

Foundations of Machine Learning (Part II)

Mohammad Emtiyaz Khan
AIP (RIKEN), Tokyo
OIST, Okinawa

<http://emtiyaz.github.io>

emtiyaz.khan@riken.jp

June 22, 2022



©Mohammad Emtiyaz Khan 2022

Goals

Learn a few methods for classification and understand why they might work and sometimes fail. ¹.

1. Classification.

2. ~~K-Nearest Neighbors (k-NN)~~.

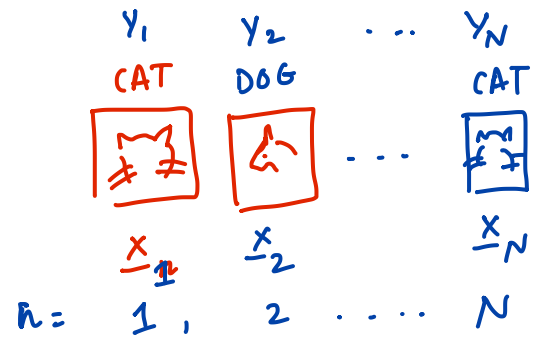
3. Logistic Regression.

4. Deep Neural Networks.

¹Some figures are taken from Hastie, Tibshirani, and Friedman's book on statistical learning. Some from Chris Bishop's Machine learning book and one from Kevin Murphy's book. The CNN picture taken from <http://www.mdpi.com/1099-4300/19/6/242>

1 Classification

Similar to regression, classification relates input variables \mathbf{x} to the output variable y , but now y can take only discrete values, i.e. y is a categorical (or nominal) variable.



$$y_n \in \{C_1, C_2\}$$

{"dog", "cat"}

Binary classification

When y can only take two discrete values, it is called binary classification. We will denote these values as $y \in \{C_1, C_2\}$. These values are also called class labels or simply classes. Other common notations are $y \in \{-1, +1\}$ or $y \in \{0, 1\}$, although there may not necessarily be any ordering between the two classes.

$$y \in \{0, 1\}$$

No ordering, just a set.



Multi-class classification

In a multi-class classification, y can take multiple discrete values i.e. $y \in \{C_0, C_1, \dots, C_{K-1}\}$ for a K -class problem. Again, there is no notion of ordering among these classes, but we may ignore this fact and may sometimes use the following notation for convenience: $y \in \{0, 1, 2, \dots, K - 1\}$.

Examples of classification problems

An online banking service must be able to determine whether or not a transaction being performed on the site is fraudulent, on the basis of the user's IP address, past transaction history, and so forth.

An example is shown below on the *default* dataset. The annual incomes and monthly credit card balances of a number of individuals. The individuals who defaulted on their credit card payments are shown in orange, and those who did not are shown in blue.



$$\begin{cases} x_{n1} > 1200 \\ x_{n2} > 0 \end{cases}$$

$$y_n? \leftarrow \beta_0 + x_{n1}\beta_1 + x_{n2}\beta_2$$

A person arrives at the emergency room with a set of symptoms that could possibly be attributed to one of the two medical conditions. Which of the two conditions does the individual have?

$N = 100$

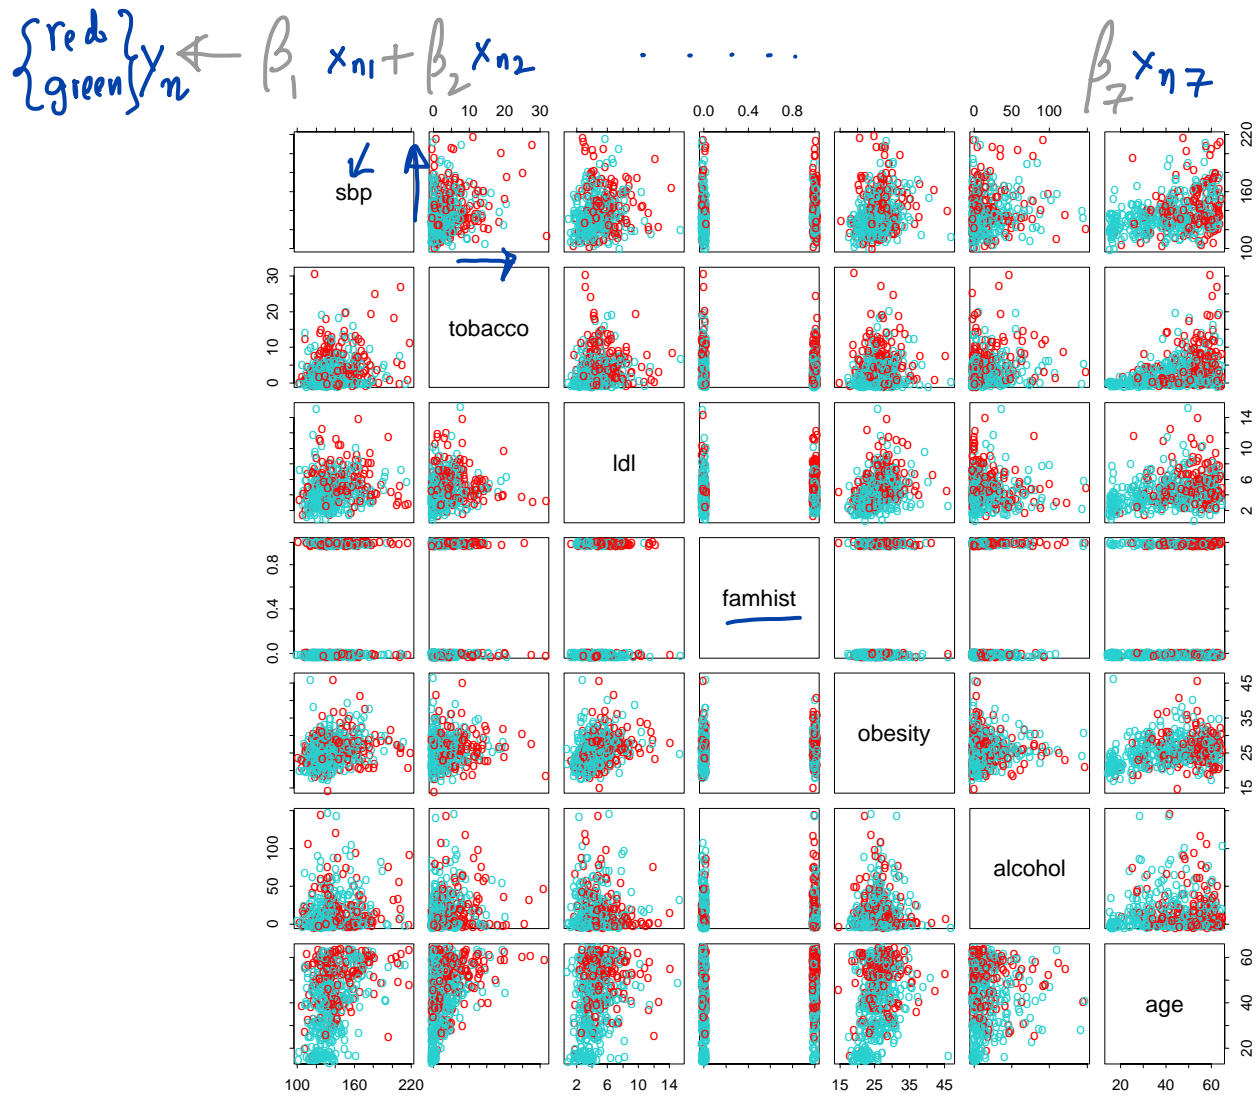


FIGURE 4.12. A scatterplot matrix of the South African heart disease data. Each plot shows a pair of risk factors, and the cases and controls are color coded (red is a case). The variable family history of heart disease (`famhist`) is binary (yes or no).

Classifier

A **classifier** will divide the input space into a collection of regions belonging to each class. The boundaries of these regions are called **decision boundaries**. A classifier can be linear or nonlinear.

Elements of Statistical Learning (2nd Ed.) ©Hastie, Tibshirani & Friedman 2009 Chap 4

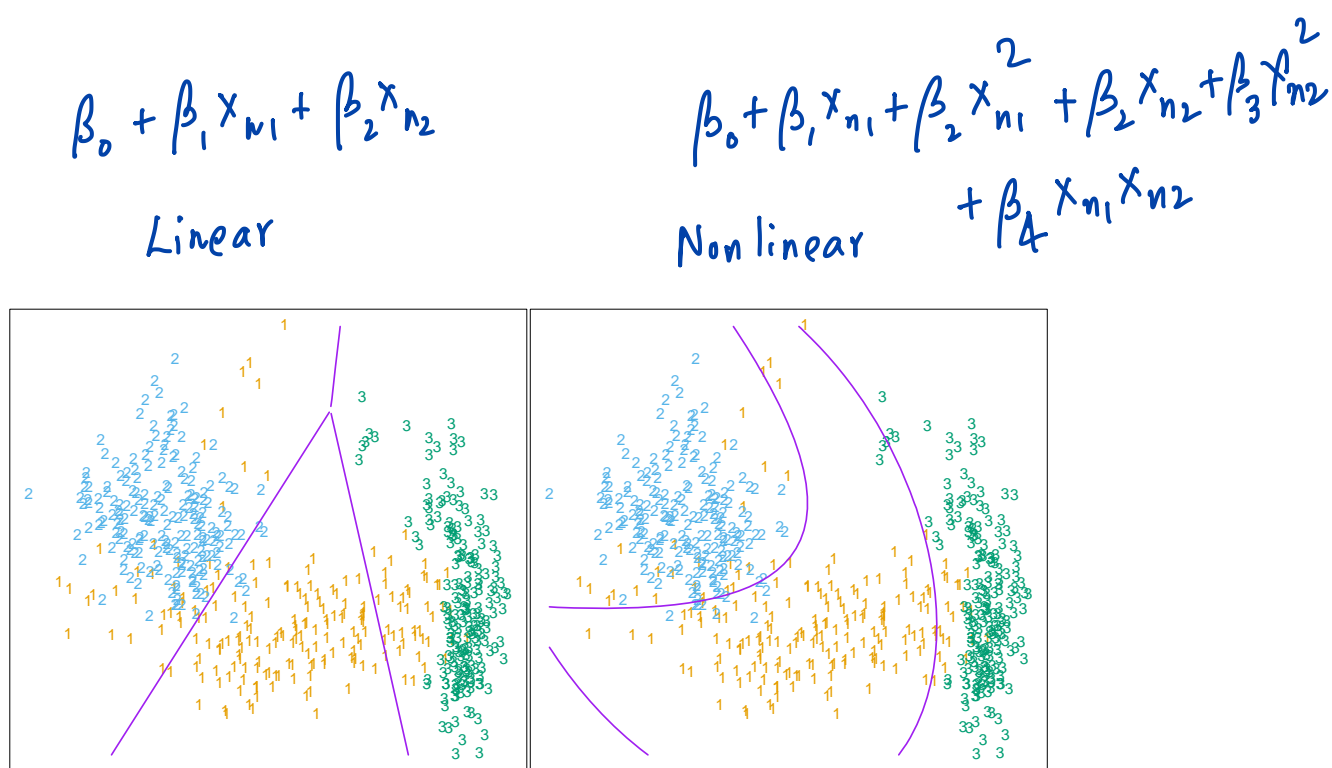


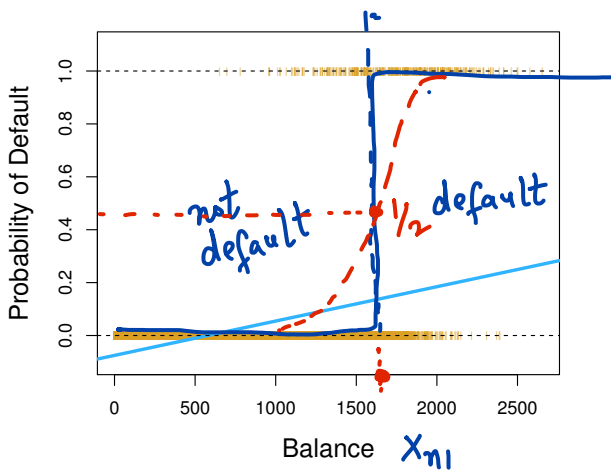
FIGURE 4.1. The left plot shows some data from three classes, with linear decision boundaries found by linear discriminant analysis. The right plot shows quadratic decision boundaries. These were obtained by finding linear boundaries in the five-dimensional space $X_1, X_2, X_1 X_2, X_1^2, X_2^2$. Linear inequalities in this space are quadratic inequalities in the original space.

2 Logistic Regression

Classification with linear regression

We can use $y = 0$ for \mathcal{C}_1 and $y = 1$ for \mathcal{C}_2 (or vice-versa), and simply use least-squares to predict \hat{y}_* given \mathbf{x}_* . We can predict \mathcal{C}_1 when $\hat{y}_* < 0.5$ and \mathcal{C}_2 when $\hat{y}_* > 0.5$.

$$\begin{matrix}
 y_1 \\
 y_2 \\
 \vdots \\
 y_N
 \end{matrix}
 \begin{bmatrix}
 0 \\
 0 \\
 1 \\
 0 \\
 0 \\
 1 \\
 0 \\
 1
 \end{bmatrix}
 \begin{matrix}
 \leftarrow x_1 \rightarrow \\
 \leftarrow x_2 \rightarrow \\
 \vdots \\
 \leftarrow x_N \rightarrow
 \end{matrix}$$



Q1: What is the simplest model?

Q2: the 2nd " " ?
 (0) $y_n = 1$... with 50% prob. 100
 (1) $y_n = \begin{cases} 0 \\ 1 \end{cases}$ " " ← 50 CATS 50 DOGS

whatis 10 90

$$\frac{1}{N} \sum y_n$$

$$\frac{50}{100} = \frac{1}{2}, \frac{10}{100} = \frac{1}{10}$$

$$(2) P(y_n = \text{CAT}) = \sigma(\beta_0)$$

$$\beta_0 = 0.5 \quad 0.1$$

Any problems with this approach?





FIGURE 2.1. A classification example in two dimensions. The classes are coded as a binary variable (BLUE = 0, ORANGE = 1), and then fit by linear regression. The line is the decision boundary defined by $x^T \hat{\beta} = 0.5$. The orange shaded region denotes that part of input space classified as ORANGE, while the blue region is classified as BLUE.

$$(3) y_n = \text{Sign} \left[\beta_1 x_{n1} + \beta_0 \right]$$

non diff.


$$P_{\text{rob}}(y_n=1) = \sigma \left(\frac{x_n^T \beta}{\sigma} \right)$$

$$\text{regression } \frac{1}{2N} \sum_{n=1}^N (y_n - \beta_0)^2$$

$$\min \rightarrow \beta_0^* = \frac{1}{N} \sum_{n=1}^N y_n$$

Logistic regression

We need to model $p(y = \mathcal{C}_1|\mathbf{x})$ and $p(y = \mathcal{C}_2|\mathbf{x})$ such that they both are > 0 and also sum to 1. For a new input \mathbf{x}_* , we can classify to \mathcal{C}_1 when $p(\hat{y}_*|\mathbf{x}_*) < 0.5$.



$$P(y = \mathcal{C}_1|\mathbf{x}) + P(y = \mathcal{C}_2|\mathbf{x}) = 1$$

$$f(\underline{x}_n) = \beta_0 + \beta_1 x_{n1} + \dots + \beta_D x_{nD} = \underline{\tilde{x}}_n^T \underline{\beta}$$

We will use the **sigmoid** **linear regression function** **logistic function**.

$$\sigma(\underline{\tilde{x}}_n^T \underline{\beta}) = \frac{1}{1 + e^{-\underline{\tilde{x}}_n^T \underline{\beta}}}$$

"linear" predictor

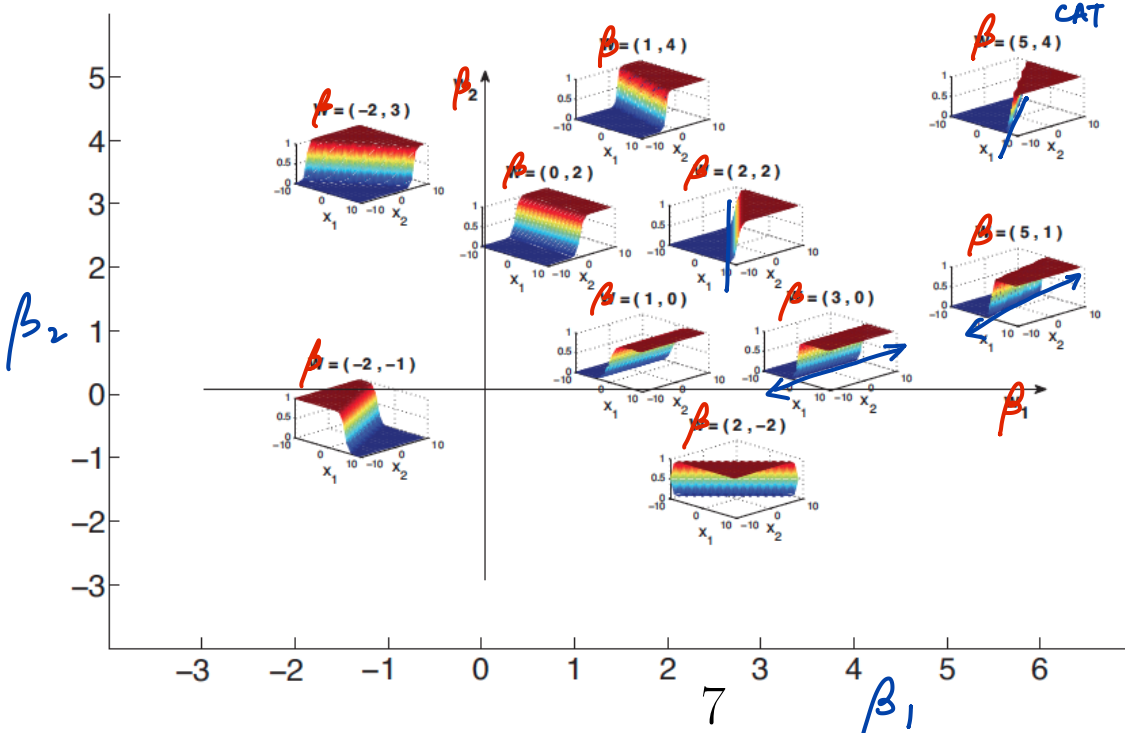

$$1 - \sigma(x) = \frac{1}{1 + \exp(x)} = 1 - \frac{e^x}{1+e^x} = \frac{1}{1+e^x}$$

We pass the linear-regression model $\eta_n = \tilde{\mathbf{x}}^T \boldsymbol{\beta}$ through the logistic function to get the probabilities.

$$\left. \begin{array}{l} \text{Lin Reg: } \hat{y}_n = \tilde{\mathbf{x}}_n^T \underline{\beta} \\ \text{Logistic Reg: } \hat{p}(y_n) = \sigma(\tilde{\mathbf{x}}_n^T \underline{\beta}) \end{array} \right\}$$

$$p(y_n = \mathcal{C}_1|\mathbf{x}_n) = \sigma(\eta_n), \quad p(y_n = \mathcal{C}_2|\mathbf{x}_n) = 1 - \sigma(\eta_n)$$

This figure visualizes the probabilities obtained for a 2-D problem (taken from KPM Chapter 7).



The probabilistic model

Assuming that each y_n is independent of others, we can define the probability of \mathbf{y} given \mathbf{X} and β :

$$p(\mathbf{y}|\mathbf{X}, \beta) = \prod_{n=1}^N p(y_n|\mathbf{x}_n)$$

$$= \prod_{n:y_n=C_1} p(y_n=C_1|\mathbf{x}_n) \prod_{n:y_n=C_2} p(y_n=C_2|\mathbf{x}_n)$$

A better way to write this is to use the coding $y_n \in \{0, 1\}$.

$$p(\mathbf{y}|\mathbf{X}, \beta) = \prod_{n=1}^N \sigma(\eta_n)^{y_n} [1 - \sigma(\eta_n)]^{1-y_n}$$

The log-likelihood is given as follows:

$$\mathcal{L}_{MLE}(\beta) = \log p(\mathbf{y}|\mathbf{X}, \beta)$$

$$= \sum_{n=1}^N y_n \log \sigma(\tilde{\mathbf{x}}_n^T \beta) + (1 - y_n) \log [1 - \sigma(\tilde{\mathbf{x}}_n^T \beta)]$$

$$\sigma(x) = \frac{e^x}{1+e^x}$$

$$\log \sigma(x) = x - \log(1+e^x)$$

Derivation:


$$= \sum_{n=1}^N y_n \tilde{\mathbf{x}}_n^T \beta - \log [1 + \exp(\tilde{\mathbf{x}}_n^T \beta)]$$

$$y_n=1, f_n - \log(1+e^{f_n}) = y_n \frac{\tilde{\mathbf{x}}_n^T \beta}{f_n} - \log(1+e^{\frac{\tilde{\mathbf{x}}_n^T \beta}{f_n}})$$

$$y_n=0, -\log(1+e^{f_n}) = y_n f_n - \log(1+e^{f_n})$$

$$f_n \rightarrow \infty$$

$$\log(1+e^{f_n}) \rightarrow f_n$$



$$\text{MSE: } \sum_{n=1}^N \frac{1}{2} (y_n - \tilde{\mathbf{x}}_n^T \beta)^2$$

$$= -\sum_{n=1}^N \log e^{-\frac{1}{2} (y_n - \tilde{\mathbf{x}}_n^T \beta)^2}$$

$$= -\log \prod_{n=1}^N e^{-\frac{1}{2} \dots}$$

$$= -\log \prod_{n=1}^N \mathcal{N}(y_n | \tilde{\mathbf{x}}_n^T \beta, 1)$$

+ const

Multiply probability of each y_n , take negative logarithm

logistic Reg:

$$-\log \prod_{n=1}^N P(y_n | \tilde{\mathbf{x}}_n^T \beta)$$

$$= \sum_{n=1}^N \log P(y_n | \tilde{\mathbf{x}}_n^T \beta)$$

$$= \sum_{\substack{n: y_n=1}} \log \sigma(\tilde{\mathbf{x}}_n^T \beta) + \sum_{\substack{n: y_n=0}} \log [1 - \sigma(\tilde{\mathbf{x}}_n^T \beta)]$$

$$\log(1 - \sigma(x)) = \log \frac{1}{1+e^x} = -\log(1+e^x)$$

$$y_n [\tilde{\mathbf{x}}_n^T \beta - \log(1+e^{\tilde{\mathbf{x}}_n^T \beta})] - (1-y_n) [\log(1+e^{\tilde{\mathbf{x}}_n^T \beta})]$$


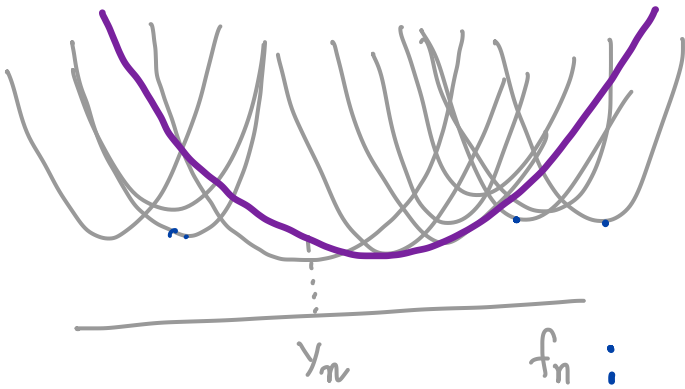
Lin Reg:

$$-\frac{1}{2} (y_n - f_n)^2 = -\frac{1}{2} (y_n^2 + f_n^2 - 2y_n f_n) = y_n f_n - \frac{1}{2} f_n^2 + \text{const}$$




Log Reg:

$$y_n f_n - \log(1 + e^{f_n})$$

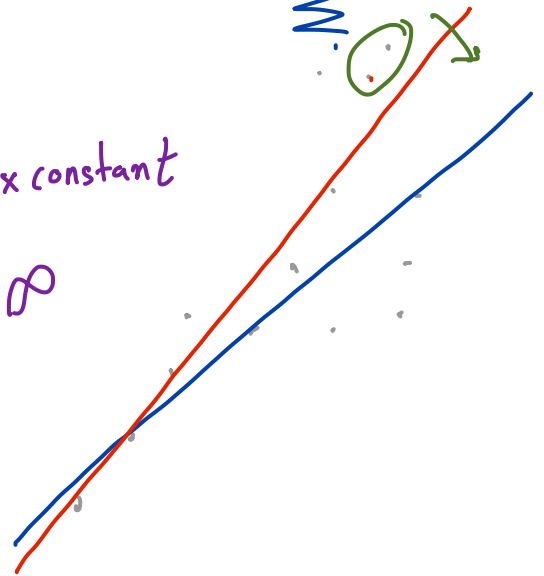
linear in f_n
log-partition function



$f_n^{(k)}$ ↓



$[\beta_0, \beta_1] \times \text{constant}$
 $\rightarrow \infty$



Optimization using SGD

We will use the following fact to derive the gradient.

$$\frac{\partial}{\partial x} \log[1 + \exp(x)] = \sigma(x)$$

Taking the gradient of the log-likelihood, we get the following stochastic gradient estimate:

$$\frac{\partial \mathcal{L}}{\partial \beta} \approx N \left[\sigma(\tilde{\mathbf{x}}_i^T \beta) - y_i \right] \tilde{\mathbf{x}}_i \quad \swarrow \text{SGD}$$

This is very similar to SGD for linear regression, i.e., the gradient is in the direction of $\tilde{\mathbf{x}}_i$ with magnitude equal to prediction error for the i 'th example.

$$l(f_n) = y_n f_n - \log(1 + e^{f_n})$$

$$\begin{aligned} \frac{\partial l}{\partial f_n} &= y_n - \frac{\partial}{\partial f_n} \log(1 + e^{f_n}) \\ &= y_n - \frac{e^{f_n}}{1 + e^{f_n}} \\ &= y_n - \underbrace{\sigma(f_n)}_{\text{error}} \end{aligned}$$

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \beta} &= \frac{\partial}{\partial \beta} \sum_{n=1}^N y_n \left(\tilde{\mathbf{x}}_n^T \beta - \log(1 + e^{\tilde{\mathbf{x}}_n^T \beta}) \right) \\ &= \sum_{n=1}^N \frac{\partial \left(\tilde{\mathbf{x}}_n^T \beta \right)}{\partial \beta} \frac{\partial \left[y_n f_n - \log(1 + e^{f_n}) \right]}{\partial f_n} \\ &= \sum_{n=1}^N \tilde{\mathbf{x}}_n \left(y_n - \underbrace{\sigma(\tilde{\mathbf{x}}_n^T \beta)}_{\text{error}} \right) \end{aligned}$$

Penalized Logistic Regression

The cost-function can be unbounded when the data is linearly separable.

Wikipedia

$$\min_{\beta} \sum_n y_n (\tilde{\mathbf{x}}_n^T \beta) - \log(1 + e^{\dots})$$


To get a well-defined problem, we will regularize.

at infinity

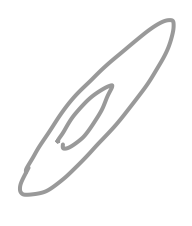
β^* x (really large number)

$$\min_{\beta} - \sum_{n=1}^N \log p(y_n | \mathbf{x}_n^T, \beta) + \lambda \sum_{d=1}^D \beta_d^2$$

$$\frac{\partial}{\partial \beta} \frac{\partial}{\partial \beta^T} = \tilde{X}^T \tilde{X} + \lambda \mathbf{I}$$

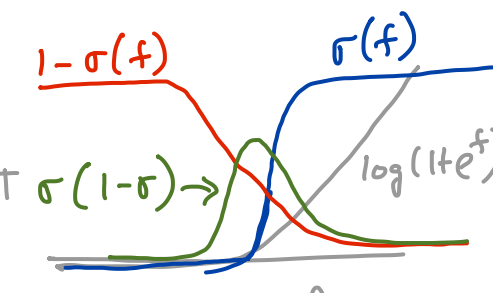


$$\sum_{n=1}^N \frac{\partial}{\partial \beta} (y_n - \tilde{x}_n^T \beta)^2 = \sum_{n=1}^N \tilde{x}_n \frac{\partial}{\partial \beta} (y_n - \tilde{x}_n^T \beta)$$

$$\sum_{n=1}^N \frac{\partial}{\partial \beta} \frac{\partial}{\partial \beta^T} \dots = \sum_{n=1}^N \tilde{x}_n \tilde{x}_n^T = \tilde{X}^T \tilde{X} = U S^2 U^T$$


$$\sum_{n=1}^N \frac{\partial}{\partial \beta} \left[y_n (\tilde{x}_n^T \beta) - \log(1 + e^{\tilde{x}_n^T \beta}) \right]$$

$$\frac{\partial}{\partial f} \log(1 + e^f) = \sigma(f)$$

$$\frac{\partial^2}{\partial f^2} \log(1 + e^f) = \sigma(f) [1 - \sigma(f)]$$


$$\frac{\partial}{\partial \beta} \frac{\partial}{\partial \beta^T} = \sum_{n=1}^N \underbrace{\sigma(\tilde{x}_n^T \beta) [1 - \sigma(\tilde{x}_n^T \beta)]}_{\text{Scalar}} \tilde{x}_n \tilde{x}_n^T$$

$$= \tilde{X}^T \Lambda \tilde{X} = U (S V^T \Lambda V S) U^T$$

$D \times N$ $N \times N$ $N \times D$

Say $\tilde{X}^T = U S V^T$

Define $\Lambda = \begin{bmatrix} \sigma(f_1)(1-\sigma(f_1)) & & 0 \\ & \ddots & \\ 0 & & \sigma(f_n)(1-\sigma(f_n)) \end{bmatrix}$

$$\begin{bmatrix} \tilde{x}_1^T \\ \vdots \\ \tilde{x}_N^T \end{bmatrix} \begin{bmatrix} \lambda_1 & & & 0 \\ & \lambda_2 & & \\ & & \lambda_3 & \\ & & & \ddots \\ 0 & & & & \lambda_N \end{bmatrix} \begin{bmatrix} \tilde{x}_1 \\ \vdots \\ \tilde{x}_N \end{bmatrix}$$


$$y_n = \tilde{x}_n^T \beta + \epsilon_n$$

\downarrow
 $\mathcal{N}(0, 1/\lambda_n)$

This can be seen as the "curvature" of a "pseudo" linear problem

$$\frac{\partial}{\partial \beta} \frac{\partial}{\partial \beta^T} \sum_{n=1}^N \frac{1}{2(1/\lambda_n)} (y_n - \tilde{x}_n^T \beta)^2 = \tilde{X}^T \Lambda \tilde{X}$$


\rightarrow noise variance



From Binary to Multi-Class

$$Y_n \in \{0, 1, 2\}$$

$$P(Y_n = k | \underline{x}_n) = S(f_n)$$



$$f_{n0} = \tilde{\underline{x}}_n^T \underline{\beta}^{(0)}$$

$$f_{n1} = \tilde{\underline{x}}_n^T \underline{\beta}^{(1)}$$

$$f_{n2} = \tilde{\underline{x}}_n^T \underline{\beta}^{(2)}$$

$$S \begin{pmatrix} f_{n0} \\ f_{n1} \\ f_{n2} \end{pmatrix} = \frac{e^{f_{nk}}}{e^{f_{n0}} + e^{f_{n1}} + e^{f_{n2}}}$$

Repeat

$$- \log \prod_{n=1}^N P(Y_n = k | \underline{x}_n)$$

$$= - \log \prod P(Y_n = 0 | \underline{x}_n) \prod P(Y_n = 1 | \underline{x}_n) \prod P(Y_n = 2 | \underline{x}_n)$$

$$= \sum_{n=1}^N \left[\underbrace{Y_n f_{nj}}_{\text{linear term}} - \underbrace{\log \sum_{k=1,2,3} e^{f_{nk}}}_{\text{log-partition}} \right]$$

$$\frac{\partial}{\partial \underline{\beta}} = \sum \dots \tilde{\underline{x}}_n (Y_n - S(f_n)) \dots$$

something like that

$$\frac{\partial}{\partial \underline{\beta} \partial \underline{\beta}^T} = \sum_{n=1}^N \tilde{\underline{x}}_n \tilde{\underline{x}}_n^T \dots$$


3 K-Nearest Neighbors (k-NN)

k-Nearest Neighbor (k-NN)

The k -NN prediction for an \mathbf{x}_* is,

$$f_k(\mathbf{x}_*) = \frac{1}{k} \sum_{\mathbf{x}_n \in nbh_k(\mathbf{x}_*)} y_n ,$$

where $nbh_k(\mathbf{x})$ is the neighborhood of \mathbf{x} defined by the k closest points \mathbf{x}_n in the training data.



We show results for $k = 1$ and $k = 15$ respectively.





FIGURE 2.3. The same classification example in two dimensions as in Figure 2.1. The classes are coded as a binary variable (BLUE = 0, ORANGE = 1), and then predicted by 1-nearest-neighbor classification.



Bias-variance revisited

How should train and test error vary with k ?




FIGURE 2.4. *Misclassification curves for the simulation example used in Figures 2.1, 2.2 and 2.3. A single training sample of size 200 was used, and a test sample of size 10,000. The orange curves are test and the blue are training error for k -nearest-neighbor classification. The results for linear regression are the bigger orange and blue squares at three degrees of freedom. The purple line is the optimal Bayes error rate.*

Curse of dimensionality

According to Pedro Domingos: “Intuitions fail in high dimensions”. This is also known as the [curse of dimensionality](#) (Bellman, 1961).

Claim 1: “Generalizing correctly becomes exponentially harder as the dimensionality grows because fixed-size training sets cover a dwindling fraction of the input space.”

The expected edge length is $e_D(r) = r^{1/D}$, e.g.

$$e_{10}(0.01) = 0.63, e_{10}(0.1) = 0.80$$

i.e. to capture 1% or 10% of the data, we must cover 63% or 80% of the range of each input variable.




FIGURE 2.6. The curse of dimensionality is well illustrated by a subcubical neighborhood for uniform data in a unit cube. The figure on the right shows the side-length of the subcube needed to capture a fraction r of the volume of the data, for different dimensions p . In ten dimensions we need to cover 80% of the range of each coordinate to capture 10% of the data.

As a result, the sampling density is proportional to $N^{1/D}$, i.e. if $N_1 = 100$ is the sample size for a 1-input problem, then $N_{10} = 100^{10}$ is required for the same sampling density with 10 inputs.

Claim 2: In high-dimension, data-points are far from each other. Consequently, “as the dimensionality increases, the choice of nearest neighbor becomes effectively random.”

Consider N data points uniformly distributed in a D -dimensional unit ball centered at the origin. We consider a nearest-neighbor estimate at the origin. The median distance from the origin to the closest data point is,

$$\left(1 - \frac{1}{2}\right)^{1/D}$$

For $N = 500$, $D = 10$, this number is 0.52, more than halfway to the boundary.

Discussion

(Taken from HTF). We will see (in the next few lectures) that there is a whole spectrum of models between the rigid linear models and the extremely flexible 1-NN model. Each model comes with their own assumptions and biases.

(Based on Domingos). You might think that gathering more input variables never hurts, since at the worst they provide no new information about the output. But in fact their benefits may be outweighed by the curse of dimensionality.

4 Deep Neural Networks

Multi-Layer Perceptron (MLP)

Also known as [feed-forward neural network](#),

$$\mathbf{x}_n \rightarrow \mathbf{a}_n^{(1)} \rightarrow \mathbf{z}_n^{(1)} \rightarrow \mathbf{a}_n^{(2)} \rightarrow \mathbf{z}_n^{(2)} \rightarrow \dots \rightarrow \mathbf{z}_n^{(K-1)} \rightarrow \mathbf{y}_n$$

where $\{y_n, \mathbf{x}_n\}$ is the n 'th input-output pair, $\mathbf{z}_n^{(k)}$ is the k 'th hidden vector, $\mathbf{a}_n^{(k)}$ is the corresponding activation. There are a total of K layers.

For the k 'th layer, we obtain the m 'th activation $a_{mn}^{(k)}$ and the corresponding hidden variable $z_{mn}^{(k)}$, as shown below:

$$a_{mn}^{(k)} = \left(\boldsymbol{\beta}_m^{(k)}\right)^T \mathbf{z}_n^{(k-1)}, \quad z_{mn}^{(k)} = h\left(a_{mn}^{(k)}\right)$$

where $\mathbf{z}_n^{(k-1)}$ is the hidden vector for the previous layer. For the first layer, we set $\mathbf{z}_n^{(0)} = \mathbf{x}_n$. For the last layer, we use a link function to map $\mathbf{z}_n^{(K-1)}$ to the output \mathbf{y}_n .

Note that a 1-Layer MLP is simply a generalization of linear/logistic regression.

Defining $\mathbf{B}^{(k)}$ as a matrix with rows $(\boldsymbol{\beta}_m^{(k)})^T$, we can express the computation of activation and hidden vectors as follows:


$$\mathbf{a}_n^{(k)} = \mathbf{B}^{(k)} \mathbf{z}_n^{(k-1)}, \quad \mathbf{z}_n^{(k)} = h\left(\mathbf{a}_n^{(k)}\right)$$

In a more compact notation, we can express the input-output relationship as follows:

$$\hat{y}_n = g\left((\boldsymbol{\beta}^{(K-1)})^T * h(\mathbf{B}^{(K-2)} * h(* \dots * h(\mathbf{B}^{(1)} * \mathbf{x}_n)))\right),$$

where g is an appropriate link function to match the output.


An illustration below shows reconstruction of the function $|x|$ at $N = 50$ data points sampled at the blue dots. The trained network has 2 layers and 3-hidden variables with $\tanh()$ activation function.



Optimization and Back-propagation

We can learn parameters \mathbf{B} using stochastic gradient-descent.

Gradient computation can be complicated due to the *deep* structure of the network. We can use [back-propagation](#) to simplify the computation. The key-idea is to express the derivatives in terms of activations $\mathbf{a}_n^{(k)}$ and hidden variables $\mathbf{z}_n^{(k)}$ using the chain rule. Below is the outline of the algorithm:



Step 1: Compute $\mathbf{a}_n^{(k)}$ and $\mathbf{z}_n^{(k)}$ using forward propagation.

Step 2: Compute $\boldsymbol{\delta}_n^{(k)} := \partial\mathcal{L}/\partial\mathbf{a}_n^{(k)}$ using backward propagation:

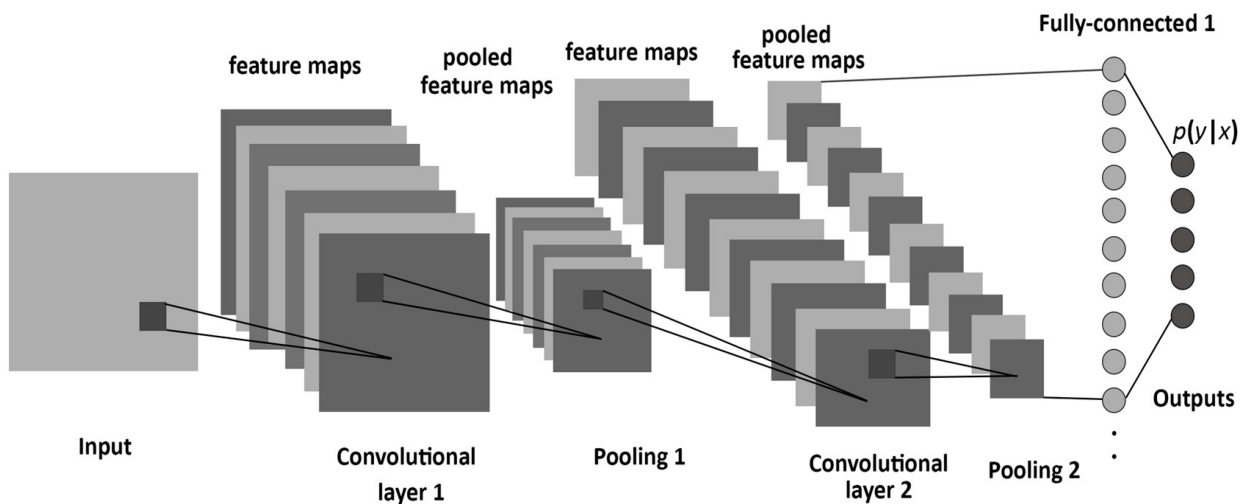
$$\boldsymbol{\delta}_n^{(k-1)} = \text{diag} \left[\mathbf{h}'(\mathbf{a}_n^{(k)}) \right] \left(\mathbf{B}^{(k)} \right)^T \boldsymbol{\delta}_n^{(k)}$$

Step 3: Compute $\partial\mathcal{L}/\partial\mathbf{B}^{(k)}$ using the above derivatives.

$$\frac{\partial\mathcal{L}}{\partial\mathbf{B}^{(k)}} = \sum_n \boldsymbol{\delta}_n^{(k)} \left(\mathbf{z}_n^{(k)} \right)^T$$

Convolutional Neural Network (CNN)

CNN produce the state-of-the-art results for image classification.



Tricks for Deep Learning

Obtaining a good generalization error with neural networks and avoiding overfitting requires a lot of hacks and tricks. A good summary of these are given in Bottou's paper "Stochastic gradient tricks". In addition, initialization seems to play a huge role in improving the performance. See the following paper "On the importance of initialization and momentum in deep learning" by Ilya Sutskever et. al.

Modern stochastic optimization methods, such as RMSprop and Adam, are currently heavily used by practitioners to optimizer neural networks.