Preliminary: Random Vectors

In ML, our data points are often multidimensional.

For example:

To predict housing prices, each data point might have: number of rooms, number of bathrooms, square footage, zip code, year built, …

To make movie recommendations, each data point might have: ratings of existing movies, whether you started a movie and stopped after 10 minutes,…

A single data point is a full vector

3

Preliminary: Random Vectors

A random vector X is a vector where each entry is a random variable.

$\mathbb{E}[X]$ is a vector, where each entry is the expectation of that entry.

For example, if X is a uniform vector from the sample space

$$\left\{\begin{bmatrix}1\\2\\3\end{bmatrix}, \begin{bmatrix}-1\\2\\3\end{bmatrix}, \begin{bmatrix}0\\2\\6\end{bmatrix}\right\}$$

$\mathbb{E}[X] = [0,2,4]^T$

4

Unequal Variances, Still Independent

Let's think about 2 dimensions.

Let $X = [X_1, X_2]^T$ where $X_1 \sim \mathcal{N}(0,5)$, $X_2 \sim \mathcal{N}(0,1)$ and X_1 and X_2 are independent.

What is Σ? Which of these pictures are i.i.d. samples of X?

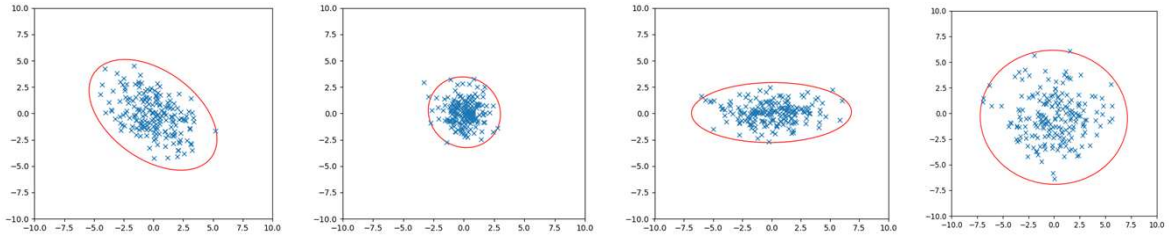

8

Dependence

Let's think about 2 dimensions.

Let $X = [X_1, X_2]^T$ where $\mathrm{Var}(X_1) = 5$, $\mathrm{Var}(X_2) = 7$ BUT X_1 and X_2 are dependent. $\mathrm{Cov}(X_1, X_2) = -2$

What is Σ? Which of these pictures are i.i.d. samples of X?

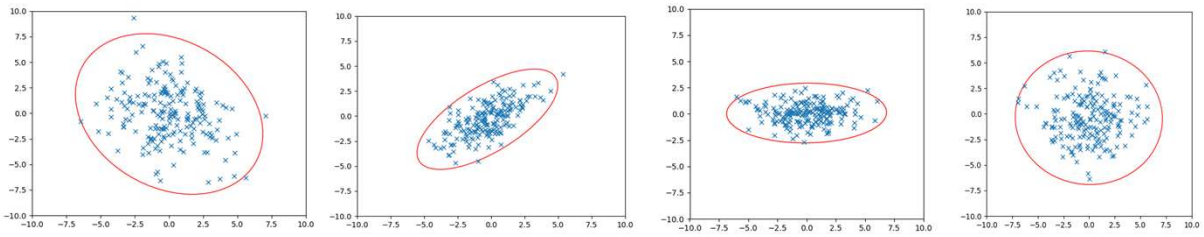

13