

Approximation of functions

A multi layer perceptron trained through BPA can be an engine for learning the non-linear I/O mappings

$$f: \mathbb{R}^{m_0} \longrightarrow \mathbb{R}^{m_L}$$

↑ Space of I/p neurons

↑ Space of o/p neurons

Qn: What is the min # of hidden layers needed for a MLP to learn the I/O mapping?

Universal Approximation Theorem (Key: 1 single layer of hidden neurons will suffice)

Let $\varphi(\cdot)$ be a non-constant, bounded and monotone non-decreasing continuous function. Let I_{m_0} denote a m_0 -dim. unit hypercube $[0, 1]^{m_0}$. The space of cont. functions on I_{m_0} is denoted by $C(I_{m_0})$. Then, given any function $f \in C(I_{m_0})$ and $\varepsilon > 0$ \exists an integer m_1 , \exists sets of constants α_i, b_i, w_{ij} $i = 1, \dots, m_1$, $j = 1, \dots, m_0$

$$F(x_1, \dots, x_{m_0}) = \sum_{i=1}^{m_1} \alpha_i \varphi\left(\sum_{j=1}^{m_0} w_{ij} x_j + b_i\right)$$

is an approximation of the true $f(\cdot)$ / $\|F(\cdot) - f(\cdot)\|_p < \varepsilon$

typically $p=2$

Bounds on Approximation errors

Barron (1993) established the approx. properties of a MLP
Suppose f is the fn. representing the data points
The n/w gives us the approx. function \hat{F} .

$$F \sim \hat{f}$$

Let $\tilde{f}(\underline{\omega})$ denote the multi dimensional Fourier transform
of $f(\underline{x})$; $\underline{x} \in \mathbb{R}^{m_0}$; $\underline{\omega}$ is the frequency
vector

$$f(\underline{x}) = \int_{\mathbb{R}^{m_0}} \tilde{f}(\underline{\omega}) e^{j \underline{\omega}^T \underline{x}} d\underline{\omega}$$

For a complex valued fn $\tilde{f}(\underline{\omega})$ for which $\int_{\mathbb{R}^{m_0}} \tilde{f}(\underline{\omega})$ is integrable, let

$$C_f = \int_{\mathbb{R}^{m_0}} |\tilde{f}(\underline{\omega})| \|\underline{\omega}\|^{\frac{1}{2}} d\underline{\omega}$$

First absolute moment of distribution of f
the Fourier mag.

C_f measures smoothness of f

Consider a ball $B_r = \{ \underline{x} : \|\underline{x}\| \leq r \}$ $r > 0$

Theorem: For every cont. function $f(\underline{x})$ with finite first moment C_f and every $m_1 \geq 1$ \exists a linear combination of sigmoid based functions $F(\underline{x})$ when $f(\underline{x})$ is observed over i.i.d. $\underline{x} \in \{ \underline{x}_i \}_{i=1}^N$ restricted to B_r , the empirical risk

$$E_{\text{ar}} = \frac{1}{N} \sum_{i=1}^N (f(\underline{x}_i) - F(\underline{x}_i))^2 \leq \frac{(2r C_f)^2}{m_1}$$

$$\mathcal{E}_{av}(N) \leq O\left(\frac{C_f^2}{m_1}\right) + O\left(\frac{m_0 m_1}{N} \log N\right)$$

size of hidden neurons
examples

m_1 is $C_f \left(\frac{N}{m_0 \log N}\right)^{\frac{1}{2}}$

$$\mathcal{E}_{av}(N) = O\left(\left(\frac{1}{N}\right)^{\frac{2s}{2s+m_0}}\right)$$

s is a measure of smoothness

$$O\left(\left(\frac{1}{N}\right)^{\frac{1}{m_0}}\right)$$

$$O\left(\frac{1}{N^{1/m_0}}\right) \leftarrow \text{density}$$

Sampling density $\propto N^{\frac{1}{m_0}}$

\Rightarrow We need to densely sample the data
points to learn the function well.

\Rightarrow Higher dim \Rightarrow Exponential growth
in the complexity of
the n/w algo.

A few other technicalities

- 1) Jacobians
- 2) Hessians

Jacobian : Let W denote the total # of free parameters i.e., the synaptic weights and biases of a MLP to form a vector \underline{W} .

Let N denote the examples presented to the n/w.
Using the BPA, we can get $F(\underline{W}, \{\underline{x}\})$

Suppose $\{\underline{x}_n\}_{n=1}^N$

$$J := \left[\frac{\partial F(\underline{w}, \underline{x}_n)}{\partial w_{ij}} \right] \quad \begin{array}{l} i = 1, \dots, W \\ j = 1, \dots, N \end{array}$$

Empirically rank(J) can decide the efficiency of the BPA

If J is not full rank (rank deficiency)
 \implies longer training times

Hessians

$$H := \left[\frac{\partial^2 \xi_{av}(\underline{w})}{\partial \underline{w}^2} \right]$$

Why do we need to study Hessians?

- 1) Eigen values of Hessian have a role in the BPA dynamics.
 - Inverse of H can provide a basis for pruning/deleting insignificant wts.
- 2) Lead to 2nd order opt. methods.

The Hessian of an error surface has :

(Empirical)

- (a) Small # of small ξ large sized eigen values
- (b) large # of med. sized eigen values

whose composition depend

- (1) non zero mean of i/p signals ξ induced neural o/p signals (non zero)
- (2) Correlations between various attributes of a data vector
- (3) Wide variations in the 2nd order derivatives of the J w.r.t W from one layer to the other.

Complexity regularization

In the context of the back prop. algo; we may want to minimize the foll. risk

$$R(\underline{w}) = \underbrace{\xi_{av}(\underline{w})}_{\substack{\downarrow \text{standard} \\ \text{performance} \\ \text{metric}}} + \lambda \underbrace{\xi_c(\underline{w})}_{\substack{\uparrow \text{complexity} \\ \text{part}}}$$

λ ↑ Scalar

$\xi_c(\underline{w})$ is the complexity penalty measured in terms of \underline{w}

$\xi_c(\underline{w}) = \|\underline{w}\|_2^2$ (Choice of $\xi_c(\underline{w})$, for some wts. to zero ξ permit larger weights)

Alternative Strategies to learning

$$\xi_{av}(\underline{w} + \Delta \underline{w}) = \xi_{av}(\underline{w}) + -g^T(\underline{w}) \Delta \underline{w} + \frac{1}{2!} \Delta \underline{w}^T H \Delta \underline{w} + \text{h.o.t.}$$

↙ 2nd order

a small perturbation $\Delta \underline{w}$ around \underline{w} .

We can set

$$\frac{\partial \xi_{av}(\underline{w} + \Delta \underline{w})}{\partial \Delta \underline{w}} = 0 \Rightarrow$$

$$\Delta \underline{w} = -H^{-1}g$$

Optimum brain Surgeon Algo

$$\min_{\Delta \underline{w}} \frac{1}{2} \Delta \underline{w}^T H \Delta \underline{w}$$

- 1) We set the i th comp. of $\Delta \underline{w}$ to zero
 ξ min. over all synaptic wt. vectors that remain
- 2) Do this min. again over all indices 'i'

Cross validation strategies

Need:

We may get a low training error over a data set, but the error may be pronounced over a different data set. To get a reasonable performance, we divide the training sets into groups, train & test over the groups. "Empirical soln" in the absence of a "theoretical" set up

- Eg:
- 1) Leave 1 out strategy: Train on $(N-1)$ samples & test on the other. Do over all $\binom{N}{1}$ choices
 - 2) K-fold cross validation: Divide N samples into k equal groups, Train on $(k-1)$ groups & test on the other. Repeat over all the choices

Convolution networks

Motivation :

Structural layout of the MLP algo. with preprocessing steps that are neurobiologically inspired.

Ex:

Work from Hubel & Wiesel related to specialized neurons in the visual cortex of a cat.

Simple cells	—	locally sensitive
Complex cells	—	Orientation <u>etc</u>

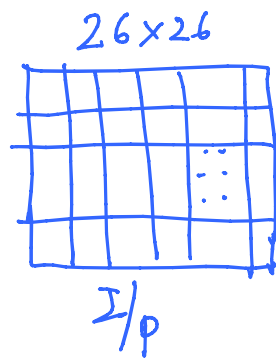
GOAL : To design a MLP to recognize 2D/3D shapes/objects with high invariance to translation/rotation/scaling.

1) Feature extraction : Each neuron takes its inputs from local receptive fields in previous layers, forcing to extract local features.
Once the features are extracted, its exact location is less important as long the relative position w.r.t. other features is preserved.

2) Feature mapping: Each computational layer is composed of multiple feature maps, with each feature map being in the form of an appropriate geometry to the signal (e.g., plane for images etc.) ξ constrained to share the same synaptic weights

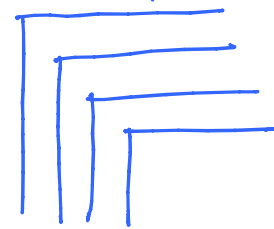
- a) Shift Invariance: Forced into the feature map through convolution with a kernel of a small size
- b) Reduction in the # of free parameters. This is ensured via wt. sharing.

3) Sub-sampling: Each conv. layer performs some local operations followed by a non-linear activation $\varphi(\cdot)$ & then sub-sampling that reduces the resolution to sensitivity of the feature maps to Jitter & distortion.



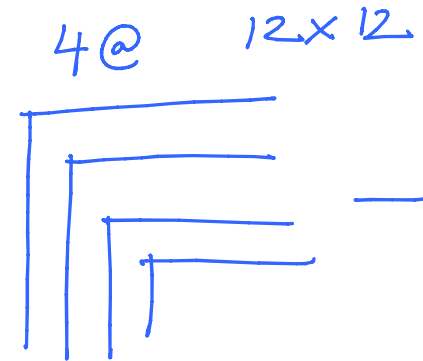
4 masks of size 3x3

4 @ 24x24 feature maps



$\varphi(\cdot)$ (applied to)

Subsampling over a 2x2 field

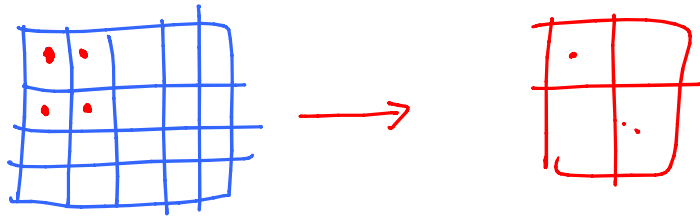


Conv
+
Subsam
⋮

Pyramidal Effect!

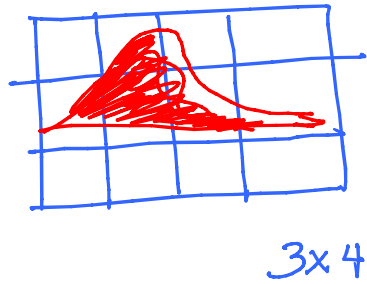
Conv. followed by subsampling is inspired by simple cells followed by complex cells as described by Hubel & Wiesel.

As # of Sub-sampling layers increase \uparrow Spatial resolution \downarrow compared to what you had in the prev. layers



Let us understand the computations through an example

1) Image:



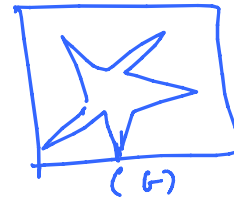
→
I/p to
an array of integers

0	1	1	0
1	1	1	1
0	0	0	0

2) Filter the image by
features (low pass)

masks / feature masks to get the
to high pass)

More
Low pass



High
Pass

Types of masks :

1	1
1	1

2x2

0	1
1	0

1	0
0	1

3x4

(X)

2x2 (Mask)
filter

2x3 array
@ o/p

1	2	3	4
5	6	7	8
9	10	11	12

(X)

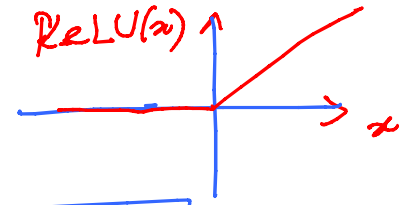
1	1
1	1

Low Pass

14	18	.
-	-	.

Eg:

3) Once the features are extracted, we feed this to a non-linear act. fn $\text{ReLU}(x) = \max(0, x)$



4) Pooling (Subsampling)

1st 2x2 subblock

1	3	10	1
4	9	6	8
1	0	1	7
8	9	3	4

2nd subblock of size 2x2

Max Pooling

over a 2x2 subblock

(Non-overlapping)

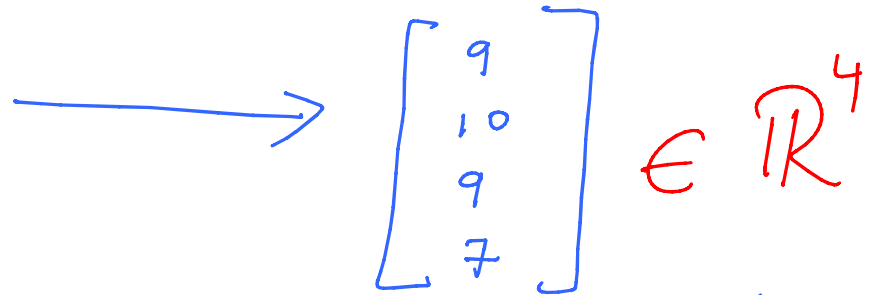
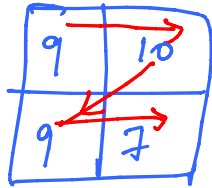
9	10
9	7

Subsampled image/map

Rect. Feat. Map

5) Iterate steps (2) - (4) in a pyramidal way to get the final size as desired.

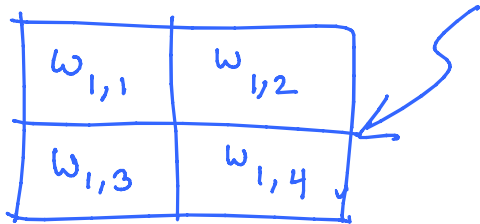
Rastering



a feature vector

desired response → ↓ fed into the MLP

Masks can be adaptive



Conditions on the elements within the mask

⇒ Regularized n/w for learning I/O mappings

Cover's Theorem

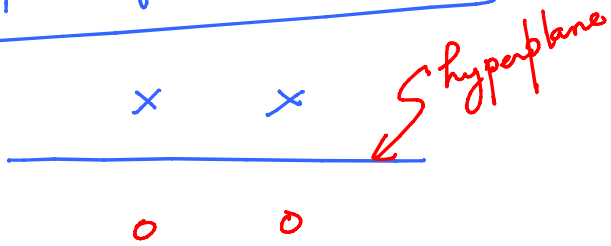
Ref 1: Thomas Cover, Feb 1964

Ref 2: Nils Nilsson, Learning Machines, 1965

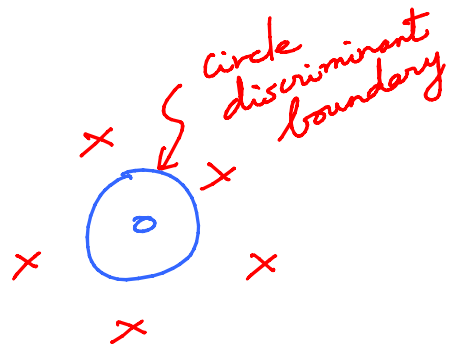
Motivation:

- a) How do we quantify the complexity of a neural network architecture?
- b) Need a counting measure for a discrete set of mappings.

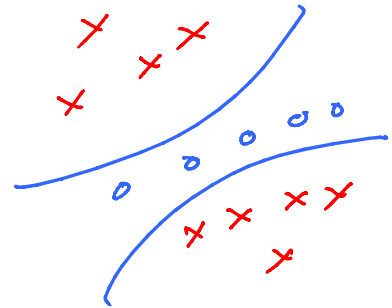
Examples of dichotomies



(a) Linear dichotomy



(b) Spherical dichotomy



(c) Quadratic dichotomy

Consider a fixed finite set of input vectors

$$\{ \underline{x}_1, \dots, \underline{x}_p \} \text{ where } \underline{x}_i \in \mathbb{R}^N$$

Can we attempt to compute the dichotomies for a perceptron?

The # of linearly realizable dichotomies on the set of points depends on a mild condition called 'general position'

General position demands that no subset of size $< N$ on $\{ \underline{x}_1, \underline{x}_2, \dots, \underline{x}_p \}$ is linearly dependent

Theorem: Let $\{\underline{x}_1, \dots, \underline{x}_p\}$ be vectors in

\mathbb{R}^N that are in a general position. The # of distinct dichotomies applied to these points can be realized by a hyperplane

is

$$C(P, N) = 2 \sum_{k=0}^N \binom{P-1}{k}$$


Proof: We start with P points in general position. Let us assume that there are $C(P, N)$ dichotomies on them. Suppose we add an extra point to this set. We need $C(P+1, N)$ dichotomies.

Idea: Set up a recursion to link $C(P+1, N)$ with $C(P, N)$. Let (b_1, b_2, \dots, b_P) be a dichotomy realizable by a hyperplane over the set of P inputs. $b_i \in \{-1, 1\}$ for every $i = 1, \dots, P$. There is a set of weights w so that

$$\left(\text{sign}(\underline{w}^T \underline{x}_1), \text{sign}(\underline{w}^T \underline{x}_2), \dots, \text{sign}(\underline{w}^T \underline{x}_p) \right) = \underline{(b_1, b_2, \dots, b_p)}$$

Using one such \underline{w} , we get a dichotomy over $P+1$ points

$$\text{i.e., } \left(\underset{\uparrow b_1}{\text{sign}(\underline{w}^T \underline{x}_1)}, \dots, \underset{\uparrow b_p}{\text{sign}(\underline{w}^T \underline{x}_p)}, \text{sign}(\underline{w}^T \underline{x}_{p+1}) \right)$$

For every linearly realized dichotomy over P points, there is at least one linearly realized dichotomy over $P+1$ points.

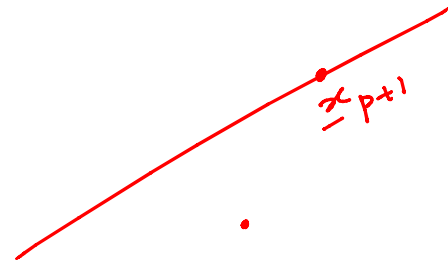
Different dichotomies over P points define different dichotomies over $P+1$ points since they differ somewhere over first P coordinates

Note that the additional dichotomy $(b_1, \dots, b_p, \underline{\text{sign}}(\underline{w}^T \underline{x}_{p+1}))$ is also possible by reversing the sign of the last coordinate.

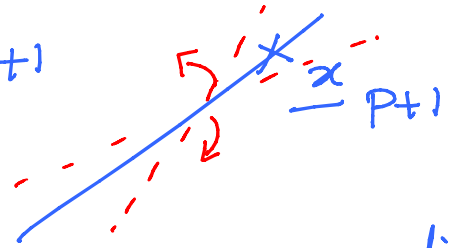
$$\Rightarrow C(p+1, N) > C(p, N)$$

Let $C(p+1, N) = C(p, N) + E$ extra dichotomies possible.

There are 2 cases to consider



Case A: One of the weight vectors \underline{w} that generates
 (b_1, b_2, \dots, b_p) passes through \underline{x}_{p+1}



By adjusting the angle of the hyperplane, we can adjust
sign $(\underline{w}^T \underline{x}_{p+1})$ to $+1$ or -1
i.e., both $(b_1 \dots b_p +1)$ and $(b_1 \dots b_p -1)$ are
also possible!

Case B: No hyperplane passes through \underline{x}_{p+1} and
generates b_1, \dots, b_p on first p vectors. \underline{x}_{p+1}

⇒ Points lies on one side of the old dichotomy

E is the # of dichotomies over p points that are
realized by a hyperplane passing through a fixed point
 \underline{x}_{p+1} . By forcing the hyperplane to pass through
 \underline{x}_{p+1} , we are going to $N-1$ dimensions instead of N

Geometrically, if a point is on the x -axis, the hyperplane has $N-1$ axes left to work on the problem. If it is not on the x -axis, then rotate the axes of the space to get the point on the x -axis and there is no effect on the geometry of the problem

$$\begin{aligned} \therefore E &= C(P, N-1) \leftarrow \\ \therefore C(P+1, N) &= C(P, N) + \underbrace{C(P, N-1)}_{\leftarrow E} \quad \text{--- } \textcircled{1} \end{aligned}$$

Let us consider the boundary conditions

$$C(1, N) = 2 \quad \checkmark \quad \left(\text{There is 1 point in } \mathbb{R}^N \text{ \& can be realized by 2 labels} \right)$$

$$C(P, 1) = 2P \quad \checkmark \quad \left(\text{There are } P \text{ points in } \mathbb{R}^1 \right)$$

Consider $P = 3$, we have

$0 - 0$
 $x - 1$

0	0	0	
0	0		x
0		x	0
0		x	x

x	0	0
x	0	x
x	x	0
x	x	x

Observe that except $0x0$ and $x0x$, we can have a hyperplane that can shatter the points!

Let us prove the result through induction.
 (We are doing induction over 'P')

Base case :
 Case 1: $C(1, N) = 2$ as expected

There is 1 point in N dim. Follows from
 one of the boundary conditions

Induction :

$$C(P+1, N) = 2 \sum_{k=0}^N \binom{P-1}{k} + 2 \sum_{k=0}^{N-1} \binom{P-1}{k}$$

$\underbrace{\hspace{15em}}_{C(P, N)} \qquad \underbrace{\hspace{15em}}_{C(P, N-1)}$

Ponder on $\binom{P-1}{0}$ case
 Meaning of '0' dim.

Given by statement of
 Cover's theorem

Now, simplifying,

$$= 2 \sum_{k=0}^N \binom{p-1}{k} + 2 \sum_{k=0}^N \binom{p-1}{k-1} \quad \left(\because \binom{p-1}{-1} = 0 \right)$$

$$= 2 \left[\sum_{k=0}^N \binom{p-1}{k} + \binom{p-1}{k-1} \right]$$
$$= 2 \sum_{k=0}^N \binom{p}{k} \quad \left(\because \binom{n}{k} = \binom{n-1}{k} + \binom{n-1}{k-1} \right)$$

Basic identity from elementary combinatorics \square

Implications :

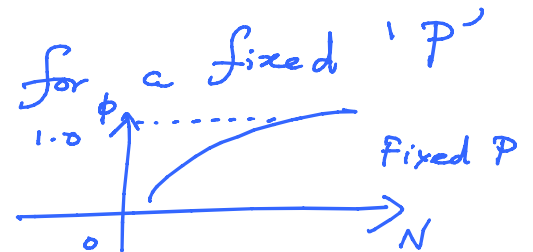
(linear decision boundary)

Let us consider the prob. of having a perceptron that can provide a linear dichotomy over P points

$$\phi = \frac{C(P, N)}{2^P} \leftarrow \text{total \# of dichotomies (not necessarily linear)}$$

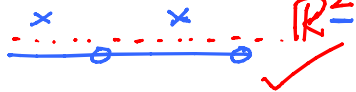
$$= 2^{1-P} \sum_{k=0}^N \binom{P-1}{k}$$

Home Work: Plot ϕ vs. dim N
Observe the concavity of ϕ



\mathbb{R}^1

(Linear decision boundary is not possible)



\mathbb{R}^2

Yes, a linear decision boundary is possible

Let us revisit the XOR problem

Recall: Cover's theorem has 2 important consequences

1) The use of a non-linear function i.e., a hidden function defined by $\varphi(\underline{x})$ acts on the input vector

2) High dimensionality of the hidden / feature space compared to the input space

$N := \text{dim. feature space}$
 $M := \text{dim. input space}$

$$N > M$$

Recall:

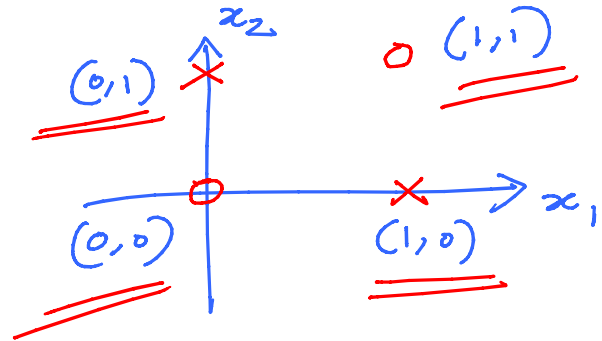
A dichotomy is φ -separable if \exists a
N-dim. vector \underline{w}

$$\underline{w}^T \phi(\underline{x}) > 0$$

$$\underline{w}^T \phi(\underline{x}) \leq 0$$

$\underline{x} \in$ Class 1
 $\underline{x} \in$ Class 2

We have the XOR problem



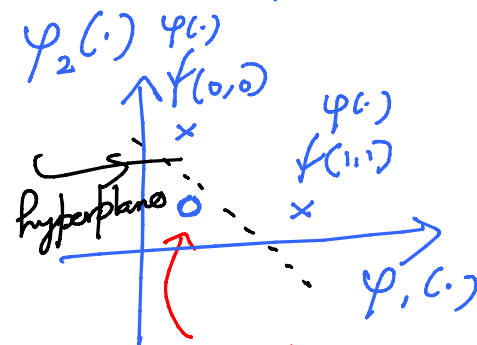
Let us consider the pair of Gaussian hidden functions.

$$\varphi_1(\underline{x}) = \exp\left(-\|\underline{x} - \underline{t}_1\|^2\right); \underline{t}_1 = [1, 1]^T$$

$$\varphi_2(\underline{x}) = \exp\left(-\|\underline{x} - \underline{t}_2\|^2\right); \underline{t}_2 = [0, 0]^T$$

Let us tabulate the comp. evaluations of \underline{x} over φ_1 & φ_2

\underline{x}	$\varphi_1(\underline{x})$	$\varphi_2(\underline{x})$
$(1, 1)$	1	0.1353
$(0, 1)$	0.3678	0.3678
$(0, 0)$	0.1353	1
$(1, 0)$	0.3678	0.3678



Both points $(0, 0)$ & $(1, 1)$ map to the same point.