Nonlinear representations

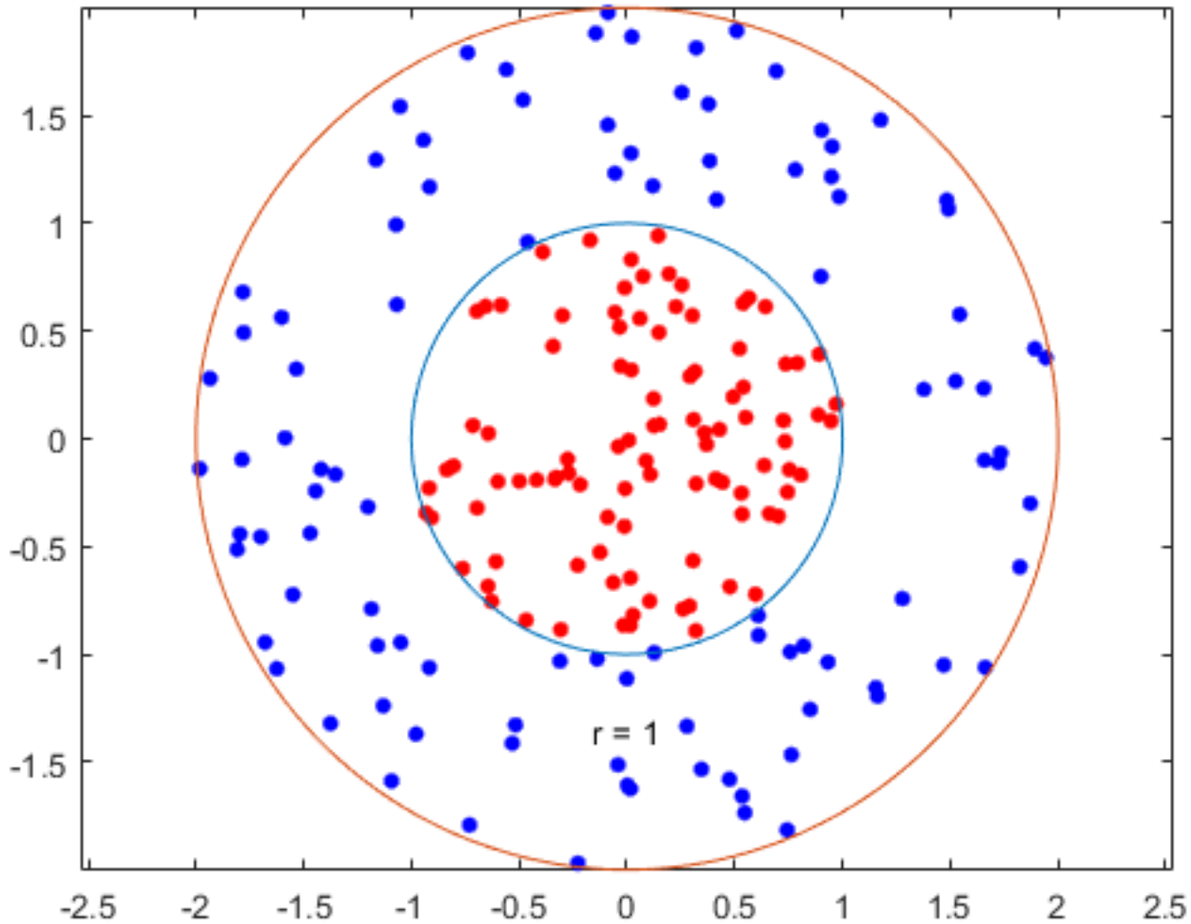

Comments

- Assignment 2 deadline extended to the end of Friday
- I do apologize for the late additions to the assignment. Any issues?

Representations for learning nonlinear functions

- Generalized linear models enabled many p(y | x) distributions
  - Still however learning a simple function for E[Y | x], i.e., f(<x,w>)

- Approach we discussed earlier: augment current features x using polynomials

- There are many strategies to augmenting x
  - fixed representations, like polynomials, wavelets
  - learned representations, like neural networks and matrix factorization

What if classes are not linearly separable?

$x_1^2 + x_2^2 = 1 \quad f(x) = x_1^2 + x_2^2 - 1$

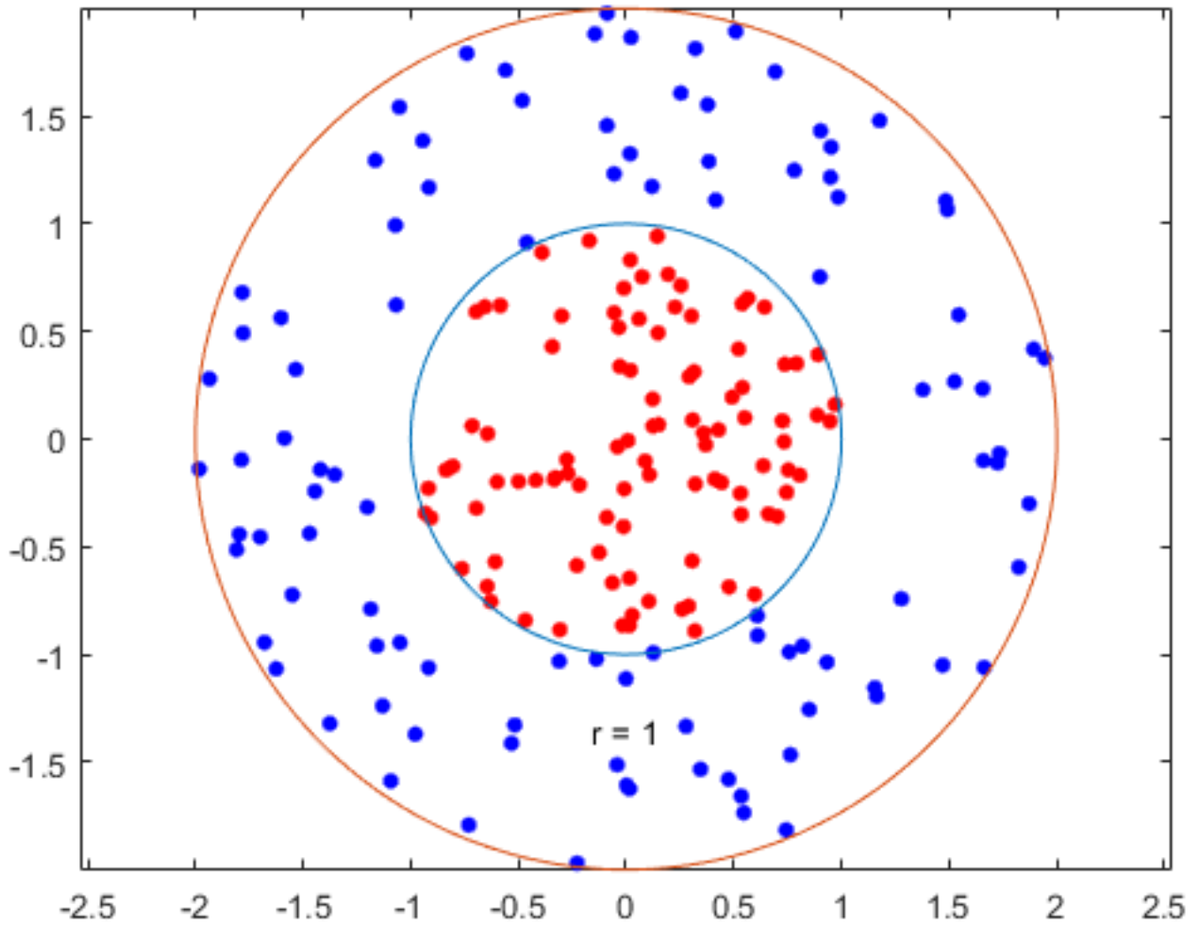

$x_1 = x_2 = 0$

$\Longrightarrow \quad f(x) = -1 < 0$

$x_1 = 2, x_2 = -1$

$\Longrightarrow \quad f(x) = 4 + 1 - 1 = 4 > 0$

How to learn f(x) such that f(x) > 0 predicts + and f(x) < 0 predicts negative?

What if classes are not linearly separable?

$x_1^2 + x_2^2 = 1 \quad f(x) = x_1^2 + x_2^2 - 1$

$$\boldsymbol{\phi}(\mathbf{x}) = \begin{bmatrix} x_1^2 \\ x_2^2 \\ 1 \end{bmatrix} \qquad f(\mathbf{x}) = \boldsymbol{\phi}(\mathbf{x})^\top \mathbf{w}$$

If use logistic regression, what is p(y=1 | x)?

How to learn f(x) such that f(x) > 0 predicts + and f(x) < 0 predicts negative?

What if classes are not linearly separable?

$x_1^2 + x_2^2 = 1 \quad f(x) = x_1^2 + x_2^2 - 1$

$$\boldsymbol{\phi}(\mathbf{x}) = \begin{bmatrix} x_1^2 \\ x_2^2 \\ 1 \end{bmatrix} \qquad f(\mathbf{x}) = \boldsymbol{\phi}(\mathbf{x})^\top \mathbf{w}$$

Imagine learned w. How do we predict on a new x?

Nonlinear transformation

$$\mathbf{x} \rightarrow \boldsymbol{\phi}(\mathbf{x}) = \begin{pmatrix} \phi_1(\mathbf{x}) \\ \dots \\ \phi_p(\mathbf{x}) \end{pmatrix}$$

$$\text{e.g., } \mathbf{x} = [x_1, x_2], \quad \boldsymbol{\phi}(\mathbf{x}) = \begin{pmatrix} x_1 \\ x_2 \\ x_1^2 \\ x_1 x_2 \\ x_2^2 \\ x_1^3 \\ x_2^3 \end{pmatrix}$$

Gaussian kernel / Gaussian radial basis function

$$k(\mathbf{x}, \mathbf{x}') = \exp\left(\frac{-\|\mathbf{x} - \mathbf{x}'\|_2^2}{\sigma^2}\right)$$

$$f(\mathbf{x}) = \sum_{i=1}^{p} w_i k(\mathbf{x}, \mathbf{x}_i)$$

$$\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \rightarrow \begin{bmatrix} k(\mathbf{x}, \mathbf{x}_1) \\ \vdots \\ k(\mathbf{x}, \mathbf{x}_p) \end{bmatrix}$$

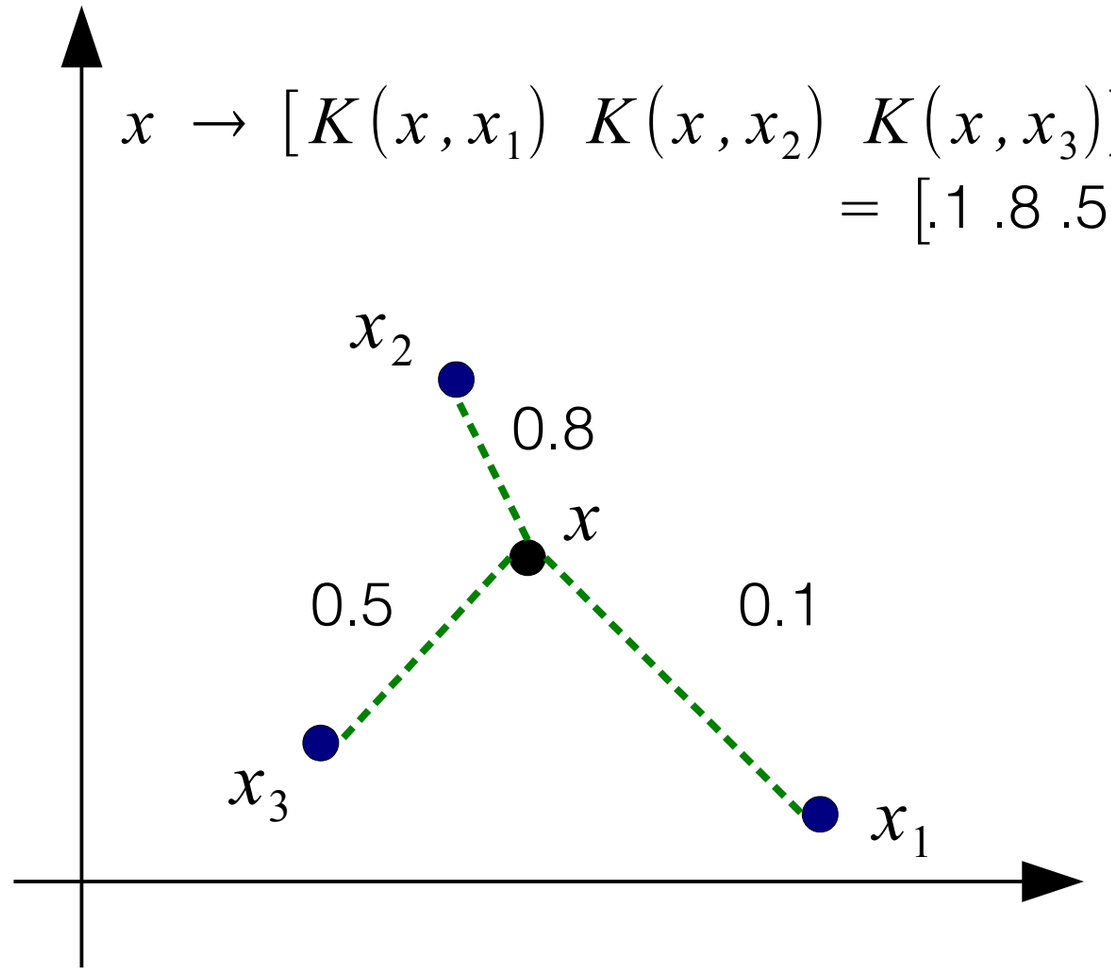

Gaussian kernel / Gaussian radial basis function

$$k(\mathbf{x}, \mathbf{x}') = \exp\left(\frac{-\|\mathbf{x} - \mathbf{x}'\|_2^2}{\sigma^2}\right)$$

Kernel

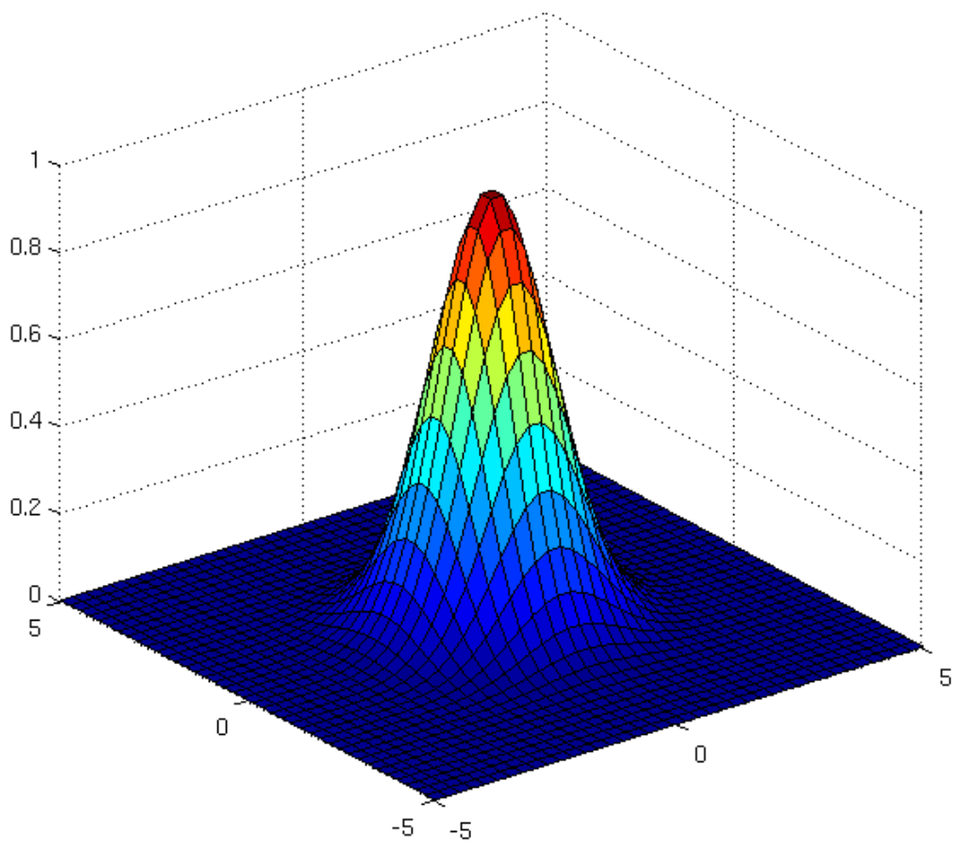

$$f(\mathbf{x}) = \sum_{i=1}^{p} w_i k(\mathbf{x}, \mathbf{x}_i)$$

Possible function f with several centers

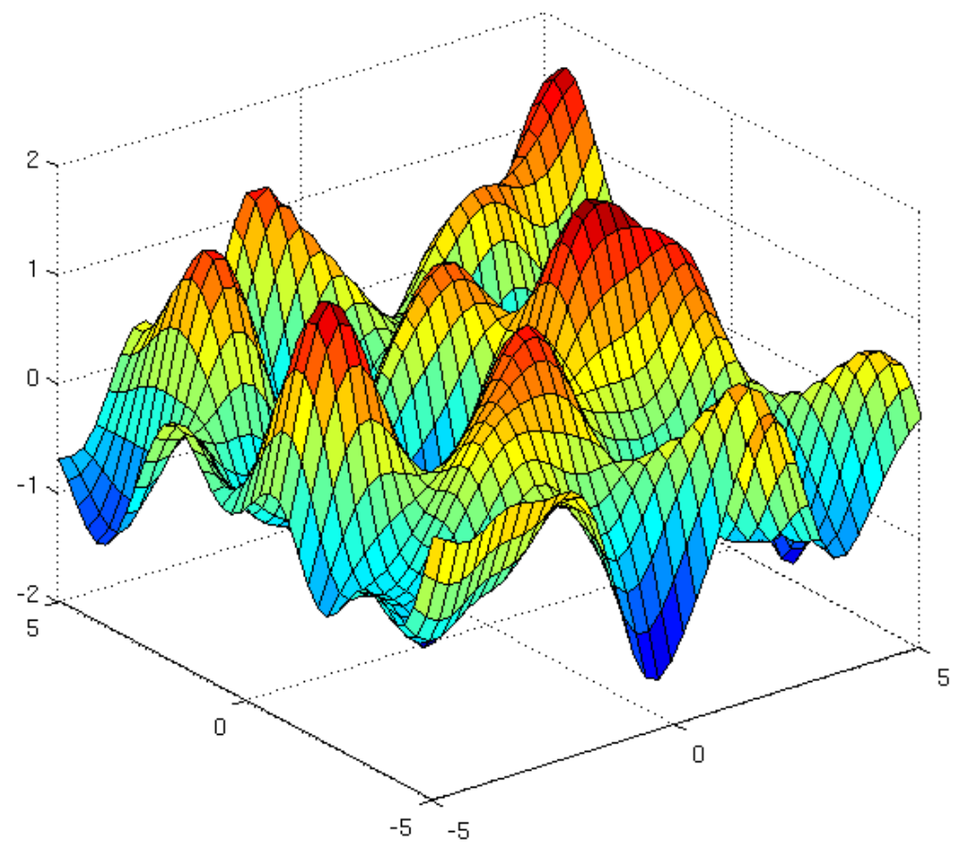

Selecting centers

- Many different strategies to decide on centers
  - many ML algorithms use kernels e.g., SVMs, Gaussian process regression
- For kernel representations, typical strategy is to select training data as centers
- Clustering techniques to find centers
- A grid of values to best (exhaustively) cover the space
- Many other strategies, e.g., using information gain, coherence criterion, informative vector machine

Covering space uniformly with centres

- Imagine has 1-d space, from range [-10, 10]
- How would we pick p centers uniformly?
- What if we have a 5-d space, in ranges [-1,1]?
  - To cover entire 5-dimensional, need to consider all possible options
  - Split up 1-d into m values, then total number of centres is m^5
  - i.e., for first value of x1, can try all other m values for x2, …, x5

Why select training data as centers?

- Observed data indicates the part of the space that we actually need to model
  - can be much more compact than exhaustively covering the space
  - imagine only see narrow trajectories in world, even if data is d-dimensional, data you encounter may lie on a much lower-dimensional manifold (space)

- Any issues with using all training data for centres?
  - How can we subselect centres from training data?

How would we use clustering to select centres?

- Clustering is taking data and finding p groups
- What distance measure should we use?
  - If k(x,c) between 0 and 1 and k(x,x) = 1 for all x, then 1-k(x,c) gives a distance

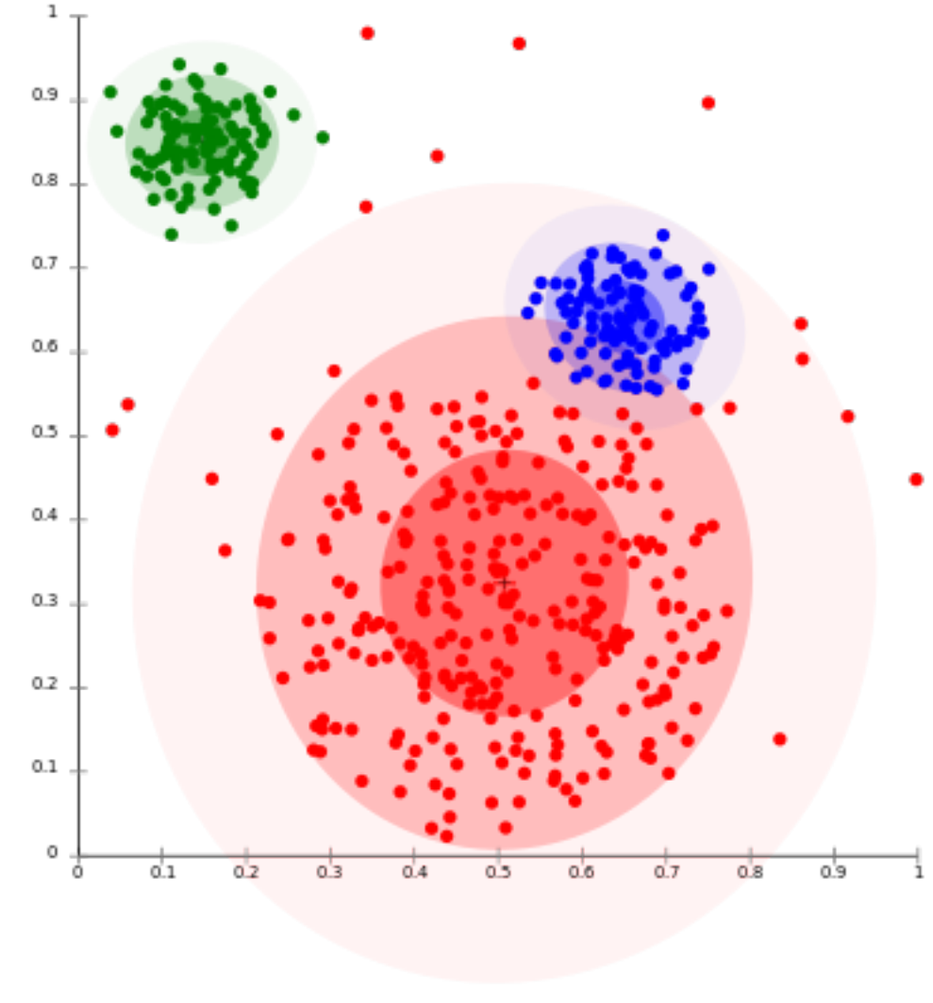

What if we select centres randomly?

- Are there any issues with the linear regression with a kernel transformation, if we select centres randomly from the data?

- If so, any suggestions to remedy the problem?

$$\sum_{i=1}^{n} \left(\sum_{j=1}^{p} k(\mathbf{x}, \mathbf{z}_j) \mathbf{w}_j - y_i \right)^2$$

Exercise

- What would it mean to use an l1 regularizer with a kernel representation?
  - Recall that l1 prefers to zero elements in w

$$\sum_{i=1}^{n} \left(\sum_{j=1}^{p} k(\mathbf{x}, \mathbf{z}_j) \mathbf{w}_j - y_i \right)^2 + \lambda \|\mathbf{w}\|_1$$

Exercise: How do we decide on the nonlinear transformation?

- We can pick a 5-th order polynomial or 6-th order, or… Which should we pick?
- We can pick p centres. How many should we pick?
- How can we avoid overparametrizing or underparameterizing?

Other similarity transforms

- Linear kernel: $k(\mathbf{x}, \mathbf{c}) = \mathbf{x}^\top \mathbf{c}$

- Laplace kernel (Laplace distribution instead of Gaussian)

$$k(\mathbf{x}, \mathbf{c}) = \exp(-b\|\mathbf{x} - \mathbf{c}\|_1)$$

- Binning transformation

$$s(\mathbf{x}, \mathbf{c}) = \begin{cases} 1 & \text{if } \mathbf{x} \text{ in box around } \mathbf{c} \\ 0 & \text{else} \end{cases}$$

Dealing with non-numeric data

- What if we have categorical features?
  - e.g., feature is the blood type of a person
- Even worse, what if we have strings describing the object?
  - e.g., feature is occupation, like “retail”

Some options

- Convert categorical feature into integers
  - e.g., {A, B, AB, O} —> {1, 2, 3, 4}
  - Any issues?
- Convert categorical feature into indicator vector
  - e.g., A —> [1 0 0 0], B —> [0 1 0 0], …
  - Any issues?

Using kernels for categorical or non-numeric data

- An alternative is to use kernels (or similarity transforms)
- If you know something about your data/domain, might have a good similarity measure for non-numeric data
- Some more generic kernel options as well
  - Matching kernel: similarity between two items is the proportion of features that are equal

Example: matching kernel

$\mathbf{x} =$

age	{15-24, 25-34, …, 65+}
gender	{F, M}
income	{Low, Medium, High}
education	{Bachelors, Trade-Sch, High-Sch, …}

Census dataset: Predict hours worked per week

Example: Matching similarity for categorical data

$\mathbf{x} =$

age	24-34	35-44
gender	F	F
income	Medium	Medium
education	Trade-Sch	Bachelors

$k(\mathbf{x}_1, \mathbf{x}_2) = k$ = 0.5

Representational properties of transformations

- Approximation properties: which transformations can approximate “any function”?
- Radial basis functions (a huge number of them)
- Polynomials and the Taylor series
- Fourier basis and wavelets

Distinction with the kernel trick

- When is the similarity actually called a “kernel”?

- Nice property of kernels: can always be written as a dot product in some feature space

$$k(\mathbf{x}, \mathbf{c}) = \boldsymbol{\psi}(\mathbf{x})^\top \boldsymbol{\psi}(\mathbf{c})$$

- In some cases, they are used to compute inner products efficiently, assuming one is actually learning with the feature expansion
  - This is called the kernel trick

- Implicitly learning with feature expansion $\boldsymbol{\psi}(\mathbf{x})$
  - not learning with expansion that is similarities to centres

Example: polynomial kernel

$$\boldsymbol{\phi}(\mathbf{x}) = \begin{bmatrix} \mathbf{x}_1^2 \\ \sqrt{2}\mathbf{x}_1\mathbf{x}_2 \\ \mathbf{x}_2^2 \end{bmatrix}$$

$$k(\mathbf{x}, \mathbf{x}') = \langle \boldsymbol{\phi}(\mathbf{x}), \boldsymbol{\phi}(\mathbf{x}') \rangle = \langle \mathbf{x}, \mathbf{x}' \rangle^2$$

In general, for order d polynomials, $k(\mathbf{x}, \mathbf{x}') = \langle \mathbf{x}, \mathbf{x}' \rangle^d$

Gaussian kernel

Infinite polynomial representation

$$\boldsymbol{\phi}(x) = \exp(-\gamma x^2) \begin{bmatrix} 1 \\ \sqrt{\frac{2\gamma}{1!}} x \\ \sqrt{\frac{(2\gamma)^2}{2!}} x^2 \\ \vdots \end{bmatrix}$$

$$k(\mathbf{x}, \mathbf{x}') = \exp\left(\frac{-\|\mathbf{x} - \mathbf{x}'\|_2^2}{\sigma^2}\right)$$

Regression with new features

$$\min_{\mathbf{w}} \sum_{i=1}^{n} (\boldsymbol{\phi}(\mathbf{x}_i)^\top \mathbf{w} - y_i)^2 = \min_{\mathbf{w}} \sum_{i=1}^{n} \left(\left(\sum_{j=1}^{p} \boldsymbol{\phi}_j(\mathbf{x}_i) \mathbf{w}_j \right) - y_i \right)^2$$

What if p is really big?

$$\min_{\mathbf{w}} \sum_{i=1}^{n} (\boldsymbol{\phi}(\mathbf{x}_i)^\top \mathbf{w} - y_i)^2 = \min_{\mathbf{a}} \sum_{i=1}^{n} \left(\left(\sum_{j=1}^{n} \langle \boldsymbol{\phi}(\mathbf{x}_i), \boldsymbol{\phi}(\mathbf{x}_j) \rangle \mathbf{a}_j \right) - y_i \right)^2$$

If can compute dot product efficiently, then can solve this regression problem efficiently

What about learning the representation?

- We have talked about fixed nonlinear transformations
  - polynomials
  - kernels

- How do we introduce learning?
  - could learn centers, for example
  - learning is quite constrained, since can only pick centres and widths

- Neural networks **learn** this transformation more from scratch

Fixed representation vs. NN

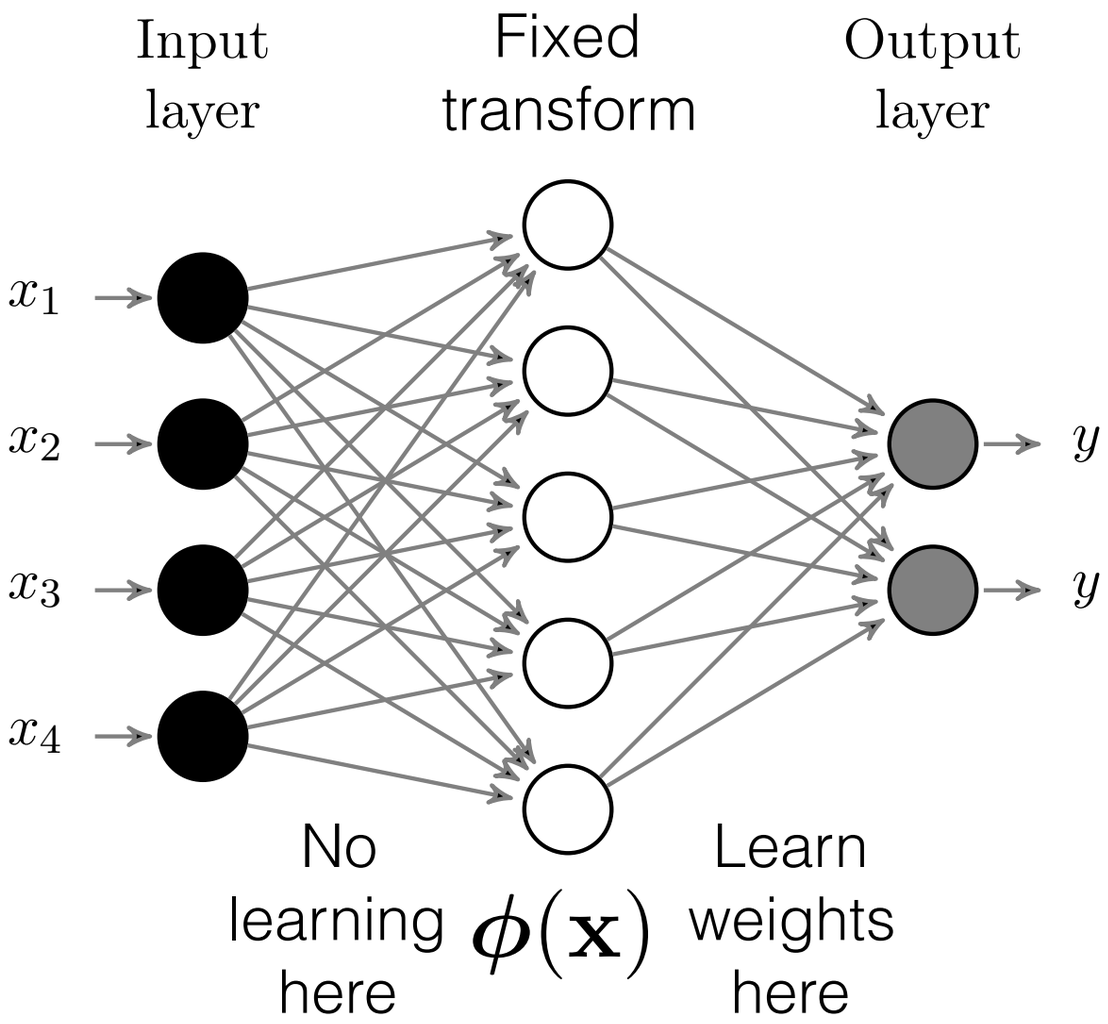

GLM with augmented fix representation

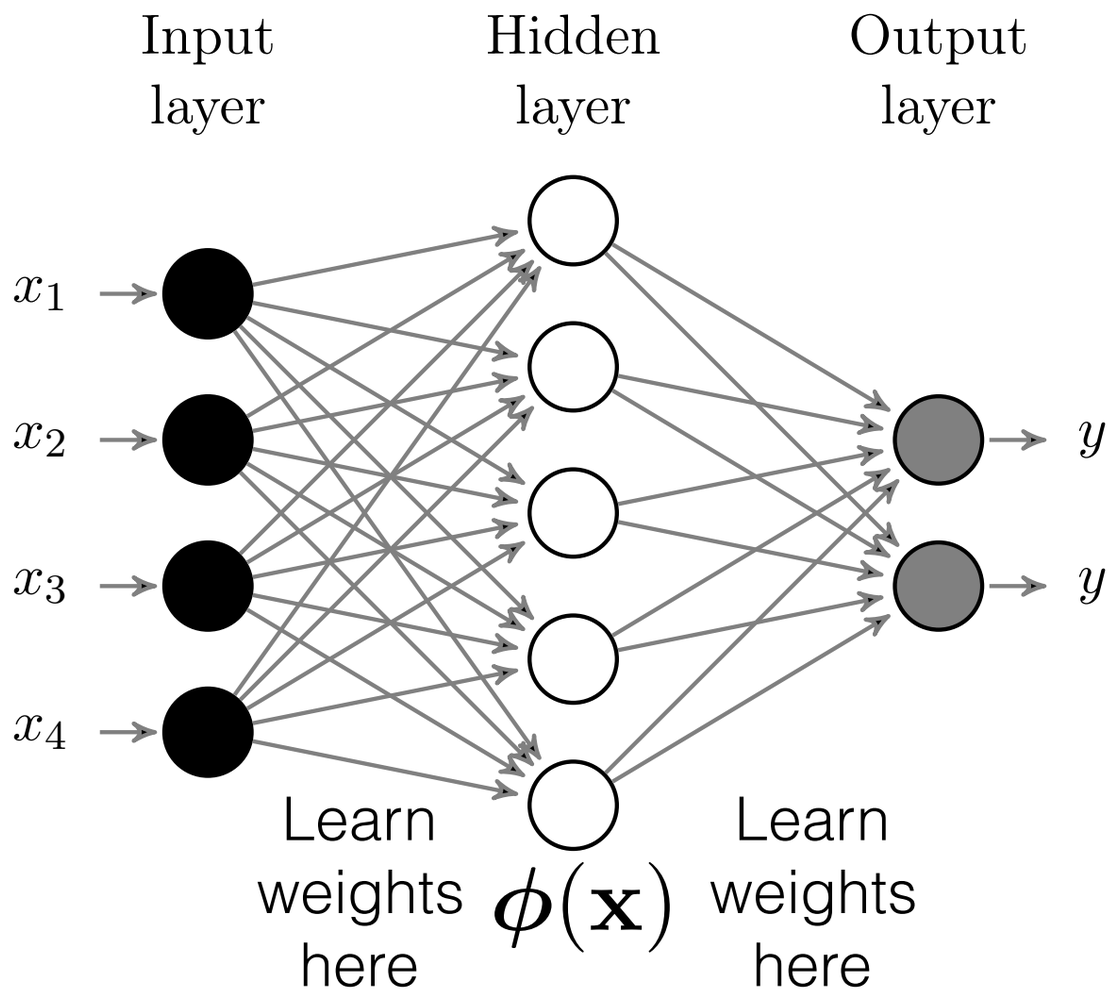

Two-layer neural network

Explicit steps in visualizations

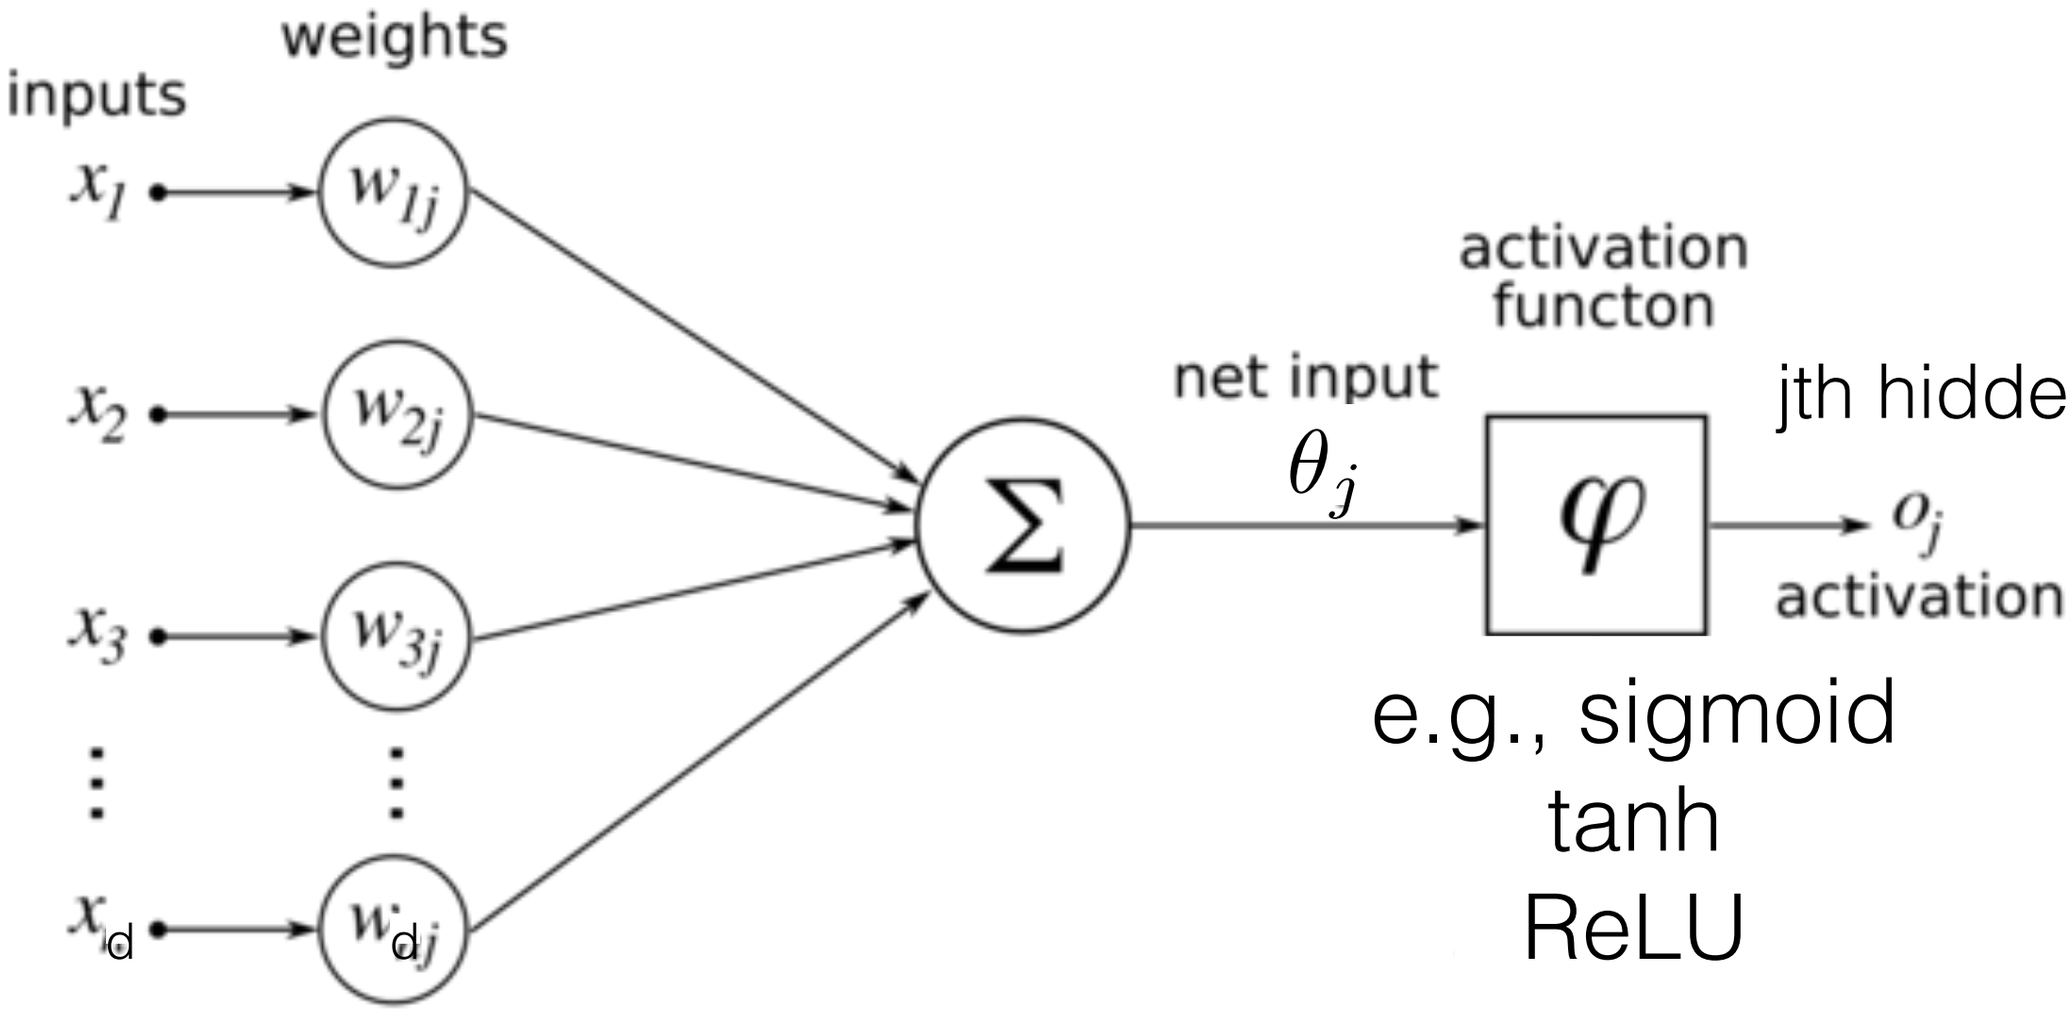

e.g., sigmoid
tanh
ReLU

Example: logistic regression versus neural network

- Both try to predict p(y = 1 | x)
- Logistic regression learns W such that

$$f(\mathbf{x}\mathbf{W}) = \sigma(\mathbf{x}\mathbf{W}) = p(y = 1|\mathbf{x})$$

- Neural network learns W1 and W2 such that

$$p(y = 1|\mathbf{x}) = \sigma(\mathbf{h}\mathbf{W}^{(1)}) = \sigma(\sigma(\mathbf{x}\mathbf{W}^{(2)})\mathbf{W}^{(1)}).$$

No representation learning vs. neural network

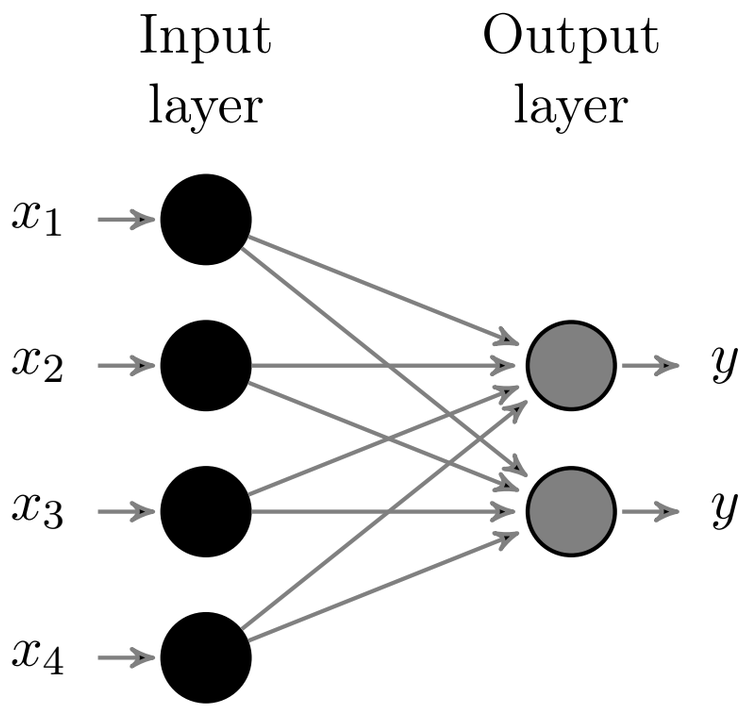

GLM
(e.g. logistic regression)

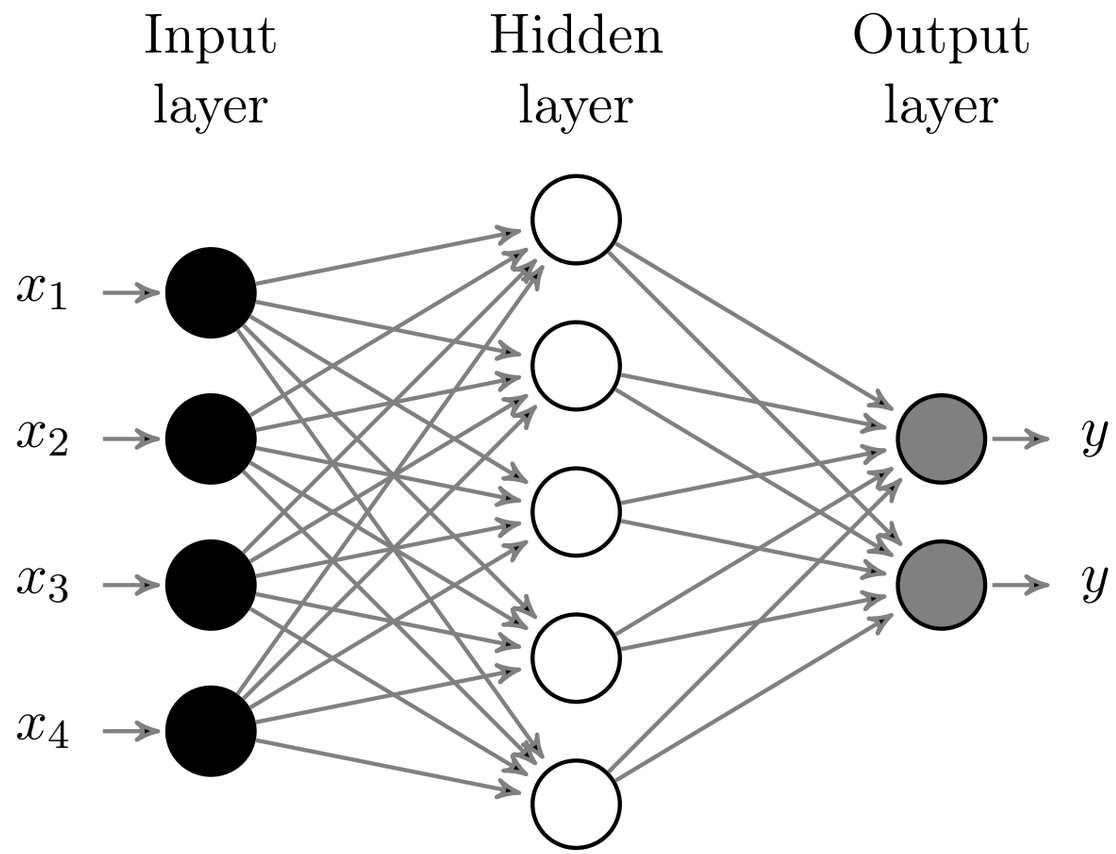

Two-layer neural network

What are the representational capabilities of neural nets?

- Single hidden-layer neural networks with sigmoid transfer can represent any continuous function on a bounded space within epsilon accuracy, for a large enough number of hidden nodes
  - see Cybenko, 1989: “Approximation by Superpositions of a Sigmoidal Function”

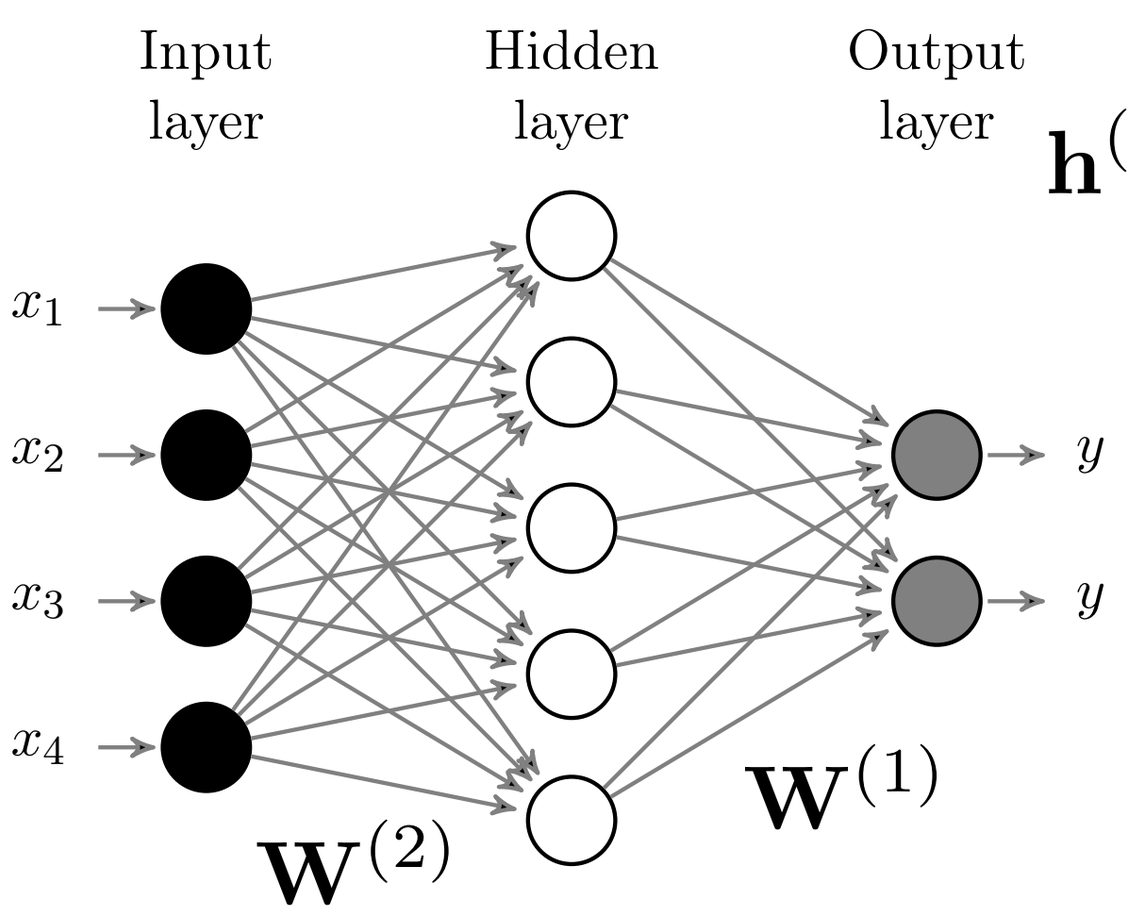

$$\mathbf{h}^{(2)} = f_2(\mathbf{x}\mathbf{W}^{(2)})$$

$$\hat{\mathbf{y}} = f_1(\mathbf{h}^{(1)}\mathbf{W}^{(1)})$$