), pp. 3011-3020, 1994