What's your type? Personalized Prediction of Facial Attractiveness

Authors: Sam Crognale, Danish Shabbir

Problem Overview:

Models that attempt to explain universal beauty by the way of symmetry, golden ratios, or measured placement of various conventional features on the face fall short in explaining the varied attraction that is actually witnessed in the world.

In this investigation, we devise an application to give a user some insight about their 'type', i.e. the faces they are attracted to. As users swipe yes or no on a large set of face images, we use dimensionality reduction and pattern recognition on the obtained labelled data to construct an image that gives the user a visual sense of what they find attractive. This can then be used to predict whether a given user will find a previously unseen face attractive.

Data

We began our investigation with MUCT [1] — a relatively small database of face images, where each face is labelled with 68 fiducial points. The faces shown in the dataset were not demographically comprehensive and hence it proved difficult to collect suitable data regarding what faces a variety of users would be drawn to.
Recently we obtained access to the 10K Adult Faces Dataset [2]— a much larger and more comprehensive data set complete with demographic labels for each image.

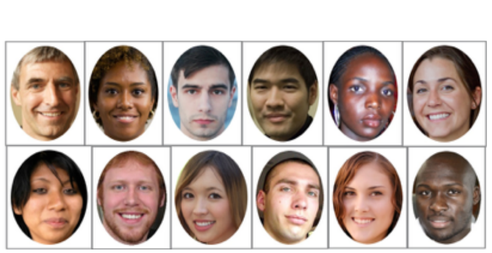

Labelling

We created a web interface with Flask to serve images from the 10K Adult Faces dataset and record a user's associated {yes}/{no} preference. We got each of our users' to label around 2000 - 3000 images in this way. Images are served at random from the dataset and the user is required to label each image as it attracts them subjectively.

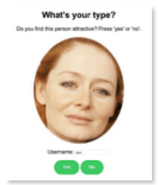

Preprocessing

We obtained a training set of images from the 10K dataset by cropping or padding the images as required so that they are uniformly sized. Once grayscaled, each image is normalized by subtracting from the mean.

Feature extraction with Eigenfaces:

It is possible find an information rich encoding of a face image that best captures the significant local / global "features" of the face, albeit these features do not necessarily correspond to usual facial features like the eyes, nose, lips, hair and so on. Such encoding can be found by capturing the variation in a collection of face images. Mathematically this translates to finding the principle components of a distribution of faces in the dataset which is achieved by finding eigenvectors of the covariance matrix of the set of face images where each image is treated as a point in a high dimensional space.

These eigenvectors — appearing below as ghostly "Eigenfaces"— characterize the variation between the faces in the 10K dataset. It follows that each individual face in the 10K dataset can be reconstructed as a linear combination of these eigenfaces.

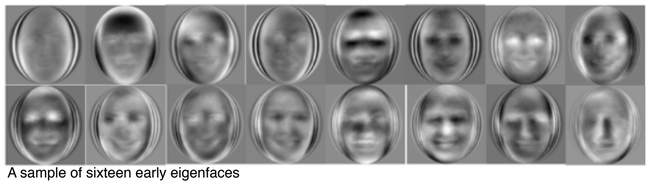

A sample of sixteen early eigenfaces

A sample of the middle-late fourteen eigenfaces

Prediction

To predict whether a given user will like a previously unseen face, we find a labelled face image with the encoding closest to the face in question. We can then predict whether the user will like this face given how they rated its nearest neighbors. These K-nearest neighbors are found using euclidean and mahalanobis distances.

Results

Reconstructed images that represent the ideal 'type' for our users given how they labelled the many varied faces in our data set. This construction is a linear combination of the eigenfaces where the coefficients are calculated by taking the average of all the weights associated with images that a user liked.

Samantha's type | Nate's type | Dan's type | Kate's | Benjamin's type

Next, for a given user, we find the form of a gaussian probability density function on the images they liked using Kernel Density Estimation. A random sample is taken from this distribution to reconstruct an image that would represent another perspective on their type.

Sources:
[1] Milborrow and J. Morkel and F. Nicolls, The MUCT Landmarked Face Database, Pattern Recognition Association of South Africa, 2010
[2] Bainbridge, W.A., Isola, P., & Oliva, A. (2013). The intrinsic memorability of face images. Journal of Experimental Psychology: General. Journal of Experimental Psychology: General, 142(4)
[3] M. Turk and A. Pentland, "Face Recognition Using Eigenfaces,"Proc. IEEE Conf. on Computer Vision and Pattern Recognition, 1991, pp. 586-591.
[4] M. Turk and A. Pentland, "Eigenfaces for Recognition," J. Cognitive Neuroscience, vol. 3, no. 1, 1991