

MIT OpenCourseWare
<http://ocw.mit.edu>

9.35 Sensation And Perception
Spring 2009

For information about citing these materials or our Terms of Use, visit: <http://ocw.mit.edu/terms>.

1. Without good correspondence across pose, size, lighting, etc, pixel-differences between images of the same person can be as large (or larger) than differences between different people. An algorithm like Retinex might do well anyways, since it attempts to estimate the illuminant and could then make a judgment based solely on estimated reflectance.
2. Here's my image distance function:

```
function out = imageDistance(im1,im2)
squaredDiff = (im1-im2).^2;
sumsquaredDiff = sum(sum(squaredDiff));
out = sqrt(sumsquaredDiff);
```

3. No "right" answer here.
4. Here are the rankings I get with my code: [24 12 2 15 17]
5. Normalization can be done by just looping through each face, taking its mean value, taking its variance, and subtracting and dividing respectively:

```
For i=1:40,i,
tmp=faces(:, :, i);
normfaces(:, :, i)=(faces(:, :, i)-mean(tmp(:)))/var(tmp(:));
end
```

And here's the new rankings I get: [24 12 11 27 8]; The differences in the ranking (both in the top 5 and elsewhere) suggest that cues like the mean luminance and global contrast (which are pretty 'dumb' cues) have a profound effect on recognition performance. I leave it to you to decide whether the different rankings mean that "dumb" cues are good or bad things to listen to.