

3.3 Facial recognition.

We have now established that it is possible to represent a face in terms of a relatively small number of eigenfaces. All the information about an individual face is contained in its eigenface expansion coefficients. Moreover, these coefficients have a very specific meaning namely, each coefficient determines the ‘presence’ of its associated eigenface in our individual. A comparison of the expansion coefficients should therefore allow a reliable comparison of faces.

Let us assume that we have collected a database of faces (independent of our training set) and have these faces represented in terms of the eigenfaces. Our goal is to match a new face with one of the faces in the database, i.e. we need to identify an individual in the database from an image provided to us. The first step is to represent the new image in terms of the eigenfaces, using (3.12).

This is not as straightforward as it sounds; again we have to ignore important detail. We may not have much control over the way the new image is collected. We may therefore be dealing with an image taken under very different lighting conditions, at a different scale, rotation and location within the image, against a different background, all of which may be confusing to our system. It is simply not feasible to do the necessary normalization by hand and commercial systems require sophisticated image processing tools to do the necessary automated normalization. We assume that the necessary normalization has already been done.

We do the most naive thing possible and calculate the L^2 norm of the difference between the the given face and all the faces in the database. Because it is not necessarily true that the given face is in the database (we may be dealing with an intruder), we cannot simply take the best match. The best match also needs to be better than a preset threshold. Let us illustrate this idea with the individual shown in Figure 3.11 (this image is not in the database). The best match from the

database is shown in Figure 3.12. Although the two images differ slightly, they are clearly of the same individual.



Figure 3.11: The given facial image.




Figure 3.12: The best match the database.

Let us also compare the the eigenface expansion coefficients directly as in Figure 3.13.




Figure 3.13: The expansion coefficients of the two images of Figures 3.11 and 3.12.

From this figure it is clear that the expansion coefficients of the two images match very closely, especially for low order coefficients.

The problem with using the L^2 norm for the comparison is that it attaches the same weight to all the different eigenfaces. This is not the case in practice. From Figure 3.4 we notice that the lower order eigenfaces contain general face-like features and not individual characteristics.

They are very good in distinguishing faces from other images, but are not as effective in distinguishing between different faces. The higher order eigenfaces, contain more detailed individual characteristics and are more suitable in identifying individuals. There is therefore much to be gained by weighing the expansion coefficients, attaching more weight to the higher order coefficients. Again the details are best pursued elsewhere.