CURSE OF DIMENSIONALITY: EXAMPLE / 2

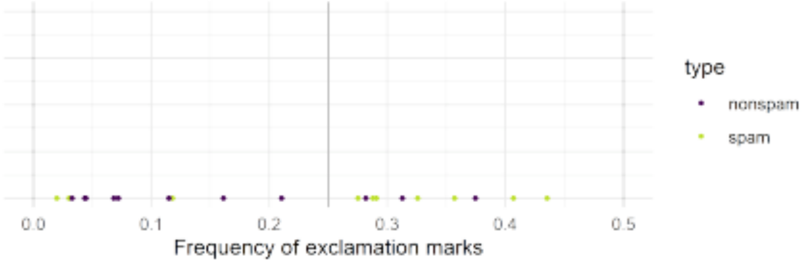

Based on the frequency of exclamation marks, we train a very simple classifier (a decision stump with split point $\mathbf{x} = 0.25$):

- We divide the input space into 2 equally sized regions.
- In the second region $[0.25, 0.5]$, 7 out of 10 are spam.
- Given that at least 0.25% of all letters are exclamation marks, an email is spam with a probability of $\frac{7}{10} = 0.7$.

CURSE OF DIMENSIONALITY: EXAMPLE / 3

Let us feed more information into our classifier. We include a feature that contains the length of the longest sequence of capital letters.

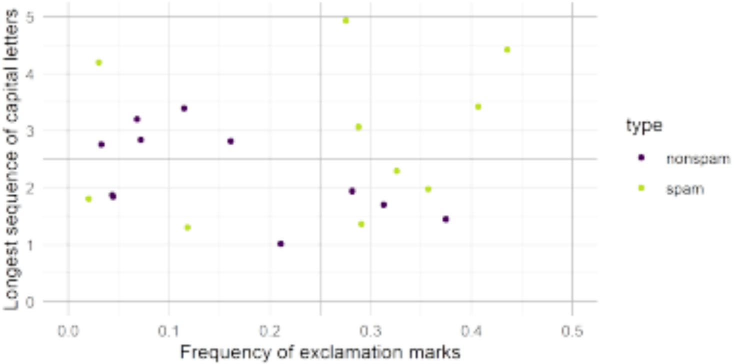

- In the 1D case we had 20 observations across 2 regions.
- The same number is now spread across 4 regions.

CURSE OF DIMENSIONALITY: EXAMPLE / 4

Let us further increase the dimensionality to 3 by using the frequency of the word “your” in an email.

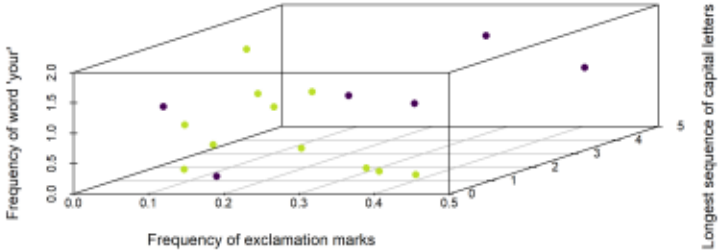

CURSE OF DIMENSIONALITY: EXAMPLE / 5

- When adding a third dimension, the same number of observations is spread across 8 regions.
- In 4 dimensions the data points are spread across 16 cells, in 5 dimensions across 32 cells and so on ...
- As dimensionality increases, the data become **sparse**; some of the cells become empty.
- There might be too few data in each of the blocks to understand the distribution of the data and to model it.

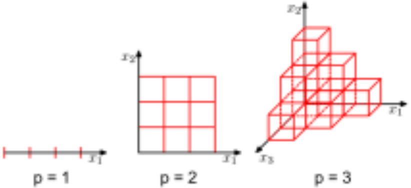

Bishop, Pattern Recognition and Machine Learning, 2006

THE HIGH-DIMENSIONAL CUBE / 2

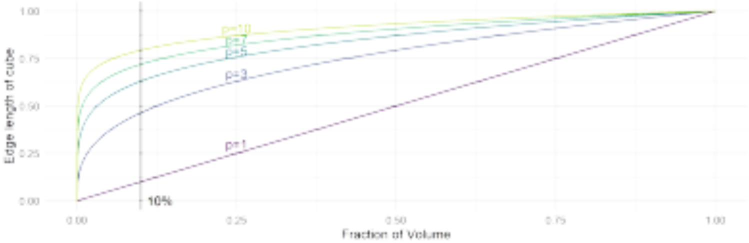

$$a^p = \frac{1}{10} \Leftrightarrow a = \frac{1}{\sqrt[p]{10}}$$

- So: covering 10% of total volume in a cell requires cells with almost 50% of the entire range in 3 dimensions, 80% in 10 dimensions.

THE HIGH-DIMENSIONAL SPHERE / 2

Consider a 20-dimensional sphere. Nearly all of the volume lies in its outer shell of thickness 0.2:

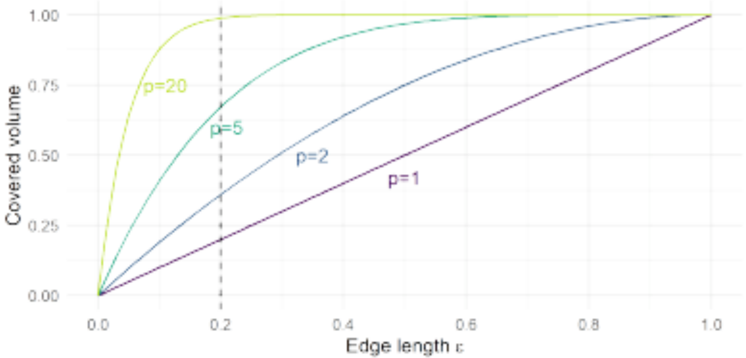

HYPHERSPHERE WITHIN HYPERCUBE / 2

Consider a 10-dimensional sphere inscribed in a 10-dimensional cube.
Nearly all of the volume lies in the corners of the cube:

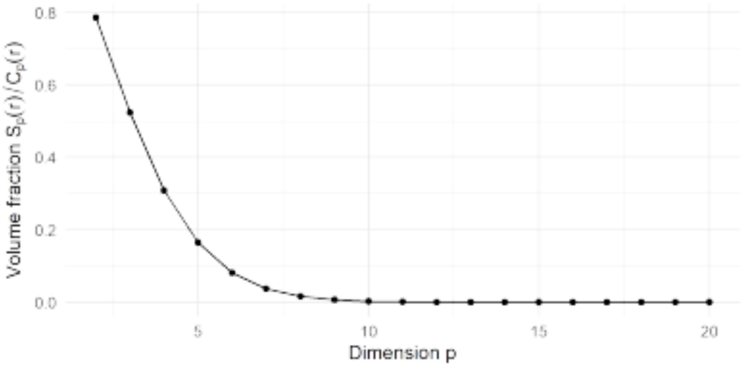

Note: For $r > 0$, the volume fraction $\frac{S_p(r)}{C_p(r)}$ is independent of r.

GAUSSIANS IN HIGH DIMENSIONS

A further manifestation of the **curse of dimensionality** appears if we consider a standard Gaussian $N_p(\mathbf{0}, \boldsymbol{I}_p)$ in p dimensions.

- After transforming from Cartesian to polar coordinates and integrating out the directional variables, we obtain an expression for the density $p(r)$ as a function of the radius r (i.e., the point's distance from the origin), s.t.

$$p(r) = \frac{S_p r^{p-1}}{(2\pi\sigma^2)^{p/2}} \exp\left(-\frac{r^2}{2\sigma^2}\right),$$

  where S_p is the surface area of the p-dimensional unit hypersphere.
- Thus $p(r)\delta r$ is the approximate probability mass inside a thin shell of thickness δr located at radius r.

GAUSSIANS IN HIGH DIMENSIONS / 2

- To verify this functional relationship empirically, we draw 10^4 points from the p-dimensional standard normal distribution and plot $p(r)$ over the histogram of the points' distances to the origin:

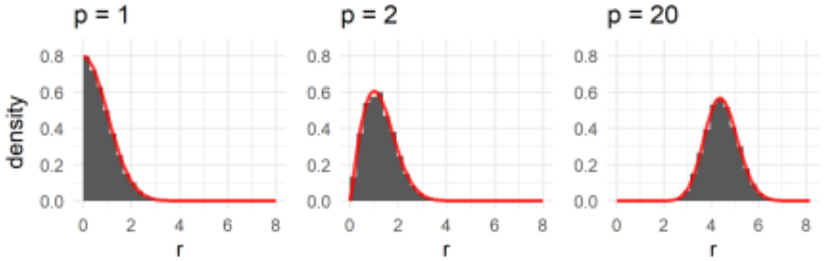

- We can see that for large p the probability mass of the Gaussian is concentrated in a fairly thin "shell" rather far away from the origin. This may seem counterintuitive, but:

GAUSSIANS IN HIGH DIMENSIONS / 3

- For the probability mass of a hyperspherical shell it follows that

$$\int_{r-\frac{\delta r}{2}}^{r+\frac{\delta r}{2}} p(\tilde{r}) d\tilde{r} = \int_{r-\frac{\delta r}{2} \leq ||\mathbf{x}||_2 \leq r+\frac{\delta r}{2}} f_p(\tilde{\mathbf{x}}) d\tilde{\mathbf{x}},$$

where $f_p(\mathbf{x})$ is the density of the p-dimensional standard normal distribution and $p(r)$ the associated radial density.

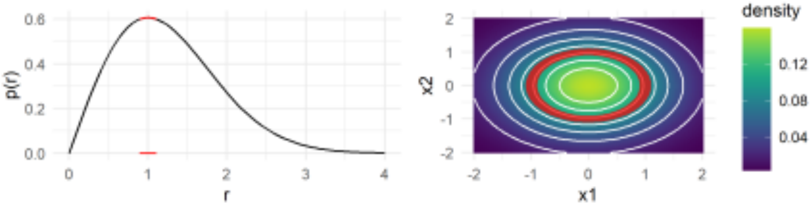

Example: 2D normal distribution

- While f_p becomes smaller with increasing r, the region of the integral -the hyperspherical shell- becomes bigger.