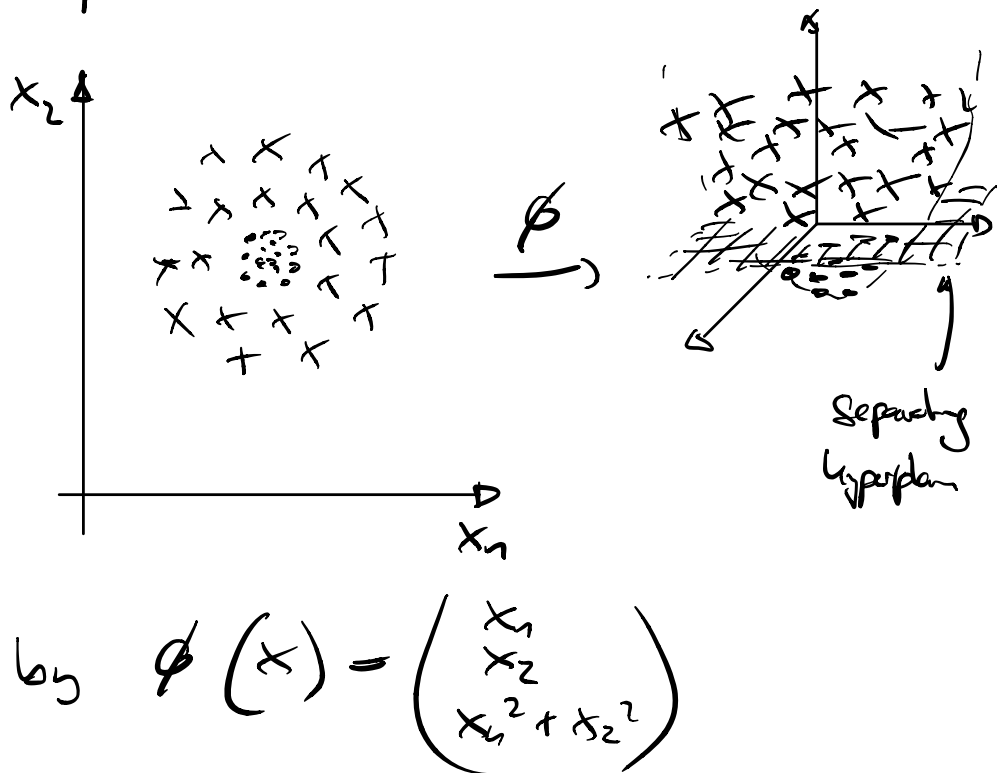


First Non-Linear Classification

A first trick to extend our linear classifiers to non-linear ones is to map the training data points $x^{(i)}$ to a higher dimensional space where it becomes linearly separable



Hence, the activation becomes

$$z(x) = w \cdot \begin{pmatrix} 1 \\ \Phi(x) \end{pmatrix}$$

for a suitable non-linear function

$$\Phi: \mathbb{R}^n \rightarrow \mathbb{R}^k$$

and the loss function

$$L(w) = \frac{1}{n} \sum_{i=1}^n \ell(y^{(i)}, z(x^{(i)}))$$

and its gradient

$$\frac{\partial L(w)}{\partial w} = \frac{1}{n} \sum_{i=1}^n \frac{\partial \ell(y^{(i)}, z)}{\partial z} \Big|_{z=z(x^{(i)})} z'(x^{(i)})$$