

Linear Subspaces - Geometry

No Invariants, so Capture Variation

- Each image = a pt. in a high-dimensional space.
 - Image: Each pixel a dimension.
 - Point set: Each coordinate of each pt. A dimension.
- Simplest rep. of variation is linear.
 - Basis (eigen) images: $x_1 \dots x_k$
 - Each image, $x = a_1 x_1 + \dots + a_k x_k$
- Useful if $k \ll n$.

When is this accurate?

- Approximately right when:
 - Variation approximately linear. Always true for small variation.
 - Some variations big, some small, can discard small.
- Exactly right sometimes.
 - Point features with scaled-orthographic projection.
 - Convex, Lambertian objects and distant lights.

Principal Components Analysis (PCA)

- All-purpose linear approximation.
- Given images (as vectors)
- Finds low-dimensional linear subspace that best approximates them.
 - Eg., minimizes distance from images to subspace.

What is PCA good for?

- I can measure how *face-like* a new image is by measuring its distance to a linear face space.
 - Project I onto face space. $a_i = \langle I, x_i \rangle$.
 - This is nearest face-point to I .
 - Measure distance: $I = \sum(a_i x_i)$.
- I can speed up comparison between faces by projecting them onto face space (Eigenfaces). This is a big win if I preprocess many faces.
- I can regularize a noisy image.
- Meaningful correspondences critical.

Derivation on whiteboard

- This is all taken from Duda, Hart and Stork *Pattern Classification* pp. 114-117. Excerpt in library.

PCA derivation

(this is all just taken from Duda, Hart and Stork)

Suppose we have a series of vectors, $x_1 \dots x_n$, and we want to approximate them with a low-dimensional subspace. What is the best way to do this? If we want to approximate them with a 0 dimensional subspace, we can do this most accurately by approximating them by their mean, m . This is probably intuitive, but if not, Duda, Hart and Stork have a very nice proof (Eq. 80, p. 115).

Next we'll consider find the best 1-dimensional subspace, written as: x_i is approximated by $m + a_i e$, where e is a unit vector indicating the direction of the space. Then our goal is to choose a_i and e to minimize:

$$\begin{aligned} J(a_1, \dots, a_n, e) &= \sum \|(m + a_i e) - x_i\|^2 \\ &= \sum \|a_i e - (x_i - m)\|^2 \\ &= \sum a_i^2 \|e\|^2 - 2 \sum a_i e(x_i - m) + \sum \|x_i - m\|^2 \\ \|e\| &= 1. \text{ Taking the derivatives w.r.t. } a_i \text{ and setting them to 0 we get:} \\ 2a_i - 2 e(x_i - m) &= 0, \\ a_i &= e(x_i - m). \end{aligned}$$

We can skip this derivation, and just say that of course we get the best choice of a_i by projecting $x_i - m$ onto e .

Now, if we set $a_i = e(x_i - m)$, we get J as a function of e

$$\begin{aligned} J(e) &= \sum a_i^2 - 2 \sum a_i^2 + \sum \|x_i - m\|^2 \\ &= - \sum [e(x_i - m)]^2 + \sum \|x_i - m\|^2 \\ &= - \sum e(x_i - m)(x_i - m)e + \sum \|x_i - m\|^2 \end{aligned}$$


So we need to maximize eSe subject to $\|e\| = 1$.

We do this with Lagrange multipliers. We set:

$U = eSe - \lambda(ee - 1)$, differentiate w.r.t. e and set this to 0. We get:

$\partial U / \partial e = 2Se - 2\lambda e$, so $Se = \lambda e$. So e is an eigenvector of S , and we can see that eSe is maximized when e is the eigenvector associated with the largest eigenvalue.

Fisher Linear Discriminant



LDA for comparison

- Capture variations that distinguish different objects, throw away variations that don't.

Suppose we have two classes. We project all points onto the direction w . Let the means of the classes be m_1 and m_2 , and their projection onto w be m_1' and m_2' . If x are the points, and y are their projection, we the variance of each class, after the projection is

$$s_i^2 = \sum (y - m_i')^2$$

We will maximize $J(w) = (m_1' - m_2')^2 / (s_1'^2 + s_2'^2)$.

That is, we maximize the separation in the means relative to the variance within the classes.

Define S_1 to be the scatter matrix for points in class 1, similarly define S_2 , and $S_w = S_1 + S_2$.

$$s_i'^2 = \sum (w x - w m_i)^2 = \sum w (x - m_i)(x - m_i) w = w S_i w.$$

$$s_1'^2 + s_2'^2 = w S_w w$$

$$\text{Similarly, } (m_1' - m_2')^2 = w S_b w.$$

$J(w) = w S_b w / w S_w w$. This is the generalized Rayleigh quotient problem. By taking the partials w.r.t. w and setting them to 0, we can show this is solved when:

$S_b w = \lambda S_w w$, a generalized eigenvalue problem.

$S_w^{-1} S_b w = \lambda w$. $S_b w$ is in direction $m_2 - m_1$. So $w = S_w^{-1} (m_2 - m_1)$.

SVD

- Scatter matrix can be big, so computation non-trivial.
- Stack data into matrix X , each row an image. SVD gives $X = U D V^T$
 - D is diagonal with non-increasing values.
 - U and V have orthonormal rows.
- $V^T(:, 1:k)$ gives first k principal components.
- *matlab*

Linear Combinations

$$\begin{array}{c}
 \begin{pmatrix} u_1^1 & u_2^1 & \cdot & \cdot & \cdot & u_n^1 \\ v_1^1 & v_2^1 & & & & v_n^1 \\ u_1^2 & u_2^2 & & & & u_n^2 \\ v_1^2 & v_2^2 & & & & v_n^2 \\ \cdot & \cdot & & & & \cdot \\ \cdot & \cdot & & & & \cdot \\ \cdot & \cdot & & & & \cdot \\ u_1^m & u_2^m & & & & u_n^m \\ v_1^m & v_2^m & \cdot & \cdot & \cdot & v_n^m \end{pmatrix} \\
 \text{I}
 \end{array}
 =
 \begin{array}{c}
 \begin{pmatrix} s_{1,1}^1 & s_{1,2}^1 & s_{1,3}^1 & t_x^1 \\ s_{2,1}^1 & s_{2,2}^1 & s_{2,3}^1 & t_y^1 \\ s_{1,1}^2 & s_{1,2}^2 & s_{1,3}^2 & t_x^2 \\ s_{2,1}^2 & s_{2,2}^2 & s_{2,3}^2 & t_y^2 \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ s_{1,1}^m & s_{1,2}^m & s_{1,3}^m & t_x^m \\ s_{2,1}^m & s_{2,2}^m & s_{2,3}^m & t_y^m \end{pmatrix} \\
 \text{S}
 \end{array}
 \begin{array}{c}
 \begin{pmatrix} x_1 & x_2 & \cdot & \cdot & \cdot & x_n \\ y_1 & y_2 & & & & y_n \\ z_1 & z_2 & & & & z_n \\ 1 & 1 & & & & 1 \end{pmatrix} \\
 \text{P}
 \end{array}$$

Immediately apparent that u and v coordinates lie in a 4D linear subspace