Basics: Orthogonal Projections

Archer Yang

McGill University

October 4, 2024

- **Reading assignment:** read chapter 3.1-3.8.2 of Mathematics for Machine Learning book.

Dimension reduction

- High-dimensional data often has only **a few** dimensions contain most information.
- When compress or visualize high-dimensional data, pick the **most informative** dimensions in the data.
- Then project the original high-dimensional data onto this lower-dimensional space.
- E.g. **principal component analysis** (PCA) by Hotelling (1933) and **deep auto-encoders** (Deng et al., 2010).
- Both techniques involves **orthogonal projection**, so let's study this concept.

Dimension reduction

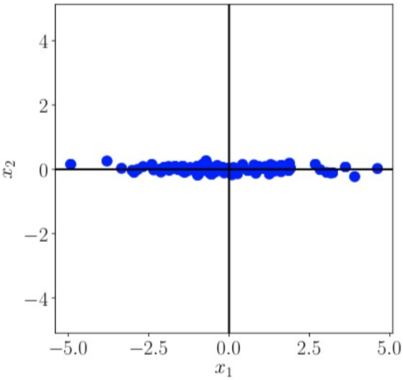

(a) Dataset with x_1 and x_2 coordinates.

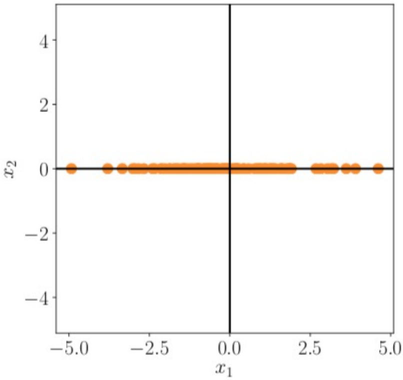

(b) Compressed dataset where only the x_1 coordinate is relevant.

Projection

Definition (Projection)

Let V be a vector space and $U \subseteq V$ a subspace of V.

A linear mapping $\pi : V \to U$ is called a **projection** if

$$\pi^2 = \pi \circ \pi = \pi$$

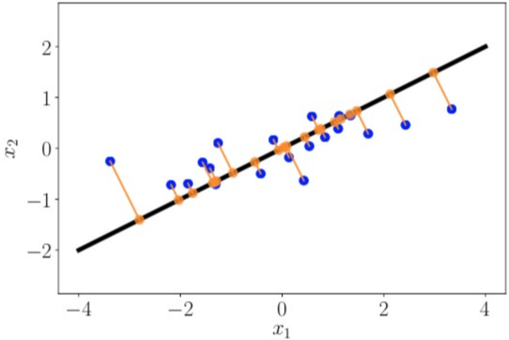

Projection onto One-dimensional Subspaces

Projection onto one-dimensional subspaces (lines)

- Given a line through the origin with basis vector $\boldsymbol{b} = (b_1, \cdots, b_D)^\top \in \mathbb{R}^D$.
- Note that $\boldsymbol{b}$ is a **column vector**.
- The line through the origin creates a one-dimensional subspace U **spanned by** $\boldsymbol{b}$

$$U = \{z\boldsymbol{b} : \forall z \in \mathbb{R}\}$$

- Now we project $\boldsymbol{x} = (x_1, \cdots, x_D)^\top \in \mathbb{R}^D$ onto U, denote this **projection** as $\pi_U(\boldsymbol{x})$

Projection onto one-dimensional subspaces (lines)

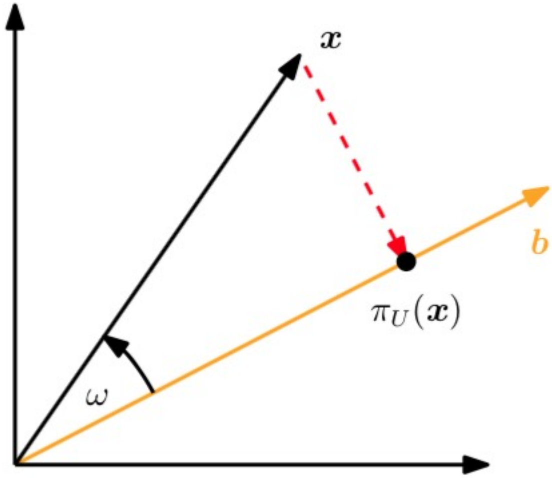

Projection onto one-dimensional subspaces (lines)

- The projection $\pi_U(\boldsymbol{x})$ of $\boldsymbol{x}$ onto U **must be an element** of U, i.e.

$$\pi_U(\boldsymbol{x}) = z\boldsymbol{b}$$

  for **some** $z \in \mathbb{R}$.

- How to determine the coordinate z?

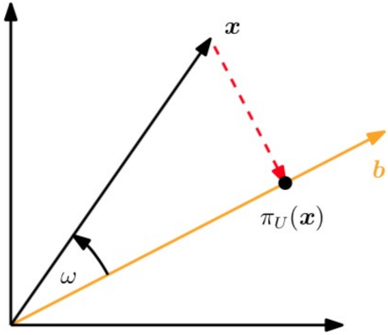

Finding the coordinate z

It's very easy:

- The orthogonality condition yields

$$\langle \boldsymbol{x} - \pi_U(\boldsymbol{x}), \boldsymbol{b} \rangle = 0 \overset{\pi_U(\boldsymbol{x}) = z\boldsymbol{b}}{\iff} \langle \boldsymbol{x} - z\boldsymbol{b}, \boldsymbol{b} \rangle = 0.$$

- By the bilinearity of the inner product

$$\langle \boldsymbol{x}, \boldsymbol{b} \rangle - z\langle \boldsymbol{b}, \boldsymbol{b} \rangle = 0 \iff z = \frac{\langle \boldsymbol{x}, \boldsymbol{b} \rangle}{\langle \boldsymbol{b}, \boldsymbol{b} \rangle} = \frac{\langle \boldsymbol{b}, \boldsymbol{x} \rangle}{||\boldsymbol{b}||^2}$$

- In a Euclidean space, $\langle \cdot, \cdot \rangle$ is the dot product

$$z = \frac{\boldsymbol{b}^\top \boldsymbol{x}}{\boldsymbol{b}^\top \boldsymbol{b}} = \frac{\boldsymbol{b}^\top \boldsymbol{x}}{||\boldsymbol{b}||^2}$$

- As a special case, if $||\boldsymbol{b}|| = 1$, then the **coordinate** z of the projection is given by

$$z = \boldsymbol{b}^\top \boldsymbol{x}$$

Key slides

- If $||\boldsymbol{b}|| = 1$, then the **coordinate** z of the projection is given by

$$z = \boldsymbol{b}^\top \boldsymbol{x}$$

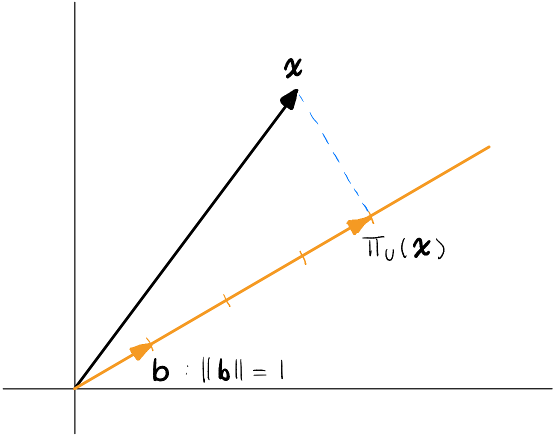

Key slides

- If $||\boldsymbol{b}|| = 1$, then the **coordinate** z of the projection is given by

$$z = \boldsymbol{b}^\top \boldsymbol{x}$$

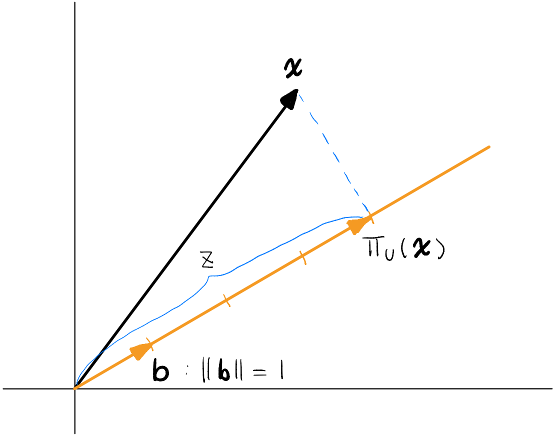

Finding the projection point

- Since $\pi_U(\boldsymbol{x}) = z\boldsymbol{b}$, we get

$$\pi_U(\boldsymbol{x}) = z\boldsymbol{b} = \frac{\langle \boldsymbol{x}, \boldsymbol{b} \rangle}{||\boldsymbol{b}||^2}\boldsymbol{b} = \frac{\boldsymbol{b}^\top \boldsymbol{x}}{||\boldsymbol{b}||^2}\boldsymbol{b}$$

- The length of $\pi_U(\boldsymbol{x})$ is

$$||\pi_U(\boldsymbol{x})|| = ||z\boldsymbol{b}|| = |z|\,||\boldsymbol{b}||$$

- As a special case, when $||\boldsymbol{b}|| = 1$,

$$\pi_U(\boldsymbol{x}) = z\boldsymbol{b} = \boldsymbol{b}\boldsymbol{b}^\top \boldsymbol{x}$$

and $||\pi_U(\boldsymbol{x})|| = |z|$.

Key slides

- If $||\boldsymbol{b}|| = 1$, then the **projection** $\pi_U(\boldsymbol{x})$ is

$$\pi_U(\boldsymbol{x}) = z\boldsymbol{b} = \boldsymbol{b}z = \boldsymbol{b}\boldsymbol{b}^\top \boldsymbol{x}$$

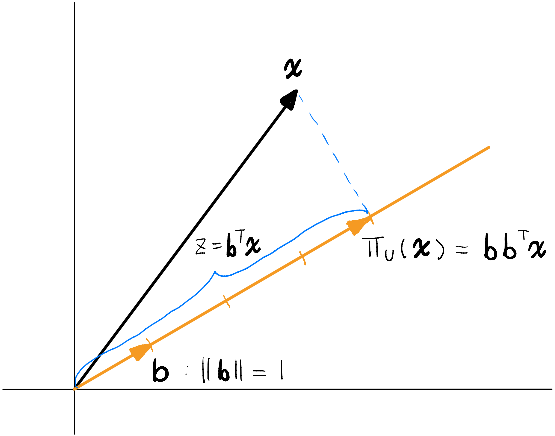

A practice

- (Homework) Given that

$$\boldsymbol{x}^\top \boldsymbol{y} = \|\boldsymbol{x}\| \|\boldsymbol{y}\| \cos \omega$$

show that

$$\|\pi_U(\boldsymbol{x})\| = |\cos \omega| \|\boldsymbol{x}\|$$

- (Homework) If $\|\boldsymbol{x}\| = 1$ (lies on the unit circle), then $\|\pi_U(\boldsymbol{x})\| = |\cos \omega|$

Finding the projection matrix P_π

- The operation of the linear projection π_U can be represented by a projection matrix $\boldsymbol{P}_\pi$, s.t.

$$\pi_U(\boldsymbol{x}) = \boldsymbol{P}_\pi \boldsymbol{x}$$

- Since

$$\pi_U(\boldsymbol{x}) = \frac{\boldsymbol{b}^\top \boldsymbol{x}}{||\boldsymbol{b}||^2}\boldsymbol{b} = \frac{\boldsymbol{b}\boldsymbol{b}^\top \boldsymbol{x}}{||\boldsymbol{b}||^2} = \frac{\boldsymbol{b}\boldsymbol{b}^\top}{||\boldsymbol{b}||^2}\boldsymbol{x}$$

  we see that

$$\boldsymbol{P}_\pi = \frac{\boldsymbol{b}\boldsymbol{b}^\top}{||\boldsymbol{b}||^2} \tag{1}$$

- Note that $\boldsymbol{b}\boldsymbol{b}^\top$ is a symmetric matrix (of rank 1), and $||\boldsymbol{b}||^2$ is a scalar.
- As a special case, if $||\boldsymbol{b}|| = 1$, then $\boldsymbol{P}_\pi = \boldsymbol{b}\boldsymbol{b}^\top$.

A practice

- (Homework) Show that a projection matrix satisfies

$$\boldsymbol{P}_\pi^2 \boldsymbol{x} = \boldsymbol{P}_\pi \boldsymbol{x}$$

this verifies that mapping $\boldsymbol{P}_\pi$ indeed satisfies the definition of projection

Projection onto General Subspaces

Projection onto general subspaces

Now consider orthogonal projections of $\boldsymbol{x} \in \mathbb{R}^D$ onto lower-dimensional subspaces $U \subseteq \mathbb{R}^D$ with $\dim(U) = M \geq 1$.

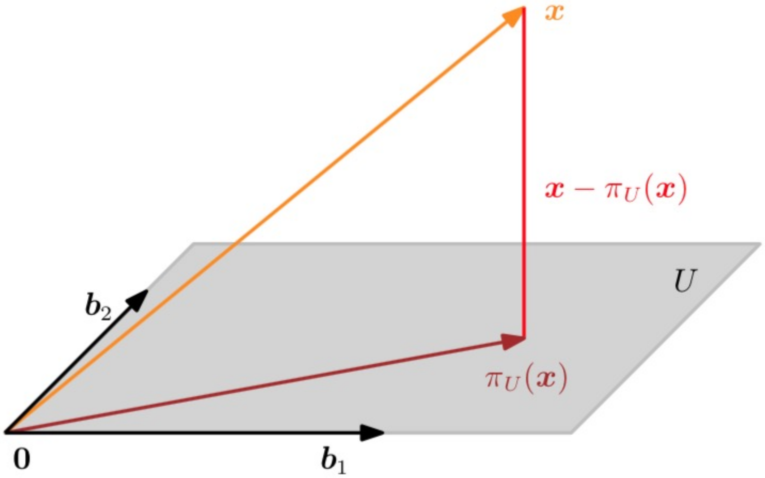

Projection onto general subspaces

- Assume $\boldsymbol{b}_1, \cdots, \boldsymbol{b}_M$ is an ordered basis of U.
- The projection $\pi_U(\boldsymbol{x})$ of $\boldsymbol{x}$ onto U **must be an element** of U, i.e. $\pi_U(\boldsymbol{x})$ can be represented as linear combinations of the basis vectors

$$\boldsymbol{b}_1, \cdots, \boldsymbol{b}_M$$

of U,

$$\pi_U(\boldsymbol{x}) = \sum_{i=1}^{M} z_i \boldsymbol{b}_i$$

- How to determine $z_1, \cdots, z_M$?

Finding the coordinates $z_1, \cdots, z_M$

- First we use the matrix notation

$$\pi_U(\boldsymbol{x}) = \sum_{i=1}^{M} z_i \boldsymbol{b}_i = \boldsymbol{B}\boldsymbol{z}$$

where

$$\boldsymbol{B} = \begin{pmatrix} | & | & & | \\ \boldsymbol{b}_1 & \boldsymbol{b}_2 & \cdots & \boldsymbol{b}_M \\ | & | & & | \end{pmatrix} \in \mathbb{R}^{D \times M}, \qquad \boldsymbol{z} = (z_1, \cdots, z_M)^\top \in \mathbb{R}^M$$

Finding the coordinates $z_1, \cdots, z_M$

- Need to find $\pi_U(\boldsymbol{x})$ in U that is closest to $\boldsymbol{x} \in \mathbb{R}^D$.
- This implies that $\boldsymbol{x} - \pi_U(\boldsymbol{x})$ must be orthogonal to all basis vectors of U:

$$\langle \boldsymbol{b}_1, \boldsymbol{x} - \pi_U(\boldsymbol{x}) \rangle = \boldsymbol{b}_1^\top (\boldsymbol{x} - \pi_U(\boldsymbol{x})) = 0$$

$$\vdots$$

$$\langle \boldsymbol{b}_M, \boldsymbol{x} - \pi_U(\boldsymbol{x}) \rangle = \boldsymbol{b}_M^\top (\boldsymbol{x} - \pi_U(\boldsymbol{x})) = 0$$

Finding the coordinates $z_1, \cdots, z_M$

- With $\pi_U(\boldsymbol{x}) = \boldsymbol{B}\boldsymbol{z}$

$$\boldsymbol{b}_1^\top(\boldsymbol{x} - \boldsymbol{B}\boldsymbol{z}) = 0$$
$$\vdots$$
$$\boldsymbol{b}_M^\top(\boldsymbol{x} - \boldsymbol{B}\boldsymbol{z}) = 0$$

- In a matrix format

$$\begin{bmatrix} \boldsymbol{b}_1^\top \\ \vdots \\ \boldsymbol{b}_M^\top \end{bmatrix} [\boldsymbol{x} - \boldsymbol{B}\boldsymbol{z}] = \boldsymbol{0} \Longleftrightarrow \boldsymbol{B}^\top(\boldsymbol{x} - \boldsymbol{B}\boldsymbol{z}) = \boldsymbol{0}$$

$$\Longleftrightarrow \boldsymbol{B}^\top\boldsymbol{B}\boldsymbol{z} = \boldsymbol{B}^\top\boldsymbol{x}.$$

- Since $\boldsymbol{b}_1, \cdots, \boldsymbol{b}_M$ are a basis of U and therefore linearly independent, therefore

$$\text{rank}(\boldsymbol{B}) = M \qquad \text{where } M \leq D$$

- Meanwhile $\boldsymbol{B}^\top \boldsymbol{B} \in \mathbb{R}^{M \times M}$ is invertible since

$$\text{rank}(\boldsymbol{B}^\top \boldsymbol{B}) = \text{rank}(\boldsymbol{B}) = M \tag{2}$$

we get

$$\boldsymbol{z} = (\boldsymbol{B}^\top \boldsymbol{B})^{-1} \boldsymbol{B}^\top \boldsymbol{x}$$

- The matrix $(\boldsymbol{B}^\top \boldsymbol{B})^{-1} \boldsymbol{B}^\top$ is also called the **pseudo-inverse** of non-square matrix $\boldsymbol{B}$.

Proof of Eq. (2)

- First we see that nullity$(\boldsymbol{B}^\top\boldsymbol{B})$ = nullity$(\boldsymbol{B})$, this is because for any $\boldsymbol{x}$ that satisfies

$$\boldsymbol{B}^\top\boldsymbol{B}\boldsymbol{x} = \mathbf{0}$$

we also have

$$||\boldsymbol{B}\boldsymbol{x}||^2 = \boldsymbol{x}\boldsymbol{B}^\top\boldsymbol{B}\boldsymbol{x} = 0$$

hence $\boldsymbol{B}\boldsymbol{x} = \mathbf{0}$. Therefore nullity$(\boldsymbol{B}^\top\boldsymbol{B})$ = nullity$(\boldsymbol{B})$

- By rank–nullity theorem, we have nullity$(\boldsymbol{B})$ = M − rank$(\boldsymbol{B})$ = 0, then

$$\text{rank}(\boldsymbol{B}^\top\boldsymbol{B}) = M - \text{nullity}(\boldsymbol{B}^\top\boldsymbol{B}) = M - \text{nullity}(\boldsymbol{B}) = M$$

Finding the projection point

- We already know that

$$\pi_U(\boldsymbol{x}) = \boldsymbol{B}\boldsymbol{z}$$

Therefore

$$\pi_U(\boldsymbol{x}) = \boldsymbol{B}(\boldsymbol{B}^\top \boldsymbol{B})^{-1}\boldsymbol{B}^\top \boldsymbol{x}$$

Finding the projection matrix P_π

- We see that the projection matrix that solves

$$\pi_U(\boldsymbol{x}) = \boldsymbol{P}_\pi \boldsymbol{x}$$

must be

$$\boldsymbol{P}_\pi = \boldsymbol{B}(\boldsymbol{B}^\top \boldsymbol{B})^{-1} \boldsymbol{B}^\top$$

- **Remark 1:** this result includes the 1D case as a special case: if $\dim(U) = 1$, then $M = 1$ and $\boldsymbol{B}^\top \boldsymbol{B} \in \mathbb{R}$ is a scalar, we can then write

$$\boldsymbol{P}_\pi = \boldsymbol{B}(\boldsymbol{B}^\top \boldsymbol{B})^{-1} \boldsymbol{B}^\top = \frac{\boldsymbol{B}\boldsymbol{B}^\top}{\boldsymbol{B}^\top \boldsymbol{B}}$$

which coincides with Eq. (1).

Finding the projection matrix P_π

- **Remark 2:** The projections $\pi_U(\boldsymbol{x})$ are still vectors in $\mathbb{R}^D$ although they lie in an M-dimensional subspace $U \subseteq \mathbb{R}^D$. However, to represent a projected vector we only need the M coordinates $z_1, \cdots, z_M$ with respect to the basis vectors $\boldsymbol{b}_1, \cdots, \boldsymbol{b}_M$. i.e. $\pi_U(\boldsymbol{x}) \in \mathbb{R}^D$ however $\boldsymbol{z} = (z_1, \cdots, z_M)^\top \in \mathbb{R}^M$.

- If the basis $\{\boldsymbol{b}_1, \cdots, \boldsymbol{b}_M\}$ for the subspace U is an orthonomal basis, i.e for all $i, j = 1, \cdots, M$

$$\langle \boldsymbol{b}_i, \boldsymbol{b}_j \rangle = \boldsymbol{b}_i^\top \boldsymbol{b}_j = 0 \qquad \text{for } i \neq j$$
$$\langle \boldsymbol{b}_i, \boldsymbol{b}_i \rangle = \boldsymbol{b}_i^\top \boldsymbol{b}_i = 1$$

  Then

$$\boldsymbol{B}^\top \boldsymbol{B} = (\boldsymbol{b}_1, \cdots, \boldsymbol{b}_M)^\top (\boldsymbol{b}_1, \cdots, \boldsymbol{b}_M) = \boldsymbol{I}_{M \times M}$$

- In this case, the coordinate of the orthogonal projection reduces to

$$\boldsymbol{z} = (\boldsymbol{B}^\top \boldsymbol{B})^{-1} \boldsymbol{B}^\top \boldsymbol{x} = \boldsymbol{I}_{M \times M}^{-1} \boldsymbol{B}^\top \boldsymbol{x} = \boldsymbol{B}^\top \boldsymbol{x}$$

  and the **projection** $\pi_U(\boldsymbol{x})$ reduces to

$$\pi_U(\boldsymbol{x}) = \boldsymbol{B}\boldsymbol{z} = \boldsymbol{B}\boldsymbol{B}^\top \boldsymbol{x}$$

  with the projection matrix $\boldsymbol{P}_\pi = \boldsymbol{B}\boldsymbol{B}^\top$