MIT CSAIL

6.869 Advances in Computer Vision

Spring 2011

Problem Set 4 Hints

This document summarizes some frequently asked questions we got from students regarding problem set 4.

1. Recall that we would like to compute the posterior distribution $P(x^i|y)$ over illuminant coefficients x^i and LMS observations y. With the simple color model we use, and assuming independent observations, the posterior factorizes as

$$P(x^i|y) \propto P(x^i) \prod_j P(y_j|x^i, x_j^s) P(x_j^s) \tag{1}$$

In practice, evaluating this expression for any particular illuminant x^i requires performing a large number of multiplications – in the order of the number of observations. For the images in the contest, this will result in an order of 10^5 multiplications of probability values (in range [0,1]). This calculation can quickly underflow and requires careful treatment.

One option is to work in *log-probabilities*. Our MAP estimate, for example, is defined as

$$x^i_{\text{MAP}} = \arg\max_{x^i} P(x^i|y) \tag{2}$$

for a given observed LMS image y. We can take the log of the RHS without changing the maximization

$$\begin{aligned} x^i_{\text{MAP}} &= \arg\max_{x^i} P(x^i|y) \\ &= \arg\max_{x^i} \ln P(x^i|y) \\ &= \arg\max_{x^i} \ln P(x^i) + \sum_j \ln P(y_j|x^i, x_j^s) + \ln P(x_j^s) \end{aligned} \tag{3}$$

and the multiplications are changed to summations, which are more numerically stable. A similar method can be used for the MMSE estimation, were we can convert back to probabilities for computing the weighted average.

2. For the palette images (images 1-5 and the calibration image), we might throw away illuminants due to negative powers in the surface spectra on the borders between the patches (see a previous email about this in the class mailing list). We do not want these observations to influence our algorithm or clutter our estimation. One possible solution is to use only a center patch from each surface patch in those images, as shown below. We will still have many observations to work with, and we do not have to worry about this issue in the other parts of the algorithm.

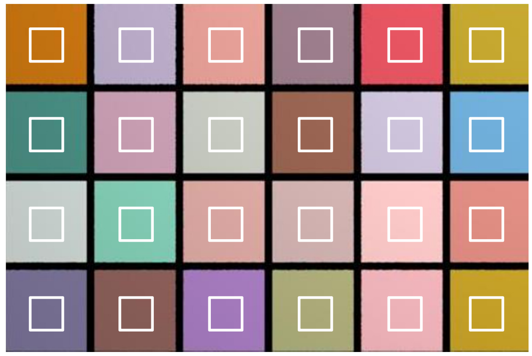

3. The time it takes to evaluate each illuminant estimation might be substantial (depending on your implementation). For either the MAP or MMSE estimates, we'd like to explore as much as possible of the huge space of possible illuminants (essentially $\mathbb{R}^3$). The above method will already improve the performance by some large factor. However, since we assume independence among the measurements, we might also consider *clustering* the observed LMS coordinates and evaluating our estimates on the cluster centers. This will help us avoid recurring evaluations of similar LMS values.
   Another option is to limit your search volume to 1-2 standard deviations around the mean for each coefficient, although this makes more sense if we are confident that the training data is representative of the test data. This need not necessarily be the case in this contest as we are unaware of the method by which the illuminants are selected.