

CSE176 Introduction to Machine Learning — Lecture notes

Miguel Á. Carreira-Perpiñán
EECS, University of California, Merced

November 28, 2016

These are notes for a one-semester undergraduate course on machine learning given by Prof. Miguel Á. Carreira-Perpiñán at the University of California, Merced. The notes are largely based on the book “Introduction to machine learning” by Ethem Alpaydm (MIT Press, 3rd ed., 2014), with some additions.

These notes may be used for educational, non-commercial purposes.

©2015–2016 Miguel Á. Carreira-Perpiñán

1 Introduction

1.1 What is machine learning (ML)?

- Data is being produced and stored continuously (“big data”):
 - science: genomics, astronomy, materials science, particle accelerators...
 - sensor networks: weather measurements, traffic...
 - people: social networks, blogs, mobile phones, purchases, bank transactions...
 - etc.
- Data is not random; it contains structure that can be used to predict outcomes, or gain knowledge in some way.

Ex: patterns of Amazon purchases can be used to recommend items.
- It is more difficult to design algorithms for such tasks (compared to, say, sorting an array or calculating a payroll). Such algorithms need data.

Ex: construct a spam filter, using a collection of email messages labelled as spam/not spam.
- Data mining: the application of ML methods to large databases.
- Ex of ML applications: fraud detection, medical diagnosis, speech or face recognition...
- ML is programming computers using data (past experience) to optimize a performance criterion.
- ML relies on:
 - Statistics: making inferences from sample data.
 - Numerical algorithms (linear algebra, optimization): optimize criteria, manipulate models.
 - Computer science: data structures and programs that solve a ML problem efficiently.
- A model:
 - is a compressed version of a database;
 - extracts knowledge from it;
 - does not have perfect performance but is a useful approximation to the data.


1.2 Examples of ML problems

- *Supervised learning*: labels provided.
 - *Classification* (pattern recognition):
 - * Face recognition. Difficult because of the complex variability in the data: pose and illumination in a face image, occlusions, glasses/beard/make-up/etc.

Training examples:




Test images:




- * Optical character recognition: different styles, slant... 0 1 2 3 4 5 6 7 8 9
- * Medical diagnosis: often, variables are missing (tests are costly).

- * Speech recognition, machine translation, biometrics...
- * Credit scoring: classify customers into high- and low-risk, based on their income and savings, using data about past loans (whether they were paid or not).
- *Regression*: the labels to be predicted are continuous:
 - * Predict the price of a car from its mileage.
 - * Navigating a car: angle of the steering.
 - * Kinematics of a robot arm: predict workspace location from angles.

if income $> \theta_1$ and savings $> \theta_2$
 then low-risk else high-risk



$$y = wx + w_0$$



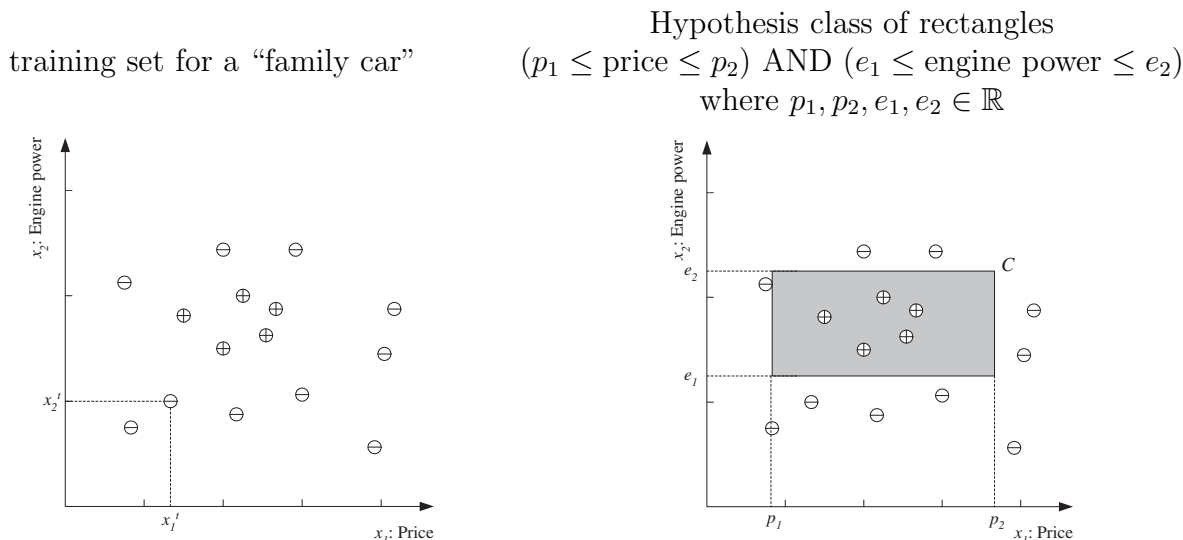
- *Unsupervised learning*: no labels provided, only input data.
 - *Learning associations*:
 - * Basket analysis: let $p(Y|X)$ = “probability that a customer who buys product X also buys product Y ”, estimated from past purchases. If $p(Y|X)$ is large (say 0.7), associate “ $X \rightarrow Y$ ”. When someone buys X , recommend them Y .
 - *Clustering*: group similar data points.
 - *Density estimation*: where are data points likely to lie?
 - *Dimensionality reduction*: data lies in a low-dimensional manifold.
 - *Feature selection*: keep only useful features.
 - *Outlier/novelty detection*.
- *Semisupervised learning*: labels provided for some points only.
- *Reinforcement learning*: find a sequence of actions (policy) that reaches a goal. No supervised output but delayed reward.

Ex: playing chess or a computer game, robot in a maze.

2 Supervised learning

2.1 Learning a class from examples: two-class problems

- We are given a training set of labeled examples (positive and negative) and want to learn a classifier that we can use to predict unseen examples, or to understand the data.
- *Input representation*: we need to decide what *attributes (features)* to use to describe the *input patterns (examples, instances)*. This implies ignoring other attributes as irrelevant.




- *Training set*: $\mathcal{X} = \{(\mathbf{x}_n, y_n)\}_{n=1}^N$ where $\mathbf{x}_n \in \mathbb{R}^D$ is the n th input vector and $y_n \in \{0, 1\}$ its class label.
- *Hypothesis (model) class \mathcal{H}* : the set of classifier functions we will use. Ideally, the true class distribution C can be represented by a function in \mathcal{H} (exactly, or with a small error).
- Having selected \mathcal{H} , learning the class reduces to finding an optimal $h \in \mathcal{H}$. We don't know the true class regions C , but we can approximate them by the *empirical error*:


$$E(h; \mathcal{X}) = \sum_{n=1}^N I(h(\mathbf{x}_n) \neq y_n) = \text{number of misclassified instances}$$

There may be more than one optimal $h \in \mathcal{H}$. In that case, we achieve better generalization by maximizing the *margin* (the distance between the boundary of h and the instances closest to it).

the hypothesis with the largest margin



noise and a more complex hypothesis



2.4 Noise

- *Noise* is any unwanted anomaly in the data. It can be due to:
 - Imprecision in recording the input attributes: \mathbf{x}_n .
 - Errors in labeling the input vectors: y_n .
 - Attributes not considered that affect the label (*hidden* or *latent* attributes, may be unobservable).
- Noise makes learning harder.
- Should we keep the hypothesis class simple rather than complex?
 - Easier to use and to train (fewer parameters, faster).
 - Easier to explain or interpret.
 - Less variance in the learned model than for a complex model (less affected by single instances), but also higher bias.

Given comparable empirical error, a simple model will generalize better than a complex one. (*Occam's razor*: simpler explanations are more plausible; eliminate unnecessary complexity.)


2.5 Learning multiple classes

- With K classes, we can code the label as an integer $y = k \in \{1, \dots, K\}$, or as a one-of- K binary vector $\mathbf{y} = (y_1, \dots, y_K)^T \in \{0, 1\}^K$ (containing a single 1 in position k).
- One approach for K -class classification: consider it as K two-class classification problems, and minimize the total empirical error:

$$E(\{h_k\}_{k=1}^K; \mathcal{X}) = \sum_{n=1}^N \sum_{k=1}^K I(h_k(\mathbf{x}_n) \neq y_{nk})$$


where \mathbf{y}_n is coded as one-of- K and h_k is the two-class classifier for problem k , i.e., $h_k(\mathbf{x}) \in \{0, 1\}$.

- Ideally, for a given pattern \mathbf{x} only one $h_k(\mathbf{x})$ is one. When no, or more than one, $h_k(\mathbf{x})$ is one then the classifier is in doubt and may reject the pattern.



regression: polynomials of order 1, 2, 6

(solve for order 1: optimal w_0, w_1)



2.6 Regression

- Training set $\mathcal{X} = \{(\mathbf{x}_n, y_n)\}_{n=1}^N$ where the label for a pattern $\mathbf{x}_n \in \mathbb{R}^D$ is a real value $y_n \in \mathbb{R}$.
In multivariate regression, $\mathbf{y}_n \in \mathbb{R}^d$ is a real vector.
- We assume the labels were produced by applying an unknown function f to the instances, and we want to learn (or estimate) that function using functions h from a hypothesis class \mathcal{H} .
Ex: \mathcal{H} = class of linear functions: $h(\mathbf{x}) = w_0 + w_1x_1 + \dots + w_Dx_D = \mathbf{w}^T\mathbf{x} + w_0$.
- *Interpolation*: we learn a function $h(\mathbf{x})$ that passes through each training pair (\mathbf{x}_n, y_n) (no noise): $y_n = h(\mathbf{x}_n)$, $n = 1, \dots, N$.
Ex: polynomial interpolation (requires a polynomial of degree $N - 1$ with N points in general position).
- *Regression*: we assume random additive noise ϵ : $y = h(\mathbf{x}) + \epsilon$.
- Empirical error: $E(h; \mathcal{X}) = \frac{1}{N} \sum_{n=1}^N (y_n - h(\mathbf{x}_n))^2$ = sum of squared errors at each instance.
Other definitions of error possible, e.g. absolute value instead of square.
- Occam's razor applies again. A line with low enough error may be preferable to a parabola with a slightly lower error.

2.7 Model selection and generalization

- Machine learning problems (classification, regression and others) are typically *ill-posed*: the observed data is finite and does not uniquely determine the classification or regression function.
- In order to find a unique solution, and learn something useful, we must make assumptions (= *inductive bias* of the learning algorithm).
Ex: the use of a hypothesis class \mathcal{H} ; the use of the largest margin; the use of the least-squares objective function.
- We can always enlarge the class of functions that can be learned by using a larger hypothesis class \mathcal{H} (higher *capacity*, or *complexity* of \mathcal{H}).
Ex: a union of rectangles; a polynomial of order N .
- How to choose the right inductive bias, in particular the right hypothesis class \mathcal{H} ? This is the *model selection problem*.
- The goal of ML is not to replicate the training data, but *to predict unseen data well*, i.e., to *generalize* well.
- For best generalization, we should match the complexity of the hypothesis class \mathcal{H} with the complexity of the function underlying the data:
 - If \mathcal{H} is less complex: *underfitting*. Ex: fitting a line to data generated from a cubic polynomial.
 - If \mathcal{H} is more complex: *overfitting*. Ex: fitting a cubic polynomial to data generated from a line.
- In summary, in ML algorithms there is a tradeoff between 3 factors:
 - the complexity $c(\mathcal{H})$ of the hypothesis class
 - the amount of training data N
 - the generalization error E

so that

- as $N \uparrow$, $E \downarrow$
- as $c(\mathcal{H}) \uparrow$, first $E \downarrow$ and then $E \uparrow$

Cross-validation Often used in practice to select among several hypothesis classes $\mathcal{H}_1, \mathcal{H}_2, \dots$. Divide the available dataset into three *disjoint* parts (say, $\frac{1}{3}$ each):

- *Training set:*

- Used to train, i.e., to fit a hypothesis $h \in \mathcal{H}_i$.
- *Optimize parameters of h given the model structure and hyperparameters.*
- Usually done with an optimization algorithm (the *learning algorithm*).

Ex: learn the weights of a neural net (with backpropagation); construct a k -nearest-neighbor classifier (by storing the training set).

- *Validation set:*

- Used to minimize the generalization error.
- *Optimize hyperparameters or model structure.*
- Usually done with a “grid search”. Ex: try all values of $H \in \{10, 50, 100\}$ and $\lambda \in \{10^{-5}, 10^{-3}, 10^{-1}\}$.

Ex: select the number of hidden units H , the architecture, the regularization parameter λ , or how long to train for a neural net; select k for a k -nearest-neighbor classifier.

- *Test set:*

- Used to report the generalization error.
- We optimize nothing on it, we just evaluate the final model on it.

Then:

1. For each class \mathcal{H}_i , fit its optimal hypothesis h_i using the training set.
2. Of all the optimal hypotheses, pick the one that is most accurate in the validation set.
We can then train the selected one on the training and validation set together (useful if we have little data altogether).
3. Report its error in the test set.

5.3 Estimation of missing values

- For any given example \mathbf{x}_n , the values of some features x_{1n}, \dots, x_{Dn} may be missing.

Ex: survey nonresponse.

- Strategies to deal with missing values in the training set:

- Discard examples having any missing values. Easy but throws away data.
- Fill in the missing values by estimating them (*imputation*):
 - * *Mean imputation*: impute feature d as its average value over the examples where it is present (for discrete features: mode imputation). Independent of the present features in \mathbf{x}_n .
 - * *Imputation by regression*: predict the value of x_{dn} from the values of the present features in \mathbf{x}_n . The regression model is learned using present values in the training set.
 - * *Multiple imputation*: create M complete training sets, each a sample from a model of the missing data, and combine the M resulting trained models.

- Whether a feature is missing for \mathbf{x}_n may depend on the values of the other features at \mathbf{x}_n .

Ex: in a census survey, rich people may wish not to give their salary.

Then, it is better to introduce a binary random variable z_{dn} that models whether x_{dn} is missing.

8.7 Outlier (novelty, anomaly) detection

- An *outlier* is an instance that is very different from the other instances in the sample. Reasons:
 - Abnormal behaviour. Fraud in credit card transactions, intruder in network traffic, etc.
 - Recording error. Faulty sensors, etc.
- Not usually cast as a two-class classification problem because there are typically few outliers and they don't fit a consistent pattern that can be easily learned.
- Instead, “one-class classification”: fit a density $p(\mathbf{x})$ to non-outliers, then consider \mathbf{x} as an outlier if $p(\mathbf{x}) < \theta$ for some threshold $\theta > 0$ (low-probability instance).
- We can also identify outliers as points that are far away from other samples.

p. 200

2.8 Dimensions of a supervised ML algorithm

- We have a sample $\mathcal{X} = \{(\mathbf{x}_n, y_n)\}_{n=1}^N$ (usually independent and identically distributed, “iid”) drawn from an unknown distribution.
- We want to learn a useful approximation to the underlying function that generated the data.
- We must choose:
 - A model $h(\mathbf{x}; \Theta)$ (hypothesis class) with parameters Θ . A particular value of Θ determines a particular hypothesis in the class.
Ex: for linear models, $\Theta =$ slope w_1 and intercept w_0 .
 - A *loss function* $L(\cdot, \cdot)$ to compute the difference between the desired output (label) y_n and our prediction to it $h(\mathbf{x}_n; \Theta)$. *Approximation error (loss)*:

$$E(\Theta; \mathcal{X}) = \sum_{n=1}^N L(y_n, h(\mathbf{x}_n; \Theta)) = \text{sum of errors over instances}$$

Ex: 0/1 loss for classification, squared error for regression.

- An optimization procedure (learning algorithm) to find parameters Θ^* that minimize the error:

$$\Theta^* = \arg \min_{\Theta} E(\Theta; \mathcal{X})$$

Different ML algorithms differ in any of these choices.

- The model, loss and learning algorithm are chosen by the ML system designer so that:
 - The model class is large enough to contain a good approximation to the underlying function that generated the data in \mathcal{X} in a noisy form.
 - The learning algorithm is efficient and accurate.
 - We must have sufficient training data to pinpoint the right model.

3 Bayesian decision theory

Probability review: appendix A.

Joint distribution: $p(X = x, Y = y)$.
Conditioning (product rule): $p(Y = y X = x) = \frac{p(X = x, Y = y)}{p(X = x)}$.
Marginalizing (sum rule): $p(X = x) = \sum_y p(X = x, Y = y)$.
Bayes' theorem: (inverse probability) $p(X = x Y = y) = \frac{p(Y = y X = x) p(X = x)}{p(Y = y)}$.

Probability theory (and Bayes' rule) is the framework for making decisions under uncertainty. It can also be used to make rational decisions among multiple actions to minimize expected risk.

3.2 Classification


- Binary classification with random variables $\mathbf{x} \in \mathbb{R}^D$ (example) and $C \in \{0, 1\}$ (label).
 - *Joint probability* $p(\mathbf{x}, C)$: how likely it is to observe an example \mathbf{x} and a class label C .
 - *Prior probability* $p(C)$: how likely it is to observe a class label C , regardless of \mathbf{x} .
 $p(C) \geq 0$ and $p(C = 1) + p(C = 0) = 1$.
 - *Class likelihood* $p(\mathbf{x}|C)$: how likely it is that, having observed an example with class label C , the example is at \mathbf{x} .
This represents how the examples are distributed for each class. We need a model of \mathbf{x} for each class.
 - *Posterior probability* $p(C|\mathbf{x})$: how likely it is that, having observed an example \mathbf{x} , its class label is C .
 $p(C = 1|\mathbf{x}) + p(C = 0|\mathbf{x}) = 1$.
This is what we need to classify \mathbf{x} . We infer it from $p(C)$ and $p(\mathbf{x}|C)$ using Bayes' rule.
 - *Evidence* $p(\mathbf{x})$: probability of observing an example \mathbf{x} at all (regardless of its class).
Marginal distribution of \mathbf{x} : $p(\mathbf{x}) = p(\mathbf{x}|C = 0)p(C = 0) + p(\mathbf{x}|C = 1)p(C = 1)$.
 - *Bayes' rule*: posterior = $\frac{\text{prior} \times \text{likelihood}}{\text{evidence}}$

$$p(C|\mathbf{x}) = \frac{p(\mathbf{x}|C)p(C)}{p(\mathbf{x})} = \frac{p(\mathbf{x}|C)p(C)}{p(\mathbf{x}|C = 0)p(C = 0) + p(\mathbf{x}|C = 1)p(C = 1)}$$

- Making a decision on a new example: given \mathbf{x} , classify it as class 1 iff $p(C = 1|\mathbf{x}) > p(C = 0|\mathbf{x})$.

- Examples:

- Gaussian classes in 1D: $p(\mathbf{x}|C_k) = \frac{1}{\sqrt{2\pi}\sigma_k} e^{-\frac{1}{2}\left(\frac{x-\mu_k}{\sigma_k}\right)^2}$.
- Exe 3.1: disease diagnosis based on a test.
Approx. as $p(d = 1|t = 1) \approx p(d = 1) \frac{p(t=1|d=1)}{p(t=1|d=0)}$ in terms of likelihood ratio. How can we increase $p(d = 1|t = 1)$?



- K -class case: $C \in \{1, \dots, K\}$.

- Prior probability: $p(C_k) \geq 0, k = 1, \dots, K$, and $\sum_{k=1}^K p(C_k) = 1$.
- Class likelihood: $p(\mathbf{x}|C_k)$.
- Bayes' rule: $p(C_k|\mathbf{x}) = \frac{p(\mathbf{x}|C_k)p(C_k)}{p(\mathbf{x})} = \frac{p(\mathbf{x}|C_k)p(C_k)}{\sum_{i=1}^K p(\mathbf{x}|C_i)p(C_i)}$.
- Choose as class $\arg \max_{k=1, \dots, K} p(C_k|\mathbf{x})$.

- We learn the distributions $p(C)$ and $p(\mathbf{x}|C)$ for each class from data using an algorithm.
- *Naive Bayes classifier*: a very simple classifier were we assume $p(\mathbf{x}|C_k) = p(x_1|C_k) \cdots p(x_D|C_k)$ for each class $k = 1, \dots, K$, i.e., within each class the features are independent from each other.

3.3 Losses and risks

- Sometimes the decision of what class to choose for \mathbf{x} based on $p(C_k|\mathbf{x})$ has a cost that depends on the class. In that case, *we should choose the class with lowest risk*.

Ex: medical diagnosis, earthquake prediction.

- Define:

- λ_{ik} : *loss (cost)* incurred for choosing class i when the input actually belongs to class k .

- *Expected risk* for choosing class i : $R_i(\mathbf{x}) = \sum_{k=1}^K \lambda_{ik} p(C_k|\mathbf{x})$.

- Decision: we choose the class with minimum risk: $\arg \min_{k=1, \dots, K} R_k(\mathbf{x})$.

Ex: $C_1 = \text{no cancer}$, $C_2 = \text{cancer}$; $(\lambda_{ik}) = \begin{pmatrix} 0 & 1000 \\ 1 & 0 \end{pmatrix}$; $p(C_1|\mathbf{x}) = 0.2$.

- Particular case: *0/1 loss*.

- Correct decisions have no loss and all incorrect decisions are equally costly:

$\lambda_{ik} = 0$ if $i = k$, 1 if $i \neq k$.

- Then? $R_i(\mathbf{x}) = 1 - p(C_i|\mathbf{x})$, hence to minimize the risk we pick the most probable class.

- Up to now we have considered K possible decisions given an input \mathbf{x} (each of the K classes). When incorrect decisions (misclassifications) are costly, it may be useful to define an additional decision of *reject* or *doubt* (and then defer the decision to a human). For the 0/1 loss:

- Define a cost $\lambda_{\text{reject},k} = \lambda \in (0, 1)$ for rejecting an input that actually belongs to class k .

- Risk of choosing class i : $R_i(\mathbf{x}) = \sum_{k=1}^K \lambda_{ik} p(C_k|\mathbf{x}) = 1 - p(C_i|\mathbf{x})$.

Risk of rejecting: $R_{\text{reject}}(\mathbf{x}) = \sum_{k=1}^K \lambda_{\text{reject},k} p(C_k|\mathbf{x}) = \lambda$.

\Rightarrow optimal decision rule (given by the minimal risk)? $\arg \max\{p(C_1|\mathbf{x}), \dots, p(C_K|\mathbf{x}), 1 - \lambda\}$.

- Extreme cases of λ :

- * $\lambda = 0$: always reject (rejecting is less costly than a correct classification).

- * $\lambda = 1$: never reject (rejecting is costlier than any misclassification).

3.4 Discriminant functions

- Classification rule: choose $\arg \max_{k=1, \dots, K} g_k(\mathbf{x})$

where we have a set of K *discriminant functions* g_1, \dots, g_K (not necessarily probability-based).

- Examples:

- Risk based on Bayes' classifier: $g_k = -R_k$.


- With the 0/1 loss:

$g_k(\mathbf{x}) = p(C_k|\mathbf{x})$, or? $g_k(\mathbf{x}) = p(\mathbf{x}|C_k)p(C_k)$, or? $g_k(\mathbf{x}) = \log p(\mathbf{x}|C_k) + \log p(C_k)$.

- Many others: SVMs, neural nets, etc.

- They divide the feature space into K decision regions $\mathcal{R}_k = \{\mathbf{x}: k = \arg \max_i g_i(\mathbf{x})\}$, $k = 1, \dots, K$. The regions are separated by decision boundaries, where ties occur.

- With two classes we can define a single discriminant? $g(\mathbf{x}) = g_1(\mathbf{x}) - g_2(\mathbf{x})$ and choose class 1 iff $g(\mathbf{x}) > 0$.



Frequently Bought Together



Price for all three: **\$174.78**

Add all three to Cart

Add all three to Wish List

Some of these items ship sooner than the others. [Show details](#)

- This item:** Introduction to Machine Learning (Adaptive Computation and Machine Learning series) by Ethem Alpaydin Hardcover **\$49.99**
- [An Introduction to Statistical Learning: with Applications in R \(Springer Texts in Statistics\)](#) by Gareth James Hardcover **\$61.97**
- [The Elements of Statistical Learning: Data Mining, Inference, and Prediction, Second Edition \(Springer ...\)](#) by Trevor Hastie Hardcover **\$62.82**

3.5 Association rules

- *Association rule*: implication of the form $X \rightarrow Y$ (X : antecedent, Y : consequent).
- Application: basket analysis (recommend product Y when a customer buys product X).
- Define the following measures of the association rule $X \rightarrow Y$:
 - $Support = p(X, Y) = \frac{\# \text{ purchases of } X \text{ and } Y}{\# \text{ purchases}}$.
For the rule to be significant, the support should be large. High support means items X and Y are frequently bought together.
 - $Confidence = p(Y|X) = \frac{p(X, Y)}{p(X)} = \frac{\# \text{ purchases of } X \text{ and } Y}{\# \text{ purchases of } X}$.
For the rule to hold with enough confidence, should be $\gg p(Y)$ and close to 1.
- Generalization to more than 2 items: $X, Z \rightarrow Y$ has confidence $p(Y|X, Z)$, etc.
- Exe 3.7: example database of purchases.
- Given a database of purchases, we want to find all possible rules having enough support and confidence.
- Algorithm *Apriori*:
 1. Find itemsets with enough support (by finding frequent itemsets).
We don't need to enumerate all possible subsets of items: for (X, Y, Z) to have enough support, all its subsets (X, Y) , (Y, Z) , (X, Z) must have enough support[?]. Hence: start by finding frequent one-item sets (by doing a pass over the database). Then, inductively, from frequent k -item sets generate $k + 1$ -item sets and check whether they have enough support (by doing a pass over the database).
 2. Convert the found itemsets into rules with enough confidence (by splitting the itemset into antecedent and consequent).
Again, we don't need to enumerate all possible subsets: for $X \rightarrow Y, Z$ to have enough confidence, $X, Y \rightarrow Z$ must have enough confidence[?]. Hence: start by considering a single consequent and test the confidence for all possible single consequents (adding as a valid rule if it has enough confidence). Then, inductively, consider two consequents for the added rules; etc.
- A more general way to establish rules (which may include hidden variables) are graphical models.

19.7 Measuring classifier performance

Binary classification problems

- The most basic measure is the classification error (or accuracy) in %: $\frac{\# \text{ misclassified patterns}}{\# \text{ patterns}}$.

- It is often of interest to distinguish the different types of errors: confusing the positive class with the negative one, or vice versa.

Ex: voice authentication to log into a user account. A false positive (allowing an impostor) is much worse than a false negative (refusing a valid user).

True class	Predicted class		
	Positive	Negative	Total
Positive	tp: true positive	fn: false negative	p
Negative	fp: false positive	tn: true negative	n
Total	p'	n'	N

- Having trained a classifier, we can control fn vs fp through a threshold $\theta \in [0, 1]$. Let C_1 be the positive class and C_2 the negative class. If the classifier returns $p(C_1|\mathbf{x})$, let us choose the positive class if $p(C_1|\mathbf{x}) > \theta$ (so $\theta = \frac{1}{2}$ gives the usual classifier):

- $\theta \approx 1$: almost always choose C_2 , nearly no false positives but nearly no true positives.
- Decreasing θ increases the number of true positives but risks introducing false positives.

- ROC curve* (“receiver operating curve”): the pair values (fp-rate, tp-rate) as a function of $\theta \in [0, 1]$. Properties:

- It is always increasing.
- An ideal classifier is at (0, 1) (top left corner).
- Diagonal (fp-rate = tp-rate): random classifier?. This is the worst we can do. Any classifier that is below the diagonal can be improved by flipping its decision?.

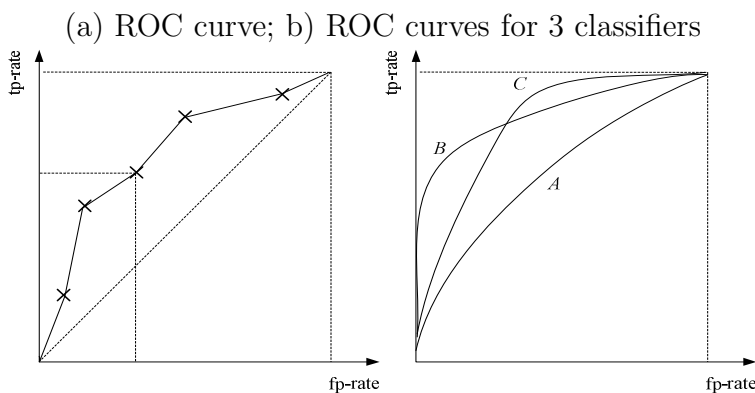
Measure	Formula
error	$(\text{fp} + \text{fn})/N$
accuracy = 1 - error	$(\text{tp} + \text{tn})/N$
tp-rate (hit rate)	tp/p
fp-rate (false alarm rate)	fp/n
precision	tp/p'
recall = tp-rate	tp/p
F-score	$\frac{\text{precision} \times \text{recall}}{(\text{precision} + \text{recall})/2}$
sensitivity = tp-rate	tp/p
specificity = 1 - fp-rate	tn/n

- Area under the curve (AUC)*: reduces the ROC curve to a number. Ideal classifier: AUC = 1.

- ROC and AUC allow us to compare classifiers over different loss conditions, and choose a value of θ accordingly. Often, there is not a dominant classifier.

- Information retrieval*: we have a database of records (e.g. documents, images) and a query, for which some of the records are relevant (“positive”) and the rest are not (“negative”). A retrieval system returns K records for the query. Its performance is measured using a *precision/recall curve*:


- Precision*: $\#(\text{relevant retrieved records}) / \#(\text{retrieved records})$.
- Recall*: $\#(\text{relevant retrieved records}) / \#(\text{relevant records})$. We can always achieve perfect recall by returning the entire database (which will contain many irrelevant records).




(a) Example ROC curve

(b) Different ROC curves for different classifiers


Precision & recall using Venn diagrams



(a) Precision and recall




(b) Precision = 1



(c) Recall = 1

$K > 2$ classes

- Again, the most basic measure is the classification error.
- *Confusion matrix*: $K \times K$ matrix where entry (i, j) contains the number of instances of class C_i that are classified as C_j .
- It allows us to identify which types of misclassification errors tend to occur, e.g. if there are two classes that are frequently confused. Ideal classifier: the confusion matrix is diagonal.



4 (Univariate) parametric methods

- How to learn probability distributions from data (in order to use them to make decisions).
- We assume such distributions follow a particular parametric form (e.g. Gaussian), so we need to estimate its parameters (μ, σ) .
- Several ways to learn them:
 - by optimizing an objective function (e.g. maximum likelihood)
 - by Bayesian estimation.
- This chapter: univariate case; next chapter: multivariate case.

Joint distribution: $p(X = x, Y = y)$.

Conditioning (product rule): $p(Y = y | X = x) = \frac{p(X = x, Y = y)}{p(X = x)}$.

Marginalizing (sum rule): $p(X = x) = \sum_y p(X = x, Y = y)$.

Bayes' theorem: (inverse probability) $p(X = x | Y = y) = \frac{p(Y = y | X = x) p(X = x)}{p(Y = y)}$.

4.2 Maximum likelihood estimation: parametric density estimation

- Problem: estimating a density $p(x)$. Assume an iid sample $\mathcal{X} = \{x_n\}_{n=1}^N$ drawn from a known probability density family $p(x; \Theta)$ with parameters Θ . We want to estimate Θ from \mathcal{X} .
- *Log-likelihood* of Θ given \mathcal{X} : $\mathcal{L}(\Theta; \mathcal{X}) = \log p(\mathcal{X}; \Theta) = \log \prod_{n=1}^N p(x_n; \Theta) = \sum_{n=1}^N \log p(x_n; \Theta)$.
- *Maximum likelihood estimate (MLE)*: $\Theta_{\text{MLE}} = \arg \max_{\Theta} \mathcal{L}(\Theta; \mathcal{X})$.
- Examples:

– Bernoulli: $\Theta = \{\theta\}$, $p(x; \theta) = \theta^x (1 - \theta)^{1-x} = \begin{cases} \theta, & x = 1 \\ 1 - \theta, & x = 0 \end{cases}$, $x \in \{0, 1\}$, $\theta \in [0, 1]$.

MLE?: $\hat{p} = \frac{1}{N} \sum_{n=1}^N x_n$ (sample average).

– Gaussian: $\Theta = \{\mu, \sigma^2\}$, $p(x; \mu, \sigma^2) = \frac{1}{\sqrt{2\pi\sigma}} e^{-\frac{1}{2}(\frac{x-\mu}{\sigma})^2}$, $x \in \mathbb{R}$, $\mu \in \mathbb{R}$, $\sigma \in \mathbb{R}^+$.

MLE?: $\hat{\mu} = \frac{1}{N} \sum_{n=1}^N x_n$ (sample average), $\hat{\sigma}^2 = \frac{1}{N} \sum_{n=1}^N (x_n - \hat{\mu})^2$ (sample variance).

For more complicated distributions, we usually need an algorithm to find the MLE.

4.4 The Bayes' estimator: parametric density estimation

- Consider the parameters Θ as random variables themselves (not as unknown numbers), and assume a *prior distribution* $p(\Theta)$ over them (based on domain information): how likely it is for the parameters to take a value *before* having observed any data.
- *Posterior distribution* $p(\Theta | \mathcal{X}) = \frac{p(\mathcal{X} | \Theta) p(\Theta)}{p(\mathcal{X})} = \frac{p(\mathcal{X} | \Theta) p(\Theta)}{\int p(\mathcal{X} | \Theta') p(\Theta') d\Theta'}$: how likely it is for the parameters to take a value *after* having observed a sample \mathcal{X} .
- Resulting estimate for the probability at a new point x ?: $p(x | \mathcal{X}) = \int p(x | \Theta) p(\Theta | \mathcal{X}) d\Theta$. Hence, rather than using the prediction of a single Θ value (“frequentist statistics”), *we average the prediction of every parameter value Θ using its posterior distribution* (“Bayesian statistics”).
- Approximations: reduce $p(\Theta | \mathcal{X})$ to a single point Θ .
 - *Maximum a posteriori (MAP) estimate*: $\Theta_{\text{MAP}} = \arg \max_{\Theta} p(\Theta | \mathcal{X})$.
Particular case: if $p(\Theta) = \text{constant}$, then $p(\Theta | \mathcal{X}) \propto p(\mathcal{X} | \Theta)$ and MAP estimate = MLE.
 - *Bayes' estimator*: $\Theta_{\text{Bayes}} = \mathbb{E} \{ \Theta | \mathcal{X} \} = \int \Theta p(\Theta | \mathcal{X}) d\Theta$.

Works well if $p(\Theta | \mathcal{X})$ is peaked around a single value.

- Example: suppose $x_n \sim \mathcal{N}(\theta, \sigma^2)$ and $\theta \sim \mathcal{N}(\mu_0, \sigma_0^2)$ where $\mu_0, \sigma^2, \sigma_0^2$ are known. Then?:

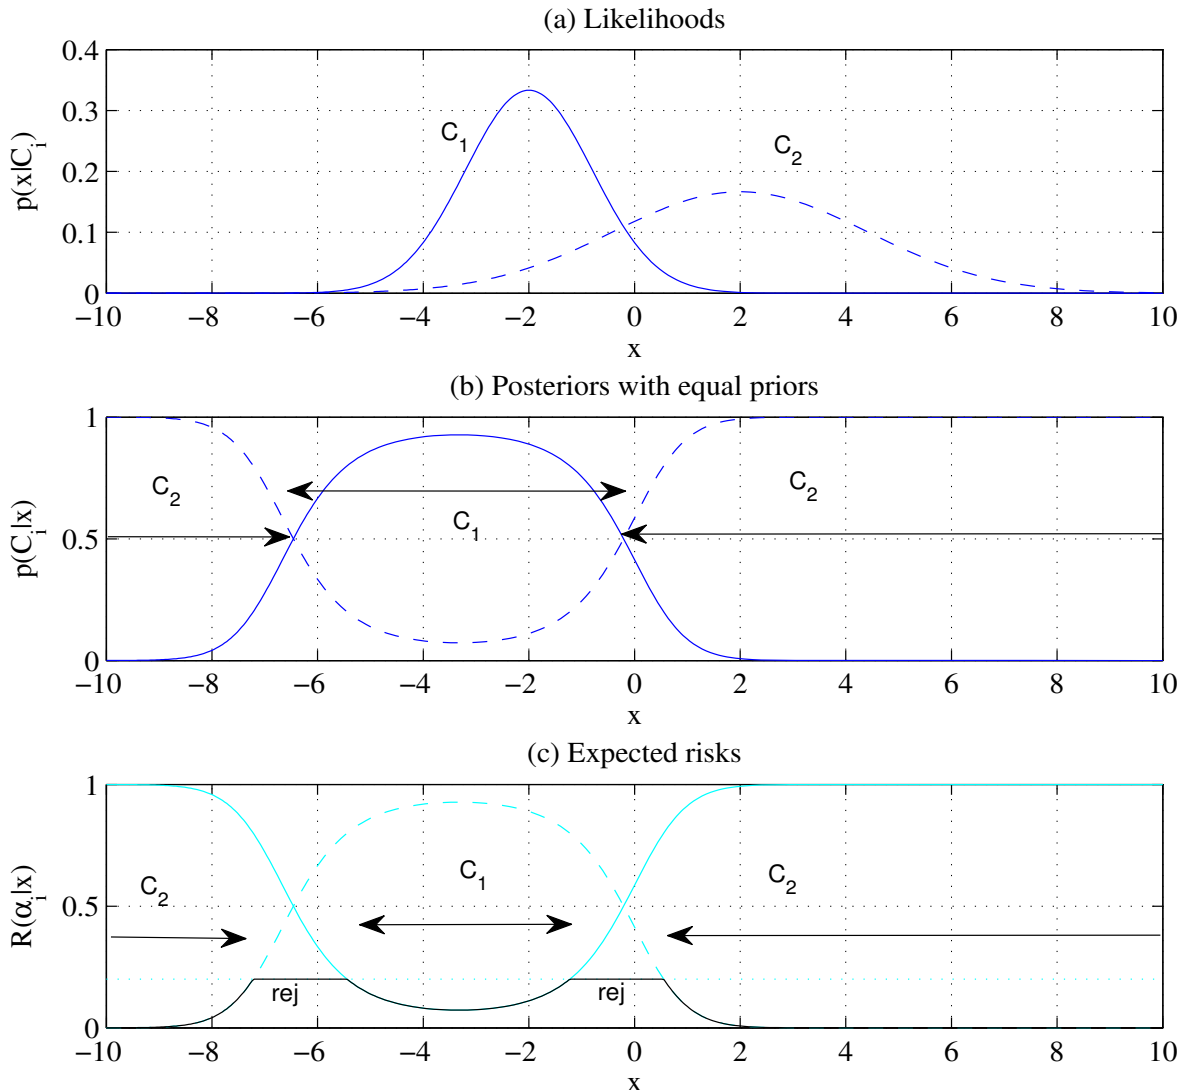
$$E\{\theta|\mathcal{X}\} = \frac{N/\sigma^2}{N/\sigma^2 + 1/\sigma_0^2} \hat{\mu} + \frac{1/\sigma_0^2}{N/\sigma^2 + 1/\sigma_0^2} \mu_0.$$
- Advantages and disadvantages of Bayesian learning:
 - ✓ works well when the sample size N is small (if the prior is helpful).
 - ✗ computationally harder (needs to compute, usually approximately, integrals or summations); needs to define a prior.

4.5 Maximum likelihood estimation: parametric classification

- From ch. 3, for classes $k = 1, \dots, K$, we use:

$$p(C_k|x) = \frac{p(x|C_k)p(C_k)}{p(x)} = \frac{p(x|C_k)p(C_k)}{\sum_{i=1}^K p(x|C_i)p(C_i)} \quad \text{discriminant function } g_k(x) = p(x|C_k)p(C_k).$$

- Ex: assume $x|C_k \sim \mathcal{N}(x; \mu_k, \sigma_k^2)$. Estimate the parameters from a data sample $\{(x_n, y_n)\}_{n=1}^N$:
 - $\hat{p}(C_k)$ = proportion of y_n that are class k .
 - $\hat{\mu}_k, \hat{\sigma}_k^2$ = as the MLE above, separately for each class k .



4.6 Maximum likelihood estimation: parametric regression



- Assume there exists an unknown function f that maps inputs x to outputs $y = f(x)$, but that what we observe as output is a noisy version $y = f(x) + \epsilon$, where ϵ is a random error. We want to estimate f by a parametric function $h(x; \Theta)$. In ch. 2 we saw the least-squares error was a good loss function to use for that purpose. We will now show that maximum likelihood estimation under Gaussian noise is equivalent to that.
- Log-likelihood of Θ given a sample $\{(x_n, y_n)\}_{n=1}^N$ drawn iid from $p(x, y)$:

$$\mathcal{L}(\Theta; \mathcal{X}) = \log \prod_{n=1}^N p(x_n, y_n) \stackrel{?}{=} \sum_{n=1}^N \log p(y_n | x_n; \Theta) + \text{constant}.$$
- Assume an error $\epsilon \sim \mathcal{N}(0, \sigma^2)$, so $p(y|x) \sim \mathcal{N}(h(x; \Theta), \sigma^2)$. Then maximizing the log-likelihood is equivalent[?] to minimizing $E(\Theta; \mathcal{X}) = \sum_{n=1}^N (y_n - h(x_n; \Theta))^2$, i.e., the least-squares error. p. 78
- Examples:

– Linear regression: $h(x; w_0, w_1) = w_1 x + w_0$. LSQ estimate[?]:


$$\mathbf{w} = \mathbf{A}^{-1} \mathbf{y}, \quad \mathbf{A} = \begin{pmatrix} 1 & \frac{1}{N} \sum_{n=1}^N x_n \\ \frac{1}{N} \sum_{n=1}^N x_n & \frac{1}{N} \sum_{n=1}^N x_n^2 \end{pmatrix}, \quad \mathbf{w} = \begin{pmatrix} w_0 \\ w_1 \end{pmatrix}, \quad \mathbf{y} = \begin{pmatrix} \frac{1}{N} \sum_{n=1}^N y_n \\ \frac{1}{N} \sum_{n=1}^N y_n x_n \end{pmatrix}.$$

– Polynomial regression: $h(x; w_0, \dots, w_k) = w_k x^k + \dots + w_1 x + w_0$. The model is still linear on the parameters. LSQ estimate also of the form $\mathbf{w} = \mathbf{A}^{-1} \mathbf{y}$. p. 79



4.3 Evaluating an estimator: bias and variance

- Statistic* $d(\mathcal{X})$: any value that is calculated from a sample \mathcal{X} (e.g. average, maximum...). It is a r.v. with an expectation (over samples) $E_{\mathcal{X}} \{d(\mathcal{X})\}$ and a variance $E_{\mathcal{X}} \{(d(\mathcal{X}) - E_{\mathcal{X}} \{d(\mathcal{X})\})^2\}$.
- $\mathcal{X} = (x_1, \dots, x_N)$ iid sample from $p(x; \theta)$. Let $d(\mathcal{X})$ be an *estimator* for θ . How good is it?
mean square error of the estimator d: $\text{error}(d, \theta) = E_{\mathcal{X}} \{(d(\mathcal{X}) - \theta)^2\}$.
- Bias* of the estimator: $b_{\theta}(d) = E_{\mathcal{X}} \{d(\mathcal{X})\} - \theta$. How much the expected value of the estimator over samples differs from the true parameter value.
 If $b_{\theta}(d) = 0$ for all θ values: *unbiased estimator*.
 Ex: the sample average $\frac{1}{N} \sum_{n=1}^N x_n$ is an unbiased estimator of the true mean μ [?].
- Variance* of the estimator: $\text{var} \{d\} = E_{\mathcal{X}} \{(d(\mathcal{X}) - E_{\mathcal{X}} \{d(\mathcal{X})\})^2\}$. How much the estimator varies around its expected value from one sample to another.
 If $\text{var} \{d\} \rightarrow 0$ as $N \rightarrow \infty$: *consistent estimator*.
 Ex: the sample average is a consistent estimator of the true mean[?].



- *Bias-variance decomposition:*

$$\text{error}(d, \theta) = \mathbb{E}_{\mathcal{X}} \{(d - \theta)^2\} = \underbrace{\mathbb{E}_{\mathcal{X}} \{(d - \mathbb{E}_{\mathcal{X}} \{d\})^2\}}_{\text{variance}} + \underbrace{(\mathbb{E}_{\mathcal{X}} \{d\} - \theta)^2}_{\text{bias}^2} = \text{var} \{d\} + b_{\theta}^2(d).$$



bias↓, var↓



bias↓, var↑



bias↑, var↓

all over the place
outside the picture

bias↑, var↑

4.7 Tuning model complexity: bias-variance dilemma

- Consider the regression setting as before, i.e., $y = f(x) + \epsilon$. Given a sample $\mathcal{X} = \{(x_n, y_n)\}_{n=1}^N$ drawn iid from $p(x, y)$, we construct a regression estimate $h(x)$. Then, for any h :

$$\text{Expected square error at } x: \mathbb{E} \{(y - h(x))^2 | x\} \stackrel{?}{=} \underbrace{\mathbb{E} \{(y - \mathbb{E} \{y|x\})^2 | x\}}_{\text{noise}} + \underbrace{(\mathbb{E} \{y|x\} - h(x))^2}_{\text{squared error}}.$$

The expectations are over the joint density $p(x, y)$ for fixed x , or, equivalently[?], over $p(y|x)$.

- *Noise term:* variance of y given x . Equal to the variance of the noise ϵ added, independent of h or \mathcal{X} . Can never be removed no matter which estimator we use.
- *Squared error term:* how much the estimate $h(x)$ deviates (at each x) from the regression function $\mathbb{E} \{y|x\}$. Depends on h and \mathcal{X} . It is zero if $h(x) = \mathbb{E} \{y|x\}$ for each x .

The ideal[?], optimal regression function is $h(x) = f(x) = \mathbb{E} \{y|x\}$ (*conditional mean* of $p(y|x)$).

- Consider $h(x)$ as an estimate (for each x) of the true $f(x)$. Bias-variance decomposition at x :

$$\mathbb{E}_{\mathcal{X}} \{(\mathbb{E} \{y|x\} - h(x))^2 | x\} = \underbrace{(\mathbb{E} \{y|x\} - \mathbb{E}_{\mathcal{X}} \{h(x)\})^2}_{\text{bias}^2} + \underbrace{\mathbb{E}_{\mathcal{X}} \{(h(x) - \mathbb{E}_{\mathcal{X}} \{h(x)\})^2\}}_{\text{variance}}.$$

The expectation $\mathbb{E}_{\mathcal{X}} \{\cdot\}$ is over samples \mathcal{X} of size N drawn from $p(x, y)$. The other expectations are over $p(y|x)$.

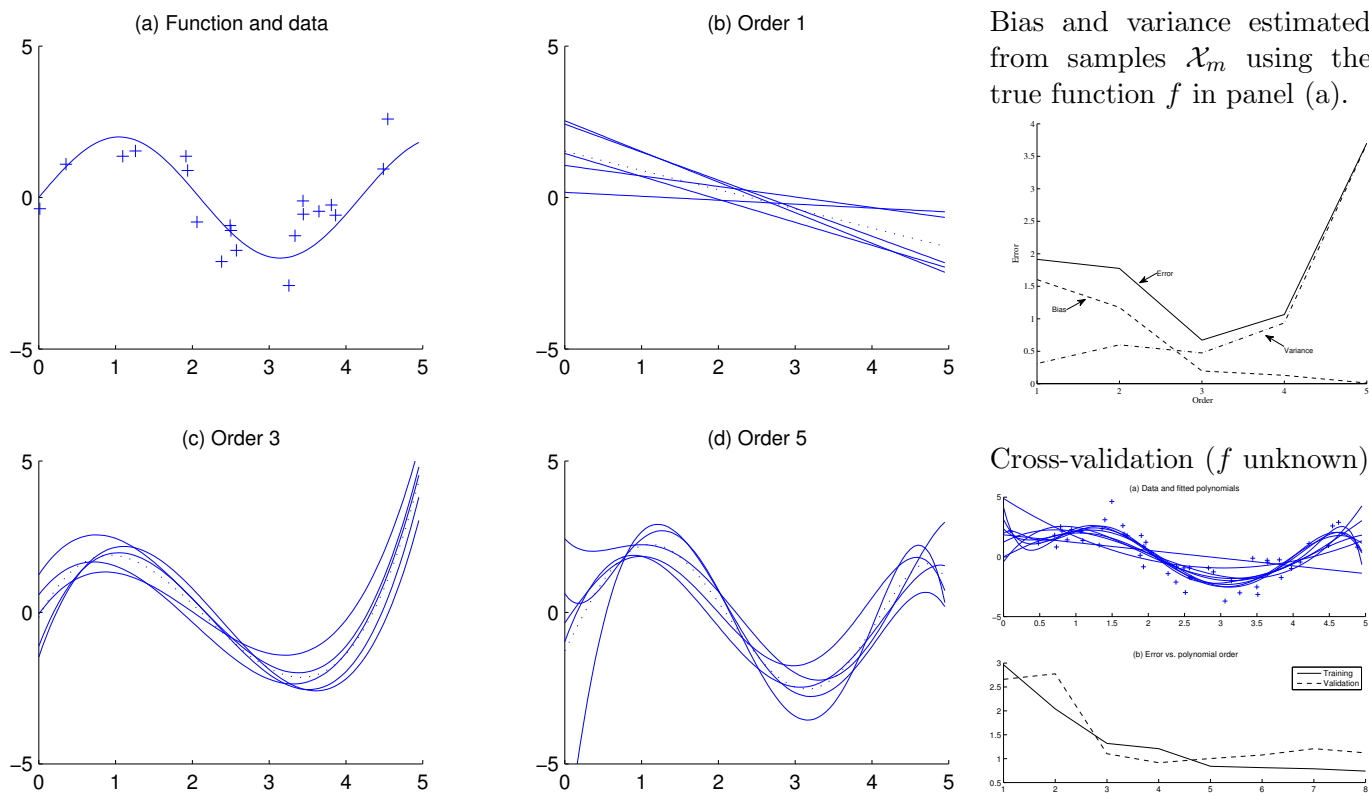
- *Bias:* how much $h(x)$ is wrong on average over different samples.
- *Variance:* how much $h(x)$ fluctuates around its expected value as the sample varies.

We want both to be small. Ex: let h_m be the estimator using a sample \mathcal{X}_m :

- $h_m(x) = 2 \forall x$ (constant fit): zero var, high bias (unless $f(x) \approx 2 \forall x$), high total error.
- $h_m(x) = \frac{1}{N} \sum_{n=1}^N y_n^{(m)}$ (average of sample \mathcal{X}_m): higher var, lower bias, lower total error.
- $h_m(x) = \text{polynomial of degree } k$: if $k \uparrow$ then bias↓, var↑.

- Low bias requires sufficiently flexible models, but flexible models have high variance. As we increase complexity, bias decreases (better fit to data) and variance increases (fit varies more with data). The optimal model has the best trade-off between bias and variance.

- *Underfitting:* model class doesn't contain the solution (it is not flexible enough) \Rightarrow bias.
- *Overfitting:* model class too general and also learns the noise \Rightarrow variance.
Also, the variance due to the sample decreases as the sample size increases.



4.8 Model selection procedures

- Methods to find the model complexity that is best suited to the data.
- *Cross-validation*: the method most used in practice. We cannot calculate the bias and variance (since we don't know the true function f), but we can estimate the total error.
 - Given a dataset \mathcal{X} , divide it into two parts as training and validation sets, \mathcal{T} and \mathcal{V} .
 - Train candidate models of different complexities on \mathcal{T} .
 - Evaluate their errors on \mathcal{V} .
 - Pick the model that gives lowest error on \mathcal{V} .

This works because, as the model complexity increases:

- the training error keeps decreasing;
- the validation error first decreases then increases (or stays about constant).
- *Regularization*: instead of minimizing just the error on the data, minimize

$$\min_h \sum_{n=1}^N (y_n - h(x_n))^2 + \lambda C(h)$$

where $C(h) \geq 0$ measures the model complexity and $\lambda \geq 0$. Ex:

- Model selection criteria (AIC, BIC...): $C(h) \propto$ number of parameters in h .
- Smoothness penalty: e.g. $C(h) = \sum_{i=1}^k w_i^2$ for polynomials of degree k .

λ is set by cross-validation.

5 Multivariate parametric methods

- Sample $\mathcal{X} = \{\mathbf{x}_n\}_{n=1}^N$ where each example \mathbf{x}_n contains D features (continuous or discrete).
- We often write the sample (or dataset) as a matrix $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_N)$ (examples = columns).
Sometimes as its transpose (examples = rows).

Review of important moments

For a sample:

- *Mean vector*: $\boldsymbol{\mu} = \frac{1}{N} \sum_{n=1}^N \mathbf{x}_n$.
- *Covariance matrix*: $\boldsymbol{\Sigma} = \frac{1}{N} \sum_{n=1}^N (\mathbf{x}_n - \boldsymbol{\mu})(\mathbf{x}_n - \boldsymbol{\mu})^T \stackrel{?}{=} \frac{1}{N} \sum_{n=1}^N \mathbf{x}_n \mathbf{x}_n^T - \boldsymbol{\mu} \boldsymbol{\mu}^T$.
Symmetric, positive definite, diagonal = variances $\sigma_{11} = \sigma_1^2, \dots, \sigma_{DD} = \sigma_D^2$.
- *Correlation matrix*: $\text{corr}(X_i, X_j) = \frac{\sigma_{ij}}{\sigma_i \sigma_j} \in [-1, 1]$.
If $\text{corr}(X_i, X_j) = 0$ (equivalently $\sigma_{ij} = 0$) then X_i and X_j are uncorrelated (but not necessarily independent).
If $\text{corr}(X_i, X_j) = \pm 1$ then X_i and X_j are linearly related.

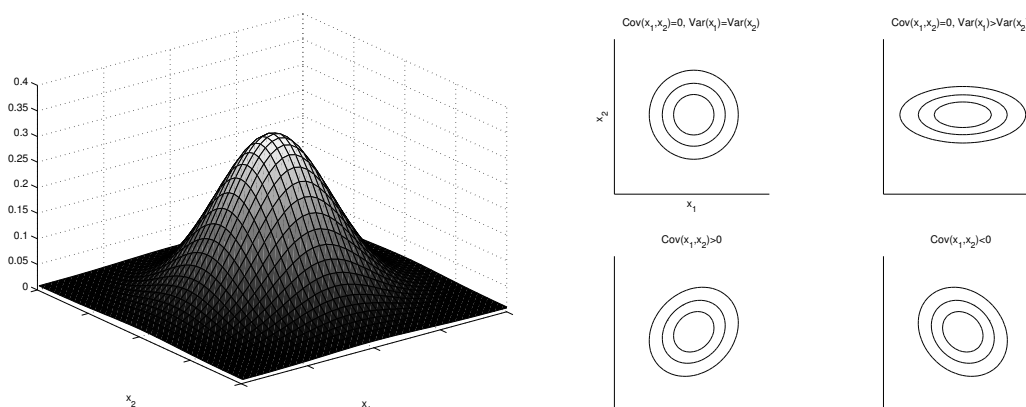
For a continuous distribution (for a discrete distribution, replace $f \rightarrow \sum$):

- *Mean vector*: $\mathbf{E}\{\mathbf{x}\} = \int \mathbf{x} p(\mathbf{x}) d\mathbf{x}$.
- *Covariance matrix*: $\text{cov}\{\mathbf{x}\} = \mathbf{E}\{(\mathbf{x} - \mathbf{E}\{\mathbf{x}\})(\mathbf{x} - \mathbf{E}\{\mathbf{x}\})^T\} \stackrel{?}{=} \mathbf{E}\{\mathbf{x}\mathbf{x}^T\} - \mathbf{E}\{\mathbf{x}\}\mathbf{E}\{\mathbf{x}\}^T$.

5.4 Review of the multivariate normal (Gaussian) distribution

- Density at $\mathbf{x} \in \mathbb{R}^D$: $p(\mathbf{x}) = |2\pi\boldsymbol{\Sigma}|^{-1/2} e^{-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1}(\mathbf{x}-\boldsymbol{\mu})}$. $\mathbf{E}\{\mathbf{x}\} = \boldsymbol{\mu}$ and $\text{cov}\{\mathbf{x}\} = \boldsymbol{\Sigma}$.
- *Mahalanobis distance*: $d_{\boldsymbol{\Sigma}}(\mathbf{x}, \mathbf{x}') = (\mathbf{x} - \mathbf{x}')^T \boldsymbol{\Sigma}^{-1}(\mathbf{x} - \mathbf{x}')$. $\boldsymbol{\Sigma} = \mathbf{I} \Rightarrow d_{\boldsymbol{\Sigma}}(\mathbf{x}, \mathbf{x}') = \|\mathbf{x} - \mathbf{x}'\|^2$ (Euclidean distance).
- Properties:
 - $\boldsymbol{\Sigma}$ diagonal \Rightarrow the D features are uncorrelated and independent: $p(\mathbf{x}) = p(x_1) \cdots p(x_D)$.
 - \mathbf{x} is Gaussian \Rightarrow each marginal $p(x_i)$ and conditional distribution $p(x_i|x_j)$ is also Gaussian.
 - $\mathbf{x} \sim \mathcal{N}_D(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, $\mathbf{W}_{D \times K} \Rightarrow \mathbf{W}^T \mathbf{x} \sim \mathcal{N}_K(\mathbf{W}^T \boldsymbol{\mu}, \mathbf{W}^T \boldsymbol{\Sigma} \mathbf{W})$. A linear projection of a Gaussian is Gaussian.
- Widely used in practice because:
 - Its special mathematical properties simplify the calculations.
 - Many natural phenomena are approximately Gaussian.
 - It models well blobby distributions centered around a prototype vector $\boldsymbol{\mu}$ with a noise characterized by $\boldsymbol{\Sigma}$.

It does make strong assumptions: symmetry, unimodality, non-heavy tails.



5.5 Multivariate classification with Gaussian classes

Assume the class-conditional densities are Gaussian: $\mathbf{x}|C_k \sim \mathcal{N}(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)$. The maths carry over from the 1D case in a straightforward way:

p. 100

- MLE[?]: class proportions for $p(C_k)$; sample mean and sample covariance for $\boldsymbol{\mu}_k$ and $\boldsymbol{\Sigma}_k$, resp.
- Discriminant function $g_k(\mathbf{x}) = \log p(\mathbf{x}|C_k) + \log p(C_k) \stackrel{?}{=} \text{quadratic form on } \mathbf{x}$.
- Number of parameters per class $k = 1, \dots, K$: $p(C_k)$: 1; $\boldsymbol{\mu}_k$: D ; $\boldsymbol{\Sigma}_k$: $\frac{D(D+1)}{2}$?

A lot of parameters for the covariances! Estimating them may need a large dataset.

Special cases, having fewer parameters:

- Equal (“shared”) covariances: $\boldsymbol{\Sigma}_k = \boldsymbol{\Sigma} \forall k$. Total $\frac{D(D+1)}{2}$ parameters (instead of $K \frac{D(D+1)}{2}$).
 - \Rightarrow MLE for $\boldsymbol{\Sigma} = \sum_{k=1}^K p(C_k) \boldsymbol{\Sigma}_k$, where $\boldsymbol{\Sigma}_k = \text{covariance matrix of class } k$.
 - \Rightarrow the discriminant $g_k(\mathbf{x})$ becomes linear on \mathbf{x} .
 - If, in addition, equal priors: $p(C_k) = \frac{1}{K} \forall k$: **Mahalanobis distance classifier**
 \Rightarrow assign \mathbf{x} to the closest centroid $\boldsymbol{\mu}_k$ in Mahalanobis distance $(\mathbf{x} - \boldsymbol{\mu}_k)^T \boldsymbol{\Sigma}^{-1} (\mathbf{x} - \boldsymbol{\mu}_k)$.

p. 103

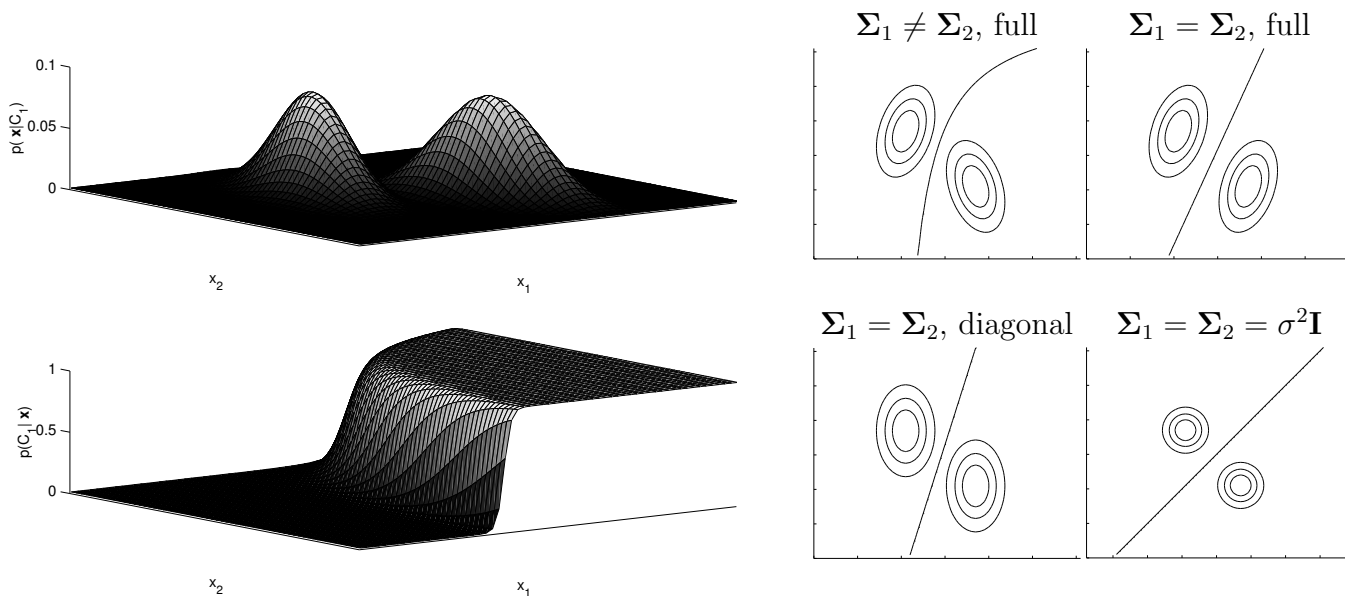
- Equal, isotropic covariances: $\boldsymbol{\Sigma}_k = \sigma^2 \mathbf{I} \forall k$. Total 1 parameter for all K covariances.
 - \Rightarrow MLE for $\sigma^2 = \frac{1}{D} \sum_{d=1}^D \sum_{k=1}^K p(C_k) \sigma_{kd}^2$, where $\sigma_{kd}^2 = d\text{th diagonal element of } \boldsymbol{\Sigma}_k$.
 - If, in addition, equal priors: **Euclidean distance classifier (“template matching”)**
 $\stackrel{?}{\Rightarrow} g_k(\mathbf{x}) = \|\mathbf{x} - \boldsymbol{\mu}_k\|^2 \Leftrightarrow g_k(\mathbf{x}) = \boldsymbol{\mu}_k^T \mathbf{x} - \frac{1}{2} \|\boldsymbol{\mu}_k\|^2$.
 If $\|\boldsymbol{\mu}_k\|^2 = 1 \Rightarrow g_k(\mathbf{x}) = \boldsymbol{\mu}_k^T \mathbf{x}$: **dot product classifier**.

p. 105

- Diagonal covariances (not shared): $\boldsymbol{\Sigma}_k = \text{diag}(\sigma_{k1}^2, \dots, \sigma_{kD}^2)$. Total KD parameters.

This assumes the features are independent within each class, and gives a naive Bayes classifier.

- \Rightarrow MLE for $\sigma_{kd}^2 = d\text{th diagonal element of } \boldsymbol{\Sigma}_k = \text{variance of class } k \text{ for feature } d$.



5.6 Tuning complexity

- Again a bias-variance dilemma:
 - Two few parameters (e.g. $\Sigma_1 = \Sigma_2 = \sigma^2 \mathbf{I}$): high bias.
 - Too many parameters (e.g. $\Sigma_1 \neq \Sigma_2$, full): high variance.
- We can use cross-validation, regularization, Bayesian priors on the covariances, etc.

5.7 Discrete features with Bernoulli distributions

- Assume that, in each class k , the features are independent and each feature $d = 1, \dots, D$ is Bernoulli with parameter θ_{kd} :

$$p(\mathbf{x}|C_k) = \prod_{d=1}^D \theta_{kd}^{x_d} (1 - \theta_{kd})^{1-x_d} \quad \mathbf{x} \in \{0, 1\}^D.$$

- This is a naive Bayes classifier. Its discriminant $g_k(\mathbf{x})$ is linear? p. 108
- MLE: class proportions for $p(C_k)$; sample mean for θ_{kd} .
- Works well with *document categorization* (e.g. classifying news reports into politics, sports and fashion) and *spam filtering* (classifying email messages into spam or non-spam). We typically represent a document as a *bag of words* vector \mathbf{x} : given a predetermined dictionary of D words, \mathbf{x} is a binary vector of dimension D where $x_d = 1$ iff word d is in the document.

5.8 Multivariate regression

- Linear regression where $\mathbf{x} \in \mathbb{R}^D$: $y = f(\mathbf{x}) + \epsilon$ with $f(\mathbf{x}) = w_D x_D + \dots + w_1 x_1 + w_0 = \mathbf{w}^T \mathbf{x}$ (where we define $x_0 = 1$).
- The maths carry over from the 1D case in a straightforward way: p. 110
 - Least-squares error (or equivalently maximum likelihood where ϵ is Gaussian with zero mean and constant variance): $E(\mathbf{w}) = \sum_{n=1}^N (y_n - \mathbf{w}^T \mathbf{x}_n)^2$.
 - The error E is quadratic on \mathbf{w} . Equating the derivatives of E wrt \mathbf{w} to $\mathbf{0}$ we obtain the *normal equations* (a linear system for \mathbf{w}):

$$\frac{\partial E}{\partial \mathbf{w}} = -2 \sum_{n=1}^N (y_n - \mathbf{w}^T \mathbf{x}_n) \mathbf{x}_n = \mathbf{0} \Rightarrow \left(\sum_{n=1}^N \mathbf{x}_n \mathbf{x}_n^T \right) \mathbf{w} = \sum_{n=1}^N y_n \mathbf{x}_n \Rightarrow (\mathbf{X}\mathbf{X}^T) \mathbf{w} = \mathbf{X}\mathbf{y}$$

where $\mathbf{X}_{D \times N} = (\mathbf{x}_1, \dots, \mathbf{x}_N)$, $\mathbf{w}_{D \times 1} = (w_0, \dots, w_D)^T$, $\mathbf{y}_{N \times 1} = (y_1, \dots, y_N)^T$.

- The resulting values of the model parameters can give insights into the data:
 - the sign of w_d tells us whether x_d has a positive or negative effect on y ;
 - the magnitude of w_d tells us how influential x_d is (if x_1, \dots, x_D all have the same range).
- With multiple outputs $\mathbf{y} = (y_1, \dots, y_{D'})^T$, the linear regression $\mathbf{y} = \mathbf{f}(\mathbf{x}) + \boldsymbol{\epsilon}$ is equivalently defined as D' independent single-output regressions.

$E(\mathbf{W}) = \sum_{n=1}^N \|\mathbf{y}_n - \mathbf{W}\mathbf{x}_n\|^2$. Taking $\frac{\partial E}{\partial \mathbf{W}} = \mathbf{0}$ gives $(\mathbf{X}\mathbf{X}^T)\mathbf{W} = \mathbf{X}\mathbf{Y}^T$ with $\mathbf{Y}_{D' \times N} = (\mathbf{y}_1, \dots, \mathbf{y}_N)$ and $\mathbf{W}_{D' \times D}$.

- We can fit a nonlinear function $h(\mathbf{x})$ similarly to a linear regression by defining additional, nonlinear features such as $x_2 = x^2$, $x_3 = e^x$, etc. (just as happens with polynomial regression). Then, using a linear model in the augmented space $\mathbf{x} = (x, x^2, e^x)^T$ will correspond to a nonlinear model in the original space x . Radial basis function networks, kernel SVMs...

6 Dimensionality reduction

- If we want to train a classifier (or regressor) on a sample $\{(\mathbf{x}_n, y_n)\}_{n=1}^N$ where the number of features D in \mathbf{x} (the dimension of \mathbf{x}) is large:
 - Training will be slow.
 - Learning a good classifier will require a large sample.
 - It is then convenient to transform each example $\mathbf{x}_n \in \mathbb{R}^D$ into a new example $\mathbf{z}_n = \mathbf{F}(\mathbf{x}_n) \in \mathbb{R}^L$ having lower dimension $L < D$ (as long as we don't lose much information).
 - Two basic ways to do this:
 - **Feature selection:** $\mathbf{F}(\mathbf{x}) = \text{a subset of } x_1, \dots, x_D$. $\mathbf{x} = (x_1, x_2, x_3, x_4, x_5)^T \Rightarrow \text{for } L = 2:$
 - $\mathbf{F}(\mathbf{x}) = \begin{pmatrix} x_2 \\ x_5 \end{pmatrix}$.

It doesn't modify the features, it simply selects L and discards the rest.

Ex: best-subset/forward/backward selection.
 - **Dimensionality reduction (DR):** $\mathbf{F}(\mathbf{x}) = \text{a l.c. or some other function of all the } x_1, \dots, x_D$. $\mathbf{F}(\mathbf{x}) = \begin{pmatrix} 1x_1+3x_2-5x_3+5x_4-4x_5 \\ 2x_1+3x_2-1x_3+0x_4+2x_5 \end{pmatrix}$.
 - It constructs L new features and discards the original D features.
 - Ex: PCA, LDA...
- If reducing to $L \leq 3$ dimensions, can visualize the dataset and look for patterns (clusters, etc.).

Review of eigenvalues and eigenvectors

For a real symmetric matrix \mathbf{A} of $D \times D$:

- *Eigenvalues* and *eigenvectors* of: $\mathbf{A}\mathbf{u} = \lambda\mathbf{u} \Rightarrow \begin{cases} \lambda \in \mathbb{R}: & \text{eigenvalue} \\ \mathbf{u} \in \mathbb{R}^D: & \text{eigenvector.} \end{cases}$
- \mathbf{A} has D eigenvalues $\lambda_1 \geq \dots \geq \lambda_D$ and D corresponding eigenvectors $\mathbf{u}_1, \dots, \mathbf{u}_D$.
- Eigenvectors of different eigenvalues are orthogonal: $\mathbf{u}_i^T \mathbf{u}_j = 0$ if $i \neq j$.
- \mathbf{A} is $\begin{cases} \text{nonsingular:} & \text{all } \lambda \neq 0 \\ \text{positive definite:} & \text{all } \lambda > 0 \quad (\Leftrightarrow \mathbf{x}^T \mathbf{A} \mathbf{x} > 0 \quad \forall \mathbf{x} \neq \mathbf{0}) \\ \text{positive semidefinite:} & \text{all } \lambda \geq 0 \quad (\Leftrightarrow \mathbf{x}^T \mathbf{A} \mathbf{x} \geq 0 \quad \forall \mathbf{x} \neq \mathbf{0}). \end{cases}$
- *Spectral theorem:* \mathbf{A} symmetric, real with normalized eigenvectors $\mathbf{u}_1, \dots, \mathbf{u}_D \in \mathbb{R}^D$ associated with eigenvalues $\lambda_1 \geq \dots \geq \lambda_D \in \mathbb{R} \Rightarrow \mathbf{A} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^T = \sum_{i=1}^D \lambda_i \mathbf{u}_i \mathbf{u}_i^T$ where $\mathbf{U} = (\mathbf{u}_1 \dots \mathbf{u}_D)$ is orthogonal and $\mathbf{\Lambda} = \text{diag}(\lambda_1, \dots, \lambda_D)$. In other words, a symmetric real matrix can be diagonalized in terms of its eigenvalues and eigenvectors.
- $\lambda_1 = \max_{\mathbf{x} \neq \mathbf{0}} \frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{\mathbf{x}^T \mathbf{x}} \Leftrightarrow \max \mathbf{x}^T \mathbf{A} \mathbf{x}$ s.t. $\|\mathbf{x}\| = 1$, achieved at $\mathbf{x} = \mathbf{u}_1$.
 $\lambda_2 = \max_{\mathbf{x} \neq \mathbf{0}} \frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{\mathbf{x}^T \mathbf{x}}$ s.t. $\mathbf{x}^T \mathbf{u}_1 = 0$, achieved at $\mathbf{x} = \mathbf{u}_2$.
 Etc.
- Covariance matrix $\mathbf{\Sigma} = \frac{1}{N} \sum_{n=1}^N (\mathbf{x}_n - \boldsymbol{\mu})(\mathbf{x}_n - \boldsymbol{\mu})^T$ positive definite (unless zero variance along some dim.). Mahalanobis distance $(\mathbf{x} - \boldsymbol{\mu})^T \mathbf{\Sigma}^{-1} (\mathbf{x} - \boldsymbol{\mu}) = 1 \Rightarrow$ ellipsoid with axes $= \sqrt{\lambda_1}, \dots, \sqrt{\lambda_D}$ (stdev along PCs).
- $\mathbf{w} \in \mathbb{R}^D$: $\text{var} \{ \mathbf{w}^T \mathbf{x}_1, \dots, \mathbf{w}^T \mathbf{x}_N \} = \mathbf{w}^T \mathbf{\Sigma} \mathbf{w}$. In general for $\mathbf{W}_{D \times L}$: $\text{cov} \{ \mathbf{W}^T \mathbf{x}_1, \dots, \mathbf{W}^T \mathbf{x}_N \} = \mathbf{W}^T \mathbf{\Sigma} \mathbf{W}$.

6.2 Feature selection: forward selection

- Problem: given a sample $\{(\mathbf{x}_n, y_n)\}_{n=1}^N$ with $\mathbf{x}_n \in \mathbb{R}^D$, determine the best subset of the D features such that the number of selected features L is as small as possible and the classification accuracy (using a given classifier, e.g. a linear SVM) is as high as possible.

Using all D features will always give the highest accuracy on the training set but not necessarily on the validation set.

- Useful when some features are unnecessary (e.g. irrelevant for classification or pure noise) or redundant (so we don't need them all).

Useful with e.g. microarray data. Not useful with e.g. image pixels.


- *Best-subset selection*: for each subset of features, train a classifier and evaluate it on a validation set. Pick the subset having highest accuracy and up to L features.

Combinatorial optimization: 2^D possible subsets of D features[?]. Ex: $D = 3$: $\{\emptyset, \{x_1\}, \{x_2\}, \{x_3\}, \{x_1, x_2\}, \{x_1, x_3\}, \{x_2, x_3\}, \{x_1, x_2, x_3\}\}$. Brute-force search only possible for small $D \Rightarrow$ approximate search.


- *Forward selection*: starting with an empty subset \mathcal{F} , sequentially add one new feature at a time. We add the feature $d \in \{1, \dots, D\}$ such that the classifier trained on $\mathcal{F} \cup \{d\}$ has highest classification accuracy in the validation set. Stop when the accuracy doesn't improve much, or when we reach L features.
- *Backward selection*: same thing but start with $\mathcal{F} = \{1, \dots, D\}$ and remove one feature at a time.
- Both are greedy algorithms and are not guaranteed to find an optimal subset.
- Forward selection is preferable when we expect the optimal subset to contain few features (faster than backward selection). It trains $\mathcal{O}(D^2)$ classifiers if we try up to D features.
- These feature selection algorithms are *supervised*: they use the labels y_n when training the classifier. The features selected depend on the classifier we use. There are also unsupervised algorithms.

Ex: forward selection on the Iris dataset ($D = 4$ features).

p. 119

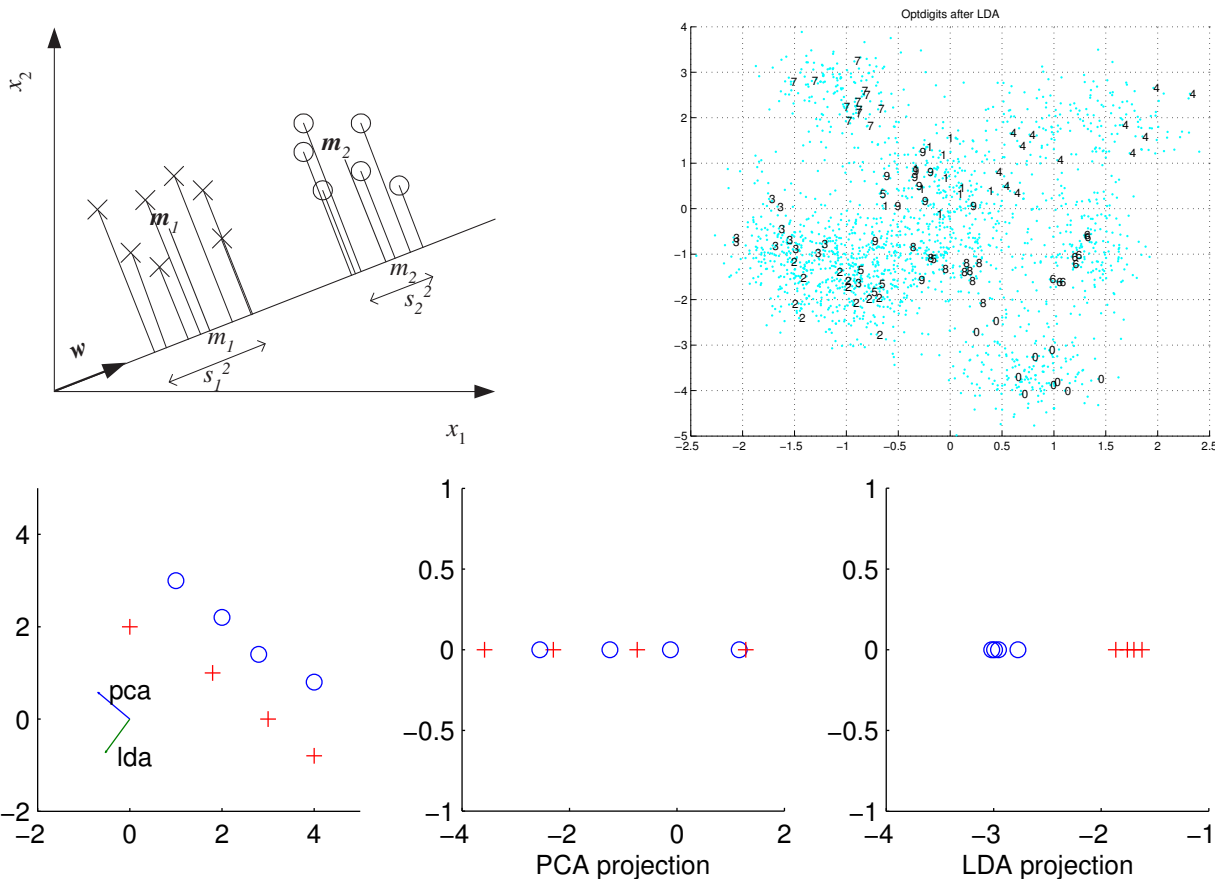


6.3 Dimensionality reduction: principal component analysis (PCA)



- Aims at preserving most of the signal information.
Find a low-dimensional space such that when \mathbf{x} is projected there, information loss is minimized.
- Which direction $\mathbf{w} \in \mathbb{R}^D$ shows most variation? $\max_{\mathbf{w}} \mathbf{w}^T \Sigma \mathbf{w}$ s.t. $\|\mathbf{w}\| = 1 \Rightarrow \mathbf{w} = \mathbf{u}_1$? p. 120
- Unsupervised linear DR method: given $\{\mathbf{x}_n\}_{n=1}^N \subset \mathbb{R}^D$ (with mean zero and covariance matrix Σ of $D \times D$), when reducing dimension to $L < D$, PCA finds:
 - a linear projection mapping $\mathbf{F}: \mathbf{x} \in \mathbb{R}^D \rightarrow \mathbf{W}^T \mathbf{x} \in \mathbb{R}^L$, and
 - a linear reconstruction mapping $\mathbf{f}: \mathbf{z} \in \mathbb{R}^L \rightarrow \mathbf{W} \mathbf{z} \in \mathbb{R}^D$,
 where $\mathbf{W}_{D \times L}$ has orthonormal columns ($\mathbf{W}^T \mathbf{W} = \mathbf{I}$), that are optimal in two equivalent senses:
 - Maximum projected variance: $\max_{\mathbf{W}} \text{tr}(\text{cov}\{\mathbf{W}^T \mathbf{x}_1, \dots, \mathbf{W}^T \mathbf{x}_N\}) \stackrel{?}{=} \text{tr}(\mathbf{W}^T \Sigma \mathbf{W})$.
 - Minimum reconstruction error: $\min_{\mathbf{W}} \frac{1}{N} \sum_{n=1}^N \|\mathbf{x}_n - \mathbf{W} \mathbf{z}_n\|^2 \stackrel{?}{=} -\text{tr}(\mathbf{W}^T \Sigma \mathbf{W}) + \text{constant}$.
- Covariance in the projected space $\text{cov}\{\mathbf{Z}\} = \mathbf{W}^T \Sigma \mathbf{W}$ is diagonal \Rightarrow uncorrelated projections.
- If the mean of the sample is $\boldsymbol{\mu} = \frac{1}{N} \sum_{n=1}^N \mathbf{x}_n \Rightarrow \mathbf{F}(\mathbf{x}) = \mathbf{W}^T (\mathbf{x} - \boldsymbol{\mu})$ and $\mathbf{f}(\mathbf{z}) = \mathbf{W} \mathbf{z} + \boldsymbol{\mu}$.
- How to compute \mathbf{W} , given Σ ? Eigenproblem $\max_{\mathbf{W}} \text{tr}(\mathbf{W}^T \Sigma \mathbf{W})$ s.t. $\mathbf{W}^T \mathbf{W} = \mathbf{I}$ whose solution is given by the spectral theorem. Decompose $\Sigma = \mathbf{U} \Lambda \mathbf{U}^T$ with eigenvectors $\mathbf{U} = (\mathbf{u}_1 \dots \mathbf{u}_D)$ and eigenvalues $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_D)$, sorted decreasingly. Then $\mathbf{W} = \mathbf{U}_{1:L} = (\mathbf{u}_1, \dots, \mathbf{u}_L)$, i.e., the eigenvectors associated with the largest L eigenvalues of the covariance matrix Σ .
- Total variance of the data: $\lambda_1 + \dots + \lambda_D = \text{tr}(\Sigma) = \sigma_1^2 + \dots + \sigma_D^2$.
Variance “explained” by the latent space: $\lambda_1 + \dots + \lambda_L = \text{tr}(\mathbf{W}^T \Sigma \mathbf{W})$.
We can use the proportion $\frac{\lambda_1 + \dots + \lambda_L}{\lambda_1 + \dots + \lambda_D} \in [0, 1]$ of explained variance to determine a good value for L (e.g. 90% of the variance, which usually will be achieved with $L \ll D$).
- In practice with high-dimensional data (e.g. images), a few principal components explain most of the variance if there are correlations among the features.
- Useful as a preprocessing step for classification/regression: $\left\{ \begin{array}{l} \text{reduce the number of features} \\ \text{partly remove noise.} \end{array} \right.$
- Basic disadvantage: it fails with nonlinear manifolds.
- Related linear DR methods:
 - Factor analysis: essentially, a probabilistic version of PCA.
 - Canonical correlation analysis (CCA): projects two sets of features \mathbf{x}, \mathbf{y} onto a common latent space \mathbf{z} .
- Related nonlinear DR methods: autoencoders (based on neural nets), etc.

6.8 Dimensionality reduction: linear discriminant analysis (LDA)




- Aims at preserving most of the signal information *that is useful to discriminate among the classes*. Find a low-dimensional space such that when \mathbf{x} is projected there, classes are well separated.
- Supervised linear DR method: given $\{(\mathbf{x}_n, y_n)\}_{n=1}^N$ where $\mathbf{x}_n \in \mathbb{R}^D$ is a high-dimensional feature vector and $y_n \in \{1, \dots, K\}$ a class label, when reducing dimension to $L < D$, LDA finds a linear projection mapping \mathbf{F} : $\mathbf{x} \in \mathbb{R}^D \rightarrow \mathbf{W}^T \mathbf{x} \in \mathbb{R}^L$ with $\mathbf{W}_{D \times L}$ that is optimal in *maximally separating the classes from each other while maximally compressing each class*.
Unlike PCA, LDA does not find a reconstruction mapping \mathbf{f} : $\mathbf{z} \in \mathbb{R}^L \rightarrow \mathbf{W} \mathbf{z} \in \mathbb{R}^D$. It only finds the projection mapping \mathbf{F} .
- Define:
 - Number of points in class k : N_k . Mean of class k : $\boldsymbol{\mu}_k = \frac{1}{N_k} \sum_{y_n=k} \mathbf{x}_n$.
 - Within-class scatter matrix for class k : $\mathbf{S}_k = \sum_{y_n=k} (\mathbf{x}_n - \boldsymbol{\mu}_k)(\mathbf{x}_n - \boldsymbol{\mu}_k)^T$.
 - Total within-class scatter matrix $\mathbf{S}_W = \sum_{k=1}^K \mathbf{S}_k$.
 - Between-class scatter matrix $\mathbf{S}_B = \sum_{k=1}^K N_k (\boldsymbol{\mu}_k - \boldsymbol{\mu})(\boldsymbol{\mu}_k - \boldsymbol{\mu})^T$ where $\boldsymbol{\mu} = \frac{1}{K} \sum_{k=1}^K \boldsymbol{\mu}_k$.
- In the latent space, the between-class and within-class scatter matrices are $\mathbf{W}^T \mathbf{S}_B \mathbf{W}$ and $\mathbf{W}^T \mathbf{S}_W \mathbf{W}$ (of $L \times L$).
- Fisher discriminant: $\max_{\mathbf{W}} J(\mathbf{W}) = \frac{|\mathbf{W}^T \mathbf{S}_B \mathbf{W}|}{|\mathbf{W}^T \mathbf{S}_W \mathbf{W}|} = \frac{\text{between-class scatter}}{\text{within-class scatter}}$.
- This is an eigenproblem whose solution is $\mathbf{W} = (\mathbf{u}_1, \dots, \mathbf{u}_L) =$ *eigenvectors associated with the largest L eigenvalues of $\mathbf{S}_W^{-1} \mathbf{S}_B$* .
rank(\mathbf{S}_B) $\leq K - 1$ (so we can only use values of L that satisfy $1 \leq L \leq K - 1$).
 \mathbf{S}_W must be invertible (if it is not, apply PCA to the data and eliminate directions with zero variance).

6.7 Dimensionality reduction: multidimensional scaling (MDS)

True distances along earth surface



Estimated positions on a 2D map



- Aims at preserving distances or similarities.
Place N points in a low-dimensional map (of dimension L) such that their distances are well preserved.
- Unsupervised DR method: given the matrix of squared Euclidean distances $d_{nm}^2 = \|\mathbf{x}_n - \mathbf{x}_m\|^2$ between N data points, MDS finds points $\mathbf{z}_1, \dots, \mathbf{z}_N \in \mathbb{R}^L$ that approximate those distances:

$$\min_{\mathbf{z}} \sum_{n,m=1}^N (d_{nm}^2 - \|\mathbf{z}_n - \mathbf{z}_m\|^2)^2.$$
- MDS does not use as training data the actual feature vectors $\mathbf{x}_n \in \mathbb{R}^D$, only the pairwise distances d_{nm} . Hence, it is applicable even when the “distances” are computed between objects that are not represented by features. Ex: perceptual distance between two different colors according to a subject.
- If $d_{nm}^2 = \|\mathbf{x}_n - \mathbf{x}_m\|^2$ where $\mathbf{x}_n \in \mathbb{R}^D$ and $D \geq L$, then MDS is equivalent to PCA on $\{\mathbf{x}_n\}_{n=1}^N$. p. 137
- MDS does not produce a projection or reconstruction mapping, only the actual L -dimensional projections $\mathbf{z}_1, \dots, \mathbf{z}_N \in \mathbb{R}^L$ for the N training points $\mathbf{x}_1, \dots, \mathbf{x}_N \in \mathbb{R}^D$.
- How to learn a projection mapping $\mathbf{F}: \mathbf{x} \in \mathbb{R}^D \rightarrow \mathbf{z} \in \mathbb{R}^L$ with parameters Θ ?

Direct fit: find the projections $\mathbf{z}_1, \dots, \mathbf{z}_N$ by MDS and then solve a nonlinear regression

$$\min_{\Theta} \sum_{n=1}^N (\mathbf{z}_n - \mathbf{F}(\mathbf{x}_n; \Theta))^2$$

Parametric embedding: requires nonlinear optimization

$$\min_{\Theta} \sum_{n,m=1}^N (d_{nm}^2 - \|\mathbf{F}(\mathbf{x}_n; \Theta) - \mathbf{F}(\mathbf{x}_m; \Theta)\|^2)^2.$$

- Generalizations of MDS:
 - *Spectral methods:* Isomap, Locally Linear Embedding, Laplacian eigenmaps...
Require solving an eigenproblem.
Isomap: define d_{nm} = geodesic distances (approximated by shortest paths in a nearest-neighbor graph of the sample).
 - *Nonlinear embeddings:* elastic embedding, t -SNE...
Require solving a nonlinear optimization.

7 Clustering

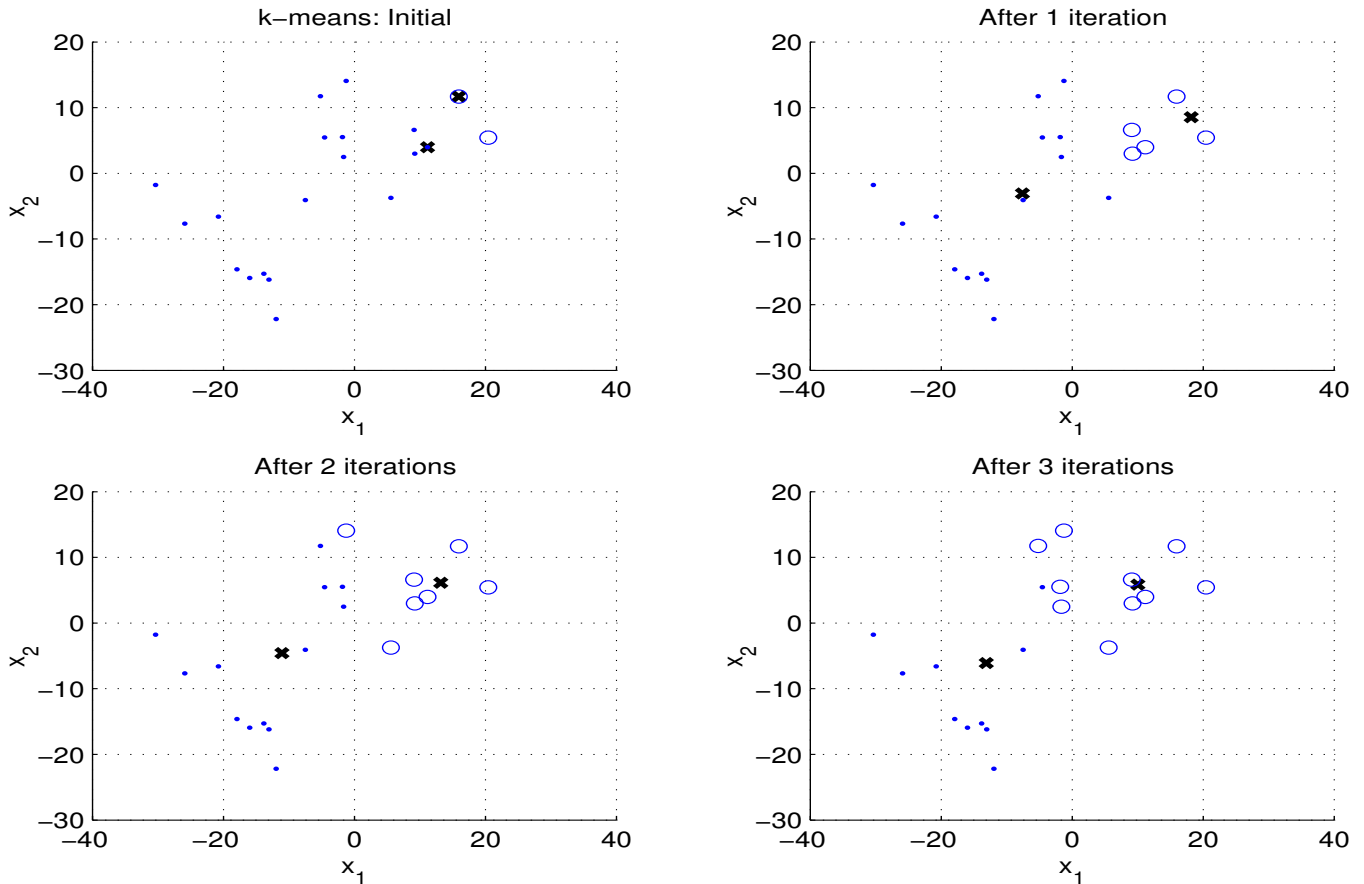
- Unsupervised problem: given a sample $\{\mathbf{x}_n\}_{n=1}^N \subset \mathbb{R}^D$, *partition it into groups such that points within each group are similar and groups are dissimilar* (“hard partition”).
Or, determine soft assignments z_{nk} of point \mathbf{x}_n to cluster k for $n = 1, \dots, N$ and $k = 1, \dots, K$, where $z_{nk} \in [0, 1]$ and $\sum_{k=1}^K z_{nk} = 1$ (“soft partition”).
- Useful to understand structure in the data, or as preprocessing for supervised learning (e.g. train a classifier for each cluster). Clustering does not need labels, which are usually costly to obtain.
- Examples:
 - Customer segmentation: group customers according to demographic and transaction attributes, then provide different strategies for each segment.
 - Image segmentation: partition pixels into meaningful objects in the image.
Represent each pixel \mathbf{x}_n by a feature vector, typically position & intensity (i, j, I) or color (i, j, L^*, u^*, v^*) .
- Basic types:
 - *Centroid-based*: find a prototype (“centroid”) for each cluster.
k-means, *k*-medoids, *k*-modes, etc.
 - Probabilistic models: *mixture densities*, where each component models one cluster.
Gaussian mixtures. Usually trained with the EM algorithm.
 - *Density-based*: find regions of high density of data.
Mean-shift clustering, level-set clustering, etc.
 - *Graph-based* (or *distance-* or *similarity-based*): construct a graph with the data points as vertices and edges weighted by distance values, and partition it.
Hierarchical clustering, connected-components clustering, spectral clustering, etc.
- Ill-defined problem: what does “similar” mean, and how similar should points be to be in the same cluster? Hence, many different definitions of clustering and many different algorithms. This is typical of exploratory tasks, where we want to understand the structure of the data before committing to specific analyses. In practice, one should try different clustering algorithms.

7.2 Centroid-based clustering: *k*-means clustering


- Two uses: *clustering* and *vector quantization*.
- User parameter: number of clusters K . Output: clusters and a centroid $\boldsymbol{\mu}_k \in \mathbb{R}^D$ per cluster.
- Objective function of *centroids* $\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_K$ and *cluster assignments* $\mathbf{Z}_{N \times K}$:

$$E(\{\boldsymbol{\mu}_k\}_{k=1}^K, \mathbf{Z}) = \sum_{n=1}^N \sum_{k=1}^K z_{nk} \|\mathbf{x}_n - \boldsymbol{\mu}_k\|^2 \quad \text{s.t.} \quad \mathbf{Z} \in \{0, 1\}^{NK}, \mathbf{Z}\mathbf{1} = \mathbf{1}.$$

- No closed-form solution. This problem is NP-complete (runtime \geq exponential on problem size). Instead, we do a local search with an iterative algorithm (*alternating optimization*):
 - *Assignment step* (over \mathbf{Z} given $\{\boldsymbol{\mu}_k\}_{k=1}^K$):
for each $n = 1, \dots, N$, **assign \mathbf{x}_n to the cluster with the closest centroid to \mathbf{x}_n .**
 - *Centroid step* (over $\{\boldsymbol{\mu}_k\}_{k=1}^K$ given \mathbf{Z}):
for each $k = 1, \dots, K$, **$\boldsymbol{\mu}_k = \frac{\sum_{n=1}^N z_{nk} \mathbf{x}_n}{\sum_{n=1}^N z_{nk}}$ = mean of the points currently assigned to cluster k .**



- Each iteration (centroid + assignment step) reduces E or leaves it unchanged. After a *finite* number of iterations, no more changes occur and we stop.
- The result depends on the initialization. We usually initialize \mathbf{Z} by a random assignment of points to clusters, run k -means from several such random \mathbf{Z} , and pick the best result.
- *Vector quantization*: compress a continuous space of dimension D , represented by a sample $\mathcal{X} = \{\mathbf{x}_n\}_{n=1}^N$, into a finite set of reference vectors (*codebook*), with minimal distortion.
 - Ex: color quantization. Represent 24-bit color images using only 256 colors with minimal distortion.
 - It can be done with k -means clustering using large K . Better than uniform quantization, because it adapts to density variations in the data and does not require a codebook of size exponential on D .
 - A new point $\mathbf{x} \in \mathbb{R}^D$ is quantized as $\boldsymbol{\mu}_{k^*}$ with $k^* = \arg \min_{k=1, \dots, K} \|\mathbf{x} - \boldsymbol{\mu}_k\|$. This partitions the space into a *Voronoi tessellation*.
 - *Lossy compression*: encode $\mathbf{x} \in \mathbb{R}^D$ (D floats) as an integer $k \in \{1, \dots, K\}$ ($\lceil \log_2 K \rceil$ bits). Also needs to store the codebook.



7.1 Mixture densities and the EM algorithm

- *Mixture density* with K components: $p(\mathbf{x}; \Theta) = \sum_{k=1}^K p(\mathbf{x}|k)p(k) \begin{cases} p(\mathbf{x}|k) & \text{component densities} \\ p(k) = \pi_k & \text{mixture proportions.} \end{cases}$

- Ex: Gaussian mixture: $\mathbf{x}|k \sim \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)$. Mixture parameters: $\Theta = \{\pi_k, \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k\}_{k=1}^K$.

- *Maximum likelihood* estimation of Gaussian mixture parameters: given a sample $\mathcal{X} = \{\mathbf{x}_n\}_{n=1}^N$:

$$\max_{\Theta} \mathcal{L}(\Theta; \mathcal{X}) = \sum_{n=1}^N \log p(\mathbf{x}_n; \Theta) = \sum_{n=1}^N \log \left(\sum_{k=1}^K p(\mathbf{x}|k)p(k) \right).$$

- \mathcal{L} cannot be maximized in closed form over Θ ; it needs an iterative optimization algorithm. Many such algorithms exist (such as gradient descent), but there is a specially convenient one for mixture models (and more generally, for maximum likelihood with missing data).

- *Expectation-Maximization (EM) algorithm*: for Gaussian mixtures:

- *E step*: given the current parameter values Θ , compute the posterior probability of component k given data point \mathbf{x}_n (for each $k = 1, \dots, K$ and $n = 1, \dots, N$):

$$z_{nk} = p(k|\mathbf{x}_n; \Theta) = \frac{p(\mathbf{x}_n|k)p(k)}{p(\mathbf{x}_n; \Theta)} = \frac{\pi_k |2\pi\boldsymbol{\Sigma}_k|^{-1/2} \exp\left(-\frac{1}{2}(\mathbf{x}_n - \boldsymbol{\mu}_k)^T \boldsymbol{\Sigma}_k^{-1}(\mathbf{x}_n - \boldsymbol{\mu}_k)\right)}{\sum_{k'=1}^K \pi_{k'} |2\pi\boldsymbol{\Sigma}_{k'}|^{-1/2} \exp\left(-\frac{1}{2}(\mathbf{x}_n - \boldsymbol{\mu}_{k'})^T \boldsymbol{\Sigma}_{k'}^{-1}(\mathbf{x}_n - \boldsymbol{\mu}_{k'})\right)} \in (0, 1).$$

- *M step*: given the posterior probabilities, estimate the parameters Θ : for $k = 1, \dots, K$:


$$\pi_k = \frac{1}{N} \sum_{n=1}^N z_{nk} \quad \boldsymbol{\mu}_k = \frac{\sum_{n=1}^N z_{nk} \mathbf{x}_n}{\sum_{n=1}^N z_{nk}} \quad \boldsymbol{\Sigma}_k = \frac{\sum_{n=1}^N z_{nk} (\mathbf{x}_n - \boldsymbol{\mu}_k)(\mathbf{x}_n - \boldsymbol{\mu}_k)^T}{\sum_{n=1}^N z_{nk}}.$$

Similar to k -means, where the assignment and centroid steps correspond to the E and M steps. But in EM the assignments are *soft* ($z_{nk} \in [0, 1]$), while in k -means they are *hard* ($z_{nk} \in \{0, 1\}$).

- If we knew which component \mathbf{x}_n came from for each $n = 1, \dots, N$, we'd not need the E step: a single M step that estimates each component's parameters on its set of points would suffice. This was the case in classification (where $\mathbf{x}|C_k$ is Gaussian): we are given (\mathbf{x}_n, y_n) .
- Each EM step increases \mathcal{L} or leaves it unchanged, but it takes an infinite number of iterations to converge. In practice, we stop when the parameters don't change much, or when the number of iterations reaches a limit.
- EM converges to a local optimum that depends on the initial value of Θ . Usually from k -means?
- User parameter: number of clusters K . Output: posterior probabilities $\{p(k|\mathbf{x}_n)\}$ and $\{\pi_k, \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k\}_{k=1}^K$.
- *Parametric clustering*: K clusters, assumed Gaussian.
- The fundamental advantage of Gaussian mixtures over k -means for clustering is that we can model the uncertainty in the assignments (particularly useful for points near cluster boundaries), and the clusters can be elliptical and have different proportions.

	k -means	EM for Gaussian mixtures
assignments z_{nk}	hard	soft, $p(k \mathbf{x}_n)$
probability model?	no	yes
number of iterations	finite	infinite
parameters	centroids $\{\boldsymbol{\mu}_k\}_{k=1}^K$	$\{\pi_k, \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k\}_{k=1}^K$

Density-based clustering: mean-shift clustering



- Define a function $p(\mathbf{x})$ that represents the density of the dataset $\{\mathbf{x}_n\}_{n=1}^N \subset \mathbb{R}^D$, then declare each mode (maximum) of p as a cluster representative and *assign each \mathbf{x}_n to a mode via the mean-shift algorithm*. Most useful with low-dimensional data, e.g. image segmentation (where \mathbf{x}_n = features of n th pixel).
- *Kernel density estimate* with bandwidth σ : a mixture having one component for each data point:

$$p(\mathbf{x}) = \sum_{n=1}^N p(\mathbf{x}|n)p(n) = \frac{1}{N\sigma^D} \sum_{n=1}^N K\left(\left\|\frac{\mathbf{x} - \mathbf{x}_n}{\sigma}\right\|\right) \quad \mathbf{x} \in \mathbb{R}^D.$$

Usually the kernel K is Gaussian: $K\left(\left\|\frac{\mathbf{x} - \mathbf{x}_n}{\sigma}\right\|\right) \propto \exp\left(-\frac{1}{2}\|(\mathbf{x} - \mathbf{x}_n)/\sigma\|^2\right)$.

- *Mean-shift algorithm*: starting from an initial value of \mathbf{x} , it iterates the following expression:

$$\mathbf{x} \leftarrow \sum_{n=1}^N p(n|\mathbf{x})\mathbf{x}_n \quad \text{where} \quad p(n|\mathbf{x}) = \frac{p(\mathbf{x}|n)p(n)}{p(\mathbf{x})} = \frac{\exp\left(-\frac{1}{2}\|(\mathbf{x} - \mathbf{x}_n)/\sigma\|^2\right)}{\sum_{n'=1}^N \exp\left(-\frac{1}{2}\|(\mathbf{x} - \mathbf{x}_{n'})/\sigma\|^2\right)}.$$

$\sum_{n=1}^N p(n|\mathbf{x})\mathbf{x}_n$ can be understood as the weighted average of the N data points using as weights the posterior probabilities $p(n|\mathbf{x})$. The mean-shift algorithm converges to a mode of $p(\mathbf{x})$. Which one it converges to depends on the initialization. By running mean-shift starting at a data point \mathbf{x}_n , we effectively assign \mathbf{x}_n to a mode. We repeat for all points $\mathbf{x}_1, \dots, \mathbf{x}_N$.

- User parameter: σ , which determines the number of clusters ($\sigma \downarrow$: N clusters, $\sigma \uparrow$: 1 cluster). Output: modes and clusters.
- *Nonparametric clustering*: no assumption on the shape of the clusters or their number.

Graph-based clustering: connected-components clustering

- Define a distance $d(\mathbf{x}, \mathbf{y})$ for pairs of instances \mathbf{x}, \mathbf{y} :

– *Minkowski distance* $d(\mathbf{x}, \mathbf{y}) = \left(\sum_{d=1}^D |x_d - y_d|^p\right)^{1/p}$, if $\mathbf{x}, \mathbf{y} \in \mathbb{R}^D$.
 $p = 1$: city-block distance; $p = 2$: Euclidean distance.

– We may use distances without \mathbf{x} and \mathbf{y} being defined by explicit feature vectors.


Number of words that two documents have in common; “edit” distance between strings (Exe. 6).

- Define a *neighborhood graph* with vertices = data points and edges $\mathbf{x}_n \sim \mathbf{x}_m$ if:

– $d(\mathbf{x}_n, \mathbf{x}_m) \leq \epsilon$ (ϵ -ball graph)


– \mathbf{x}_n and \mathbf{x}_m are among the k -nearest-neighbors of each other (k -nearest-neighbor graph).

- *Clusters = connected-components of the graph* (which can be found by depth-first search).
- User parameter: $\epsilon > 0$ or $k \in \mathbb{N}$. The number of clusters depends on that ($\epsilon \downarrow$: N , $\epsilon \uparrow$: 1).
- It can handle complex cluster shapes, but works only with non-overlapping clusters.



7.8 Graph-based clustering: hierarchical clustering

- Generates a nested sequence of clusterings.
- *Agglomerative clustering*: start with N clusters (= singleton points) and repeatedly merge pairs of clusters until there is only one cluster left containing all points.
Divisive clustering: start with one cluster containing all points and repeatedly divide it until we have N clusters (= singleton points).
- We merge the two closest clusters $\mathcal{C}_i, \mathcal{C}_j$ according to a distance $\delta(\mathcal{C}_i, \mathcal{C}_j)$ between clusters:
 - *Single-link clustering*: $\delta(\mathcal{C}_i, \mathcal{C}_j) = \min_{\mathbf{x}_n \in \mathcal{C}_i, \mathbf{x}_m \in \mathcal{C}_j} d(\mathbf{x}_n, \mathbf{x}_m)$.
 Tends to produce elongated clusters (“chaining” effect).
 Equivalent to Kruskal’s algorithm to find a minimum spanning tree of a weighted graph $G = (V, E, w)$ where $V = \{\mathbf{x}_n\}_{n=1}^N$, $E = V \times V$ and $w_{nm} = d(\mathbf{x}_n, \mathbf{x}_m)$.
 - *Complete-link clustering*: $\delta(\mathcal{C}_i, \mathcal{C}_j) = \max_{\mathbf{x}_n \in \mathcal{C}_i, \mathbf{x}_m \in \mathcal{C}_j} d(\mathbf{x}_n, \mathbf{x}_m)$.
 Tends to produce compact clusters.
- Result: *dendrogram*, a binary tree where leaves = instances \mathbf{x}_n and internal nodes = merges. We can obtain a specific clustering from the tree by allowing only distances $d(\mathbf{x}_n, \mathbf{x}_m) \leq \epsilon$ (equivalent to connected-components for single-link clustering).
- User parameter: when to stop merging. Output: dendrogram and clusters.



7.10 How to choose the number of clusters or other user parameters

- Typical user parameters:
 - K : number of clusters (k -means, mixtures).
 - σ, ϵ : “scale” in feature space (mean-shift, connected-components, hierarchical clustering).
 - k : number of neighbors (connected-components).
- How to set K, σ , etc.? Don’t trust “automatic” algorithms that select all user parameters for you!
 - Try several values (and several clustering algorithms) and explore the results:
 - * Plot the objective function (error, log-likelihood, etc.) vs K and look for an “elbow”.
 - * Project to 2D with PCA and inspect the result.
 - In some applications, K may be fixed or known. Color quantization, medical image segmentation.

8 Nonparametric methods

Parametric methods

- Examples:
 - Linear function $f(x; \Theta) = w_1x + w_0$ ($\Theta = \{w_0, w_1\}$).
 - Gaussian distribution $p(x; \Theta) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2}$ ($\Theta = \{\mu, \sigma^2\}$).
 - Gaussian mixture $p(\mathbf{x}; \Theta) = \sum_{k=1}^K p(\mathbf{x}|k)p(k)$ ($\Theta = \{\pi_k, \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k\}_{k=1}^K$).
 - etc.
- Training is the process of choosing the best parameter values Θ from a given dataset.
- The size of the parameters is separate and smaller from the size of the training set N . The model “compresses” the dataset into the parameters Θ , and the complexity of the model doesn’t grow with N . Ex: a linear function $f(x) = w_1x + w_0$ has 2 parameters (w_0, w_1) regardless of the size N of the training set.
- After training, we discard the dataset, and use only the parameters to apply the model to future data. Ex: keep the 2 parameters (w_0, w_1) of the learned linear function and discard the training set $\{(x_n, y_n)\}_{n=1}^N$.
- There are user parameters: number of components K , regularization parameter λ , etc.
- The hypothesis space is restricted: only models of a certain parametric form (linear, etc.).
- Small complexity at test time as a function of N : both memory and time are $\mathcal{O}(1)$.

Nonparametric methods

- There is no training.
- The size of the model is given by the size of the dataset N , hence the complexity of the model can grow with N .
- We keep the training set, which we need to apply the model to future data.
- There are user parameters: number of neighbors k or kernel bandwidth h .
- The hypothesis space is less restricted (fewer assumptions). Often based on a smoothness assumption: similar instances have similar outputs, and the output for an instance is a local function of its neighboring instances. *Essentially, the model behaves like a lookup table that we use to interpolate future instances.*
- Large complexity at test time as a function of N : both memory and time are $\mathcal{O}(N)$, because we have to store all training instances and look through them to make predictions.
- They are particularly preferable to parametric methods with small training sets, where they give better models and are relatively efficient computationally.
- Also called *instance-based* or *memory-based learning algorithms*.
- Examples: kernel density estimate, nearest-neighbor classifier, running mean smoother, etc.

8.2 Nonparametric density estimation

- Given a sample $\mathcal{X} = \{\mathbf{x}_n\}_{n=1}^N$ drawn iid from an unknown density, we want to construct an estimator $p(\mathbf{x})$ of the density.
- Histogram* (consider first $x \in \mathbb{R}$): split the real line into bins $[x_0 + mh, x_0 + (m + 1)h]$ of width h for $m \in \mathbb{Z}$, and count points in each bin:


$$p(x) = \frac{1}{Nh} (\text{number of } x_n \text{ in the same bin as } x) \quad x \in \mathbb{R}.$$

- We need to select the bin width h and the origin x_0 .
- x_0 has a small but annoying effect on the histogram (near bin boundaries).
- h controls the histogram smoothness: spiky if $h \downarrow$ and smooth if $h \uparrow$.
- $p(x)$ is discontinuous at bin boundaries.
- We don't have to retain the training set once we have computed the counts.
- They generalize to D dimensions, but are practically useful only for $D \lesssim 2$
In D dimensions, it requires an exponential number of bins, most of which are empty.
- Kernel density estimate (Parzen windows)*: generalization of histograms to define smooth, multivariate density estimates. Place a kernel $K(\cdot)$ on each data point and sum them:

$$p(\mathbf{x}) = \frac{1}{Nh^D} \sum_{n=1}^N K\left(\left\|\frac{\mathbf{x} - \mathbf{x}_n}{h}\right\|\right) \quad \mathbf{x} \in \mathbb{R}^D \quad \text{“sum of bumps”}.$$

- K must satisfy $K(t) \geq 0 \forall t \in \mathbb{R}$ and $\frac{1}{h^D} \int_{\mathbb{R}^D} K(t) dt = 1$, typically K is Gaussian or uniform.
Gaussian: $K\left(\left\|\frac{\mathbf{x} - \mathbf{x}_n}{h}\right\|\right) = \exp(-\frac{1}{2} \|\mathbf{x} - \mathbf{x}_n\|^2/h^2)$. The uniform kernel gives a histogram without an origin x_0 .
- Only parameter: the bandwidth $h > 0$. The KDE is spiky if $h \downarrow$, smooth if $h \uparrow$.
The KDE is not very sensitive to the choice of K .
- $p(\mathbf{x})$ is continuous and differentiable if K is continuous and differentiable.
- In practice, can take $K(\|\mathbf{x} - \mathbf{x}_n\|/h) = 0$ if $\|\mathbf{x} - \mathbf{x}_n\| > 3h$ to simplify the calculation.
We still need to find the samples \mathbf{x}_n that satisfy $\|\mathbf{x} - \mathbf{x}_n\| \leq 3h$ (neighbors at distance $\leq 3h$).
- Also possible to define a different bandwidth h_n for each data point \mathbf{x}_n (*adaptive KDE*).
- The KDE quality degrades as the dimension D increases (no matter how h is chosen).
Could be improved by using a full covariance Σ_n per point, but it is preferable to use a mixture with $K < N$ components.
- k-nearest-neighbor density estimate*: $p(\mathbf{x}) = \frac{k}{2N} \frac{1}{d_k(\mathbf{x})}$ for $\mathbf{x} \in \mathbb{R}^D$, where $d_k(\mathbf{x}) =$ (Euclidean) distance of \mathbf{x} to its k th nearest sample in \mathcal{X} .

- Like using a KDE with an adaptive bandwidth $h = 2d_k(\mathbf{x})$.
Instead of fixing h and counting how many samples fall in the bin, we fix k and compute the bin size containing k samples.
- Only parameter: the number of nearest neighbors $k \geq 1$.
- $p(\mathbf{x})$ has a discontinuous derivative. It does not integrate to 1 so it is not a pdf.



8.4 Nonparametric classification

- Given $\{(\mathbf{x}_n, y_n)\}_{n=1}^N$ where $\mathbf{x}_n \in \mathbb{R}^D$ is a feature vector and $y_n \in \{1, \dots, K\}$ a class label.
- Estimate the class-conditional densities $p(\mathbf{x}|C_k)$ with a KDE each. The resulting discriminant function can be written (ignoring constant factors) as[?]:

$$g_k(\mathbf{x}) = \sum_{n=1, y_n=k}^N K\left(\left\|\frac{\mathbf{x} - \mathbf{x}_n}{h}\right\|\right) \quad k = 1, \dots, K$$

so each data point \mathbf{x}_n votes only for its class, with a weight given by $K(\cdot)$ (instances closer to \mathbf{x} have bigger weight).

- *k-nearest-neighbor classifier*: assigns \mathbf{x} to the class k having most instances among the k nearest neighbors of \mathbf{x} .

- Most common case: $k = 1$, *nearest-neighbor classifier*.


It defines a Voronoi tessellation in \mathbb{R}^D .

It works with as few as one data point per class!

- k is chosen to be an odd number to minimize ties.

- Simple; good classification accuracy; but slow at test time.

Condensed nearest-neighbor classifier: heuristic way to approximate the nearest-neighbor classifier using a subset of the N data points (to reduce its space and time complexity).



8.8 Nonparametric regression

- Consider a sample $\{(\mathbf{x}_n, \mathbf{y}_n)\}_{n=1}^N$ with $\mathbf{x}_n \in \mathbb{R}^D$ and $\mathbf{y}_n \in \mathbb{R}^E$ drawn iid from an unknown function plus noise $\mathbf{y} = \mathbf{f}(\mathbf{x}) + \epsilon$.

We construct a *smoother* (nonparametric regression estimate) $\mathbf{g}(\mathbf{x})$ of $\mathbf{f}(\mathbf{x})$.

- *Kernel smoother*: weighted average of the labels of instances near \mathbf{x} (local average):

$$\mathbf{g}(\mathbf{x}) = \sum_{n=1}^N \frac{K(\|\mathbf{x} - \mathbf{x}_n\|/h)}{\sum_{n'=1}^N K(\|\mathbf{x} - \mathbf{x}_{n'}\|/h)} \mathbf{y}_n.$$


$\Leftrightarrow \mathbf{g}(\mathbf{x}) = \sum_{n=1}^N w_n(\mathbf{x}) \mathbf{y}_n$ is a weighted average of $\{\mathbf{y}_n\}_{n=1}^N$ with weights satisfying $w_1(\mathbf{x}), \dots, w_N(\mathbf{x}) \geq 0$, $\sum_{n=1}^N w_n(\mathbf{x}) = 1$, and $w_n(\mathbf{x})$ is large if \mathbf{x}_n is near \mathbf{x} and small otherwise

$\mathbf{g}(\mathbf{x})$ can be seen as the conditional mean $E\{\mathbf{y}|\mathbf{x}\}$ of a KDE $p(\mathbf{x}, \mathbf{y})$ constructed on the sample[?].

Also, as with KDEs:

- *Regressogram*: with an origin x_0 and bin width h , and discontinuous at boundaries. eq. (8.24)
- *Running mean smoother*: $K =$ uniform kernel. No origin x_0 , only bin width h , but still discontinuous at boundaries. eq. (8.25)
- *Running median smoother*: like the running mean but using the median instead; more robust to outliers and noise.
- *k-nearest-neighbor smoother*: fixes k instead of h , adapting to the density around \mathbf{x} .

- *Running-line smoother*: we use the data points around \mathbf{x} (as defined by h or k) to estimate a line rather than a constant, hence obtaining a local regression line.



8.9 How to choose the smoothing parameter

- *Smoothing parameter*: number of neighbors k or kernel bandwidth h .
- It controls the complexity of the model:
 - *Too small: low bias, high variance.*
The model is too rough (single instances have a large effect on the predictions).
 - *Too large: high bias, low variance.*
The model is too smooth (single instances have a small effect on the predictions).
- How to choose it?
 - Regression or classification: cross-validation.
 - Density estimation: by trial-and-error.
Some heuristic rules exist that can be used as ballpark estimates.


9 Decision trees

- Nonparametric methods applicable to classification and regression given a sample $\{(\mathbf{x}_n, y_n)\}_{n=1}^N$.
- It can use continuous and discrete features.
- Efficient at test time. May not even use all the features in the training set.
- The tree can be transformed into a set of IF-THEN rules that are *interpretable* for people.
- Widely used in practice, sometimes preferred over more accurate but less interpretable models.
- They define class regions as a sequence of recursive splits.
- The tree consists of:
 - *Internal decision nodes*, each having ≥ 2 children. Decision node m selects one of its children based on a test (a *split*) applied to the input \mathbf{x} .
 - * Continuous feature x_d : “ $x_d > s_m$ ” for some $s_m \in \mathbb{R}$.
 - * Discrete feature x_d : n -way split for the n possible values of x_d .
 - *Leaves*, each containing a value (class label or output value y). Instances \mathbf{x}_n falling in the same leaf should have identical (or similar) output values y_n .
 - * Classification: class label $y \in \{1, \dots, K\}$ (or proportion of each class p_1, \dots, p_K).
 - * Regression: numeric value $y \in \mathbb{R}$ (average of the outputs for the leaf’s instances).

The predicted output for an instance \mathbf{x} is obtained by following a path from the root to a leaf. In the best case (balanced tree) for binary trees, the path has length $\log_2 L$ if there are L leaves.

- Decision trees vs nonparametric methods:
 - Each leaf defines a “bin” (box) in input space. Different bins may have different sizes and contain different numbers of instances.
Unlike in nonparametric methods using a fixed bandwidth h or a fixed number of neighbors k .
 - The bin for an instance \mathbf{x} can be found very quickly (a few feature comparisons).
No need to scan the whole training set.
 - We can discard the training set once the tree is constructed, and keep only its nodes (and associated thresholds and output values).
Then, why do we call it nonparametric?
 - The size of the tree can still be $\mathcal{O}(N)$ if each leaf represents one instance. We stop the tree from growing before this happens, so each leaf represents several instances.
We can also grow the full tree and then prune it back to a desired size.

9.2 Univariate trees



- The test at node m compares one feature with a threshold: “ $x_d > s_m$ ” for some $d \in \{1, \dots, D\}$ and $s_m \in \mathbb{R}$. This defines a *binary split* into two regions: $\{\mathbf{x} \in \mathbb{R}^D: x_d \leq s_m\}$ and $\{\mathbf{x} \in \mathbb{R}^D: x_d > s_m\}$. The overall tree defines box-shaped, axis-aligned regions in input space.

Simplest and most often used. More complex tests exist, e.g. “ $\mathbf{w}_m^T \mathbf{x} > s_m$ ”, which define oblique regions (*multivariate trees*).

- With discrete features, the number of children equals the number n_d of values the feature can take, and the test selects the child corresponding to the value of x_d (*n-way split*).

Ex: $x_d \in \{\text{red, green, blue}\} \Rightarrow n_d = 3$ children.

- *Tree induction* is learning the tree from a training set $\{(\mathbf{x}_n, y_n)\}_{n=1}^N$, i.e., determining its nodes:

- For each internal node, its test (feature $d \in \{1, \dots, D\}$ and threshold s_m).
- For each leaf, its output value y .

- For a given sample, many (sufficiently big) trees exist that code it with zero error. We want to find the smallest such tree. This is NP-complete so an approximate, greedy algorithm is used:

- Starting at the root with the entire training set, select a best split according to a *purity criterion*, and associate each child with the subset of instances that fall in it.
- Continue splitting each child (with its subset of the training set) recursively until each child is pure (hence a leaf) and no more splits are necessary.


Ex. algorithms: CART, ID3, C4.5.

Fig. 9.3
pseudocode

Classification trees

- *Purity criterion*: a split is pure if each child it produces contains instances of the same class. Consider a node and all the training set instances that reach it, and call p_k the proportion of instances of class k , for $k = 1, \dots, K$ (so $p_k \geq 0$ and $\sum_{k=1}^K p_k = 1$). We can measure impurity as the *entropy* of $\mathbf{p} = (p_1, \dots, p_K)$: $\phi(\mathbf{p}) = -\sum_{k=1}^K p_k \log_2 p_k$ (where $0 \log_2 0 \equiv 0$). This is maximum if $p_1 = \dots = p_K = \frac{1}{K}$, and minimum (“pure”) if one $p_k = 1$ and the rest are 0.

Other measures possible: Gini index $\phi(\mathbf{p}) = \sum_{i \neq j} p_i p_j$, misclassification error $\phi(\mathbf{p}) = 1 - \max(p_1, \dots, p_K)$.



- If a node is pure, i.e., all its instances are of the same class k , there is no need to split it. *It becomes a leaf with output value k .*

We can also store the proportions $\mathbf{p} = (p_1, \dots, p_K)$ in the node (e.g. to compute risks).

- If a node m is not pure, we split it. We evaluate all possible features $d = 1, \dots, D$ and all possible split thresholds s_m for each feature and pick the one with minimum impurity ϕ .
 - If the number of instances that reach node m is N_m , there are $N_m - 1$ possible thresholds (half way between each value of x_d , assuming we have sorted them), and the best threshold must be between values of different classes.
 - For discrete features, there is no threshold but an n -way split.

Regression trees

- *Purity criterion*: the squared error $E(g) = \sum_{n \in \text{node}} (y_n - g)^2$ where $g \in \mathbb{R}$, which is minimal when g is the mean of the y_n values[?], and then $E(g)$ is the variance of the y_n values.

If there is much noise or outliers, it is preferable to set g to the median of the y_n values.

- We consider a node to be pure if $E \leq \theta$ for a user threshold $\theta > 0$. In that case, we do not split it. *It becomes a leaf with output value g .*
- If a node m is not pure, we split it. We evaluate all possible features $d = 1, \dots, D$ and all possible split thresholds s_m for each feature and pick the one with minimum impurity (= sum of the variances E of each of the children), as in the classification case.
- Rather than assigning a constant output value to a leaf, we can assign it a regression function (e.g. linear), as in a running-mean smoother.

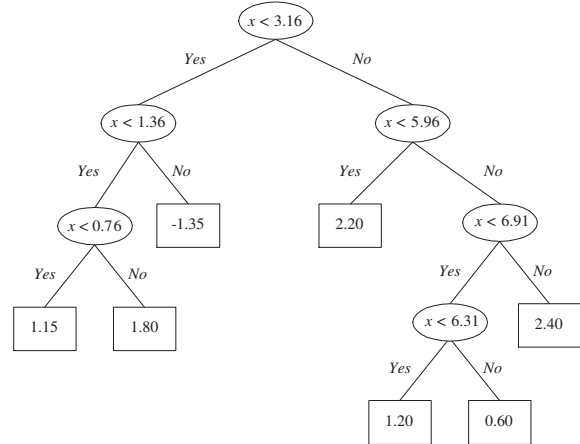
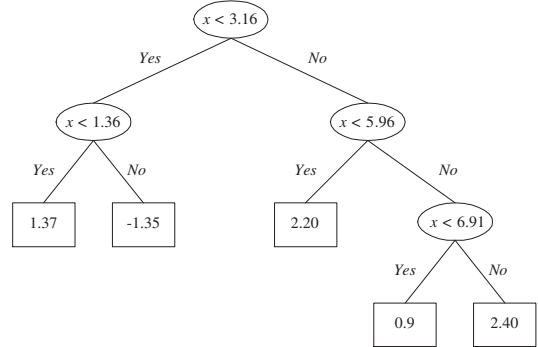

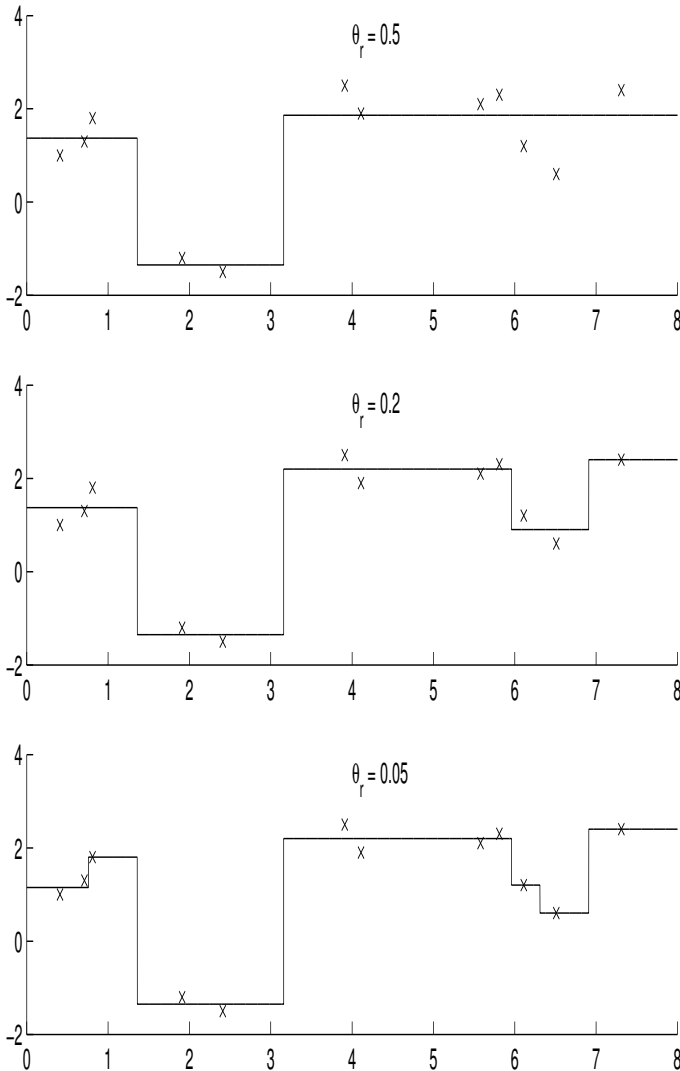
Early stopping and pruning

- Growing the tree until each leaf is pure will produce a large tree with high variance (sensitive to the training sample) that will overfit when there is noise.
How to learn smaller trees that generalize better to unseen data?
- *Early stopping*: we stop splitting if the impurity is below a user threshold $\theta > 0$.
 $\theta \downarrow$ low bias, high variance, large tree; $\theta \uparrow$ high bias, low variance, small tree.
- *Pruning*: we grow the tree in full until all leaves are pure and the training error is zero. Then, we find subtrees that cause overfitting and prune them.

Keep aside a subset from the training set (“pruning set”). For each possible subtree, try replacing it with a leaf node labeled with the training instances covered by the subtree. If the leaf node performs no worse than the subtree on the pruning set, we prune the subtree and keep the leaf node because the additional complexity of the subtree is not justified; otherwise, we keep the subtree.

- Pruning is slower than early stopping but it usually leads to trees that generalize better.


An intuitive reason why is as follows. In the greedy algorithm used to grow the tree, once we make a decision at a given node (to select a split) we never backtrack and try a different, maybe better, possibility.



9.4 Rule extraction from trees

- A decision tree does its own feature extraction: the final tree may not use all the D features.
- Features closer to the root are more important globally.
- Path root \rightsquigarrow leaf = conjunction of tests. This and the leaf's output value give a rule.
- The set of extracted rules allows us to extract knowledge from the dataset.

- R1: IF (age > 38.5) AND (years-in-job > 2.5) THEN $y = 0.8$
- R2: IF (age > 38.5) AND (years-in-job \leq 2.5) THEN $y = 0.6$
- R3: IF (age \leq 38.5) AND (job-type = 'A') THEN $y = 0.4$
- R4: IF (age \leq 38.5) AND (job-type = 'B') THEN $y = 0.3$
- R5: IF (age \leq 38.5) AND (job-type = 'C') THEN $y = 0.2$



10 Linear discrimination

- Consider learning a classifier given a sample $\{(\mathbf{x}_n, y_n)\}_{n=1}^N$ where $\mathbf{x}_n \in \mathbb{R}^D$ and $y_n \in \{1, \dots, K\}$.
- Classification works as follows:
 - Training time: **learn a set of discriminant functions** $\{g_k(\mathbf{x})\}_{k=1}^K$.
 - Test time: given a new instance \mathbf{x} , **choose** C_k if $k = \arg \max_{i=1, \dots, K} \{g_i(\mathbf{x})\}$.
- Two approaches to learning discriminant functions:
 - *Generative approach*: we learn $p(\mathbf{x}|C_k)$ and $p(C_k)$ for each class from the training data, and then use $g_k(\mathbf{x}) = p(C_k|\mathbf{x}) \propto p(\mathbf{x}|C_k)p(C_k)$ (from Bayes' rule) to predict a class. Hence, besides learning the class boundaries (where $p(C_i|\mathbf{x}) = p(C_j|\mathbf{x})$ for $i \neq j$), we model also the density of each class.
Previous chapters, using parametric and nonparametric methods for $p(\mathbf{x}|C_k)$.
 - *Discriminative approach*: we learn only the class boundaries, through discriminant functions $g_k(\mathbf{x})$, $k = 1, \dots, K$. We don't learn the class densities, and $g_k(\mathbf{x})$ need not be modeled using probabilities. Hence, it requires assumptions only about the class boundaries but not about their densities. It solves a simpler problem.
This and future chapters, using linear and nonlinear discriminant functions.
- We define a parametric model $g_k(\mathbf{x}; \Theta_k)$ for each class discriminant. In *linear discrimination*, this model is linear: $g_k(\mathbf{x}; \mathbf{w}_k) = \sum_{d=1}^D w_{kd}x_d + w_{k0} = \mathbf{w}_k^T \mathbf{x}$ (assuming an extra constant feature $x_0 = 1$). We assume that, in the training set, instances from any one class are linearly separable from instances of the other classes.
- Learning means finding parameter values $\{\Theta_k\}_{k=1}^K$ that optimize the quality of the separation.
In the parametric generative approach, learning means finding parameter values Θ_k for each class $p(\mathbf{x}|C_k; \Theta_k)$ that maximize the likelihood, separately for each class.
- Linear discriminants are simpler than nonlinear ones:
 - Faster to train.
 - Low space and time complexity at test time: $\mathcal{O}(D)$ per class. To store \mathbf{w}_k and multiply times it.
 - Simple to interpret: the output is a weighted sum of the features x_d .
 - * Magnitude of w_{kd} : importance of x_d in the decision.
 - * Sign of w_{kd} : whether the effect of x_d is positive or negative.
 - Accurate enough in many applications.
- In practice, try linear discrimination first, before trying nonlinear discrimination.


10.2 Generalizing the linear model

- When a linear model (linear in terms of \mathbf{x}) is not flexible enough, we can:
 - Add higher-order terms.
Ex: if $\mathbf{x} = (x_1, x_2)$, we can define new variables $z_1 = 1$, $z_2 = x_1$, $z_3 = x_2$, $z_4 = x_1^2$, $z_5 = x_2^2$, $z_6 = x_1x_2$ and take $\mathbf{z} = (z_1, z_2, z_3, z_4, z_5, z_6)$ as input. A linear function in the 6D space of \mathbf{z} , $\mathbf{w}^T \mathbf{z}$, corresponds to a nonlinear function in the 2D space of \mathbf{x} .
 - Write the discriminant as $g_i(\mathbf{x}) = \sum_{k=1}^K w_k \phi_{ik}(\mathbf{x})$ where $\phi_{ik}(\mathbf{x})$ are *basis functions*.
Ex. of $\phi(\mathbf{x})$: $x_1^{\alpha_1} \dots x_D^{\alpha_D}$, $\exp(-((\mathbf{x} - \boldsymbol{\mu})/\sigma)^2)$, $\sin(\mathbf{u}^T \mathbf{x})$, etc., for suitable, *fixed* values of $\boldsymbol{\alpha}$, $\boldsymbol{\mu}$, σ , \mathbf{u} , etc.

Radial basis function (RBF) networks and kernel support vector machines (SVMs) do this.

ch. 12–13

10.3 Geometry of the linear discriminant



Two classes

- One discriminant function is sufficient: $g_1(\mathbf{x}) - g_2(\mathbf{x}) \stackrel{?}{=} \mathbf{w}^T \mathbf{x} + w_0 = g(\mathbf{x})$.
Testing: choose C_1 if $g(\mathbf{x}) > 0$ and C_2 if $g(\mathbf{x}) < 0$.
- This defines a *hyperplane* where \mathbf{w} is the *weight* vector and w_0 the *threshold* (or *bias*). It divides the input space \mathbb{R}^D into two half-spaces, the decision regions \mathcal{R}_1 for C_1 (positive side) and \mathcal{R}_2 for C_2 (negative side). The hyperplane itself is the *boundary* or *decision surface*.

- The origin $\mathbf{x} = \mathbf{0}$ is on the

{	positive side	if $w_0 > 0$
	boundary	if $w_0 = 0$
	negative side	if $w_0 < 0$.

- \mathbf{w} is orthogonal to the hyperplane. Pf. Pick \mathbf{x}, \mathbf{y} on the hyperplane.
- The signed distance from $\mathbf{x} \in \mathbb{R}^D$ to the hyperplane is $r = g(\mathbf{x}) / \|\mathbf{w}\|$.
Pf. Write $\mathbf{x} = \mathbf{x}_p + r \frac{\mathbf{w}}{\|\mathbf{w}\|}$ where $\mathbf{x}_p =$ orthogonal projection of \mathbf{x} on the hyperplane and compute $g(\mathbf{x})$.
The signed distance of the origin to the hyperplane is $r_0 = w_0 / \|\mathbf{w}\|$.
- So \mathbf{w} determines the orientation of the hyperplane and w_0 its location wrt the origin.

$K > 2$ classes There are two common ways to define the decision for a point $\mathbf{x} \in \mathbb{R}^D$:

- *One-vs-all* (linear separability): we use K discriminant functions $g_k(\mathbf{x}) = \mathbf{w}_k^T \mathbf{x} + w_{k0}$.
 - *Training*: discriminant function g_k is trained to classify the training points of class C_k vs the training points of all other classes.
 - *Testing*: choose C_k if $k = \arg \max_{i=1, \dots, K} g_i(\mathbf{x})$, i.e., pick the class having the larger discriminant. Ideal case: $g_k(\mathbf{x}) > 0$ and $g_i(\mathbf{x}) < 0 \forall i \neq k$.
- *One-vs-one* (pairwise linear separability): we use $K(K-1)/2$ discriminants $g_{ij}(\mathbf{x}) = \mathbf{w}_{ij}^T \mathbf{x} + w_{ij0}$.
 - *Training*: discriminant function g_{ij} (for $i \neq j$) is trained to classify the training points of class C_i vs the training points of class C_j (training points from other classes are not used).
 - *Testing*: choose C_k if $k = \arg \max_{i=1, \dots, K} \sum_{j \neq i}^K g_{ij}(\mathbf{x})$, i.e., pick the class having the larger summed discriminant. Ideal case: $g_{k-j}(\mathbf{x}) > 0 \forall j$ and $g_{ij}(\mathbf{x}) < 0 \forall i \neq k, \forall j$.
Also possible: for each class k , count the number of times $g_{kj}(\mathbf{x}) > 0$ for $j \neq k$, and pick the class with most votes.
- Both divide the input space \mathbb{R}^D into K convex decision regions $\mathcal{R}_1, \dots, \mathcal{R}_K$.

10.6 Minimizing functions by (stochastic) gradient descent

- When the minimization problem $\min_{\mathbf{w} \in \mathbb{R}^D} E(\mathbf{w})$ cannot be solved in closed form (solution $\mathbf{w}^* = \text{some formula}$), we use iterative methods ($\mathbf{w}^{(0)} \rightarrow \mathbf{w}^{(1)} \rightarrow \dots \rightarrow \mathbf{w}^{(\infty)} = \mathbf{w}^*$). Many such methods exist (Newton's method, conjugate gradients...). One of the simplest ones, reasonably effective in many cases (although sometimes very slow), is gradient descent.
- *Gradient descent (GD)*: repeatedly iterate $\mathbf{w} \leftarrow \mathbf{w} + \Delta \mathbf{w}$ with an *update* $\Delta \mathbf{w} = -\eta \nabla E(\mathbf{w})$. Elementwise for $d = 1, \dots, D$: $w_d \leftarrow w_d + \Delta w_d$ where $\Delta w_d = -\eta \frac{\partial E}{\partial w_d}$.
 - *Gradient* (vector of partial derivatives): $\nabla E(\mathbf{w}) = \left(\frac{\partial E}{\partial w_1}, \dots, \frac{\partial E}{\partial w_D} \right)^T \in \mathbb{R}^D$.
 - *Step size* or *learning rate*: $\eta > 0$. Neither too small (slow convergence) nor too large (oscillations or divergence).
 - *Initial weight vector* $\mathbf{w}^{(0)}$: usually small random numbers. Ex: $w_d \sim \text{uniform}[-0.01, 0.01]$.
 - *Stop iterating*:
 - * When $\nabla E(\mathbf{w}) \approx \mathbf{0}$ (since $\nabla E(\mathbf{w}^*) = \mathbf{0}$ at a minimizer \mathbf{w}^* of E). But this overfits.
 - * When the error on a validation set starts increasing (*early stopping*), which will happen before the training error is minimized. The model \mathbf{w} generalizes better to unseen data.
 - It will find a *local minimizer* (not necessarily *global*).
- *Stochastic gradient descent (SGD)*: applicable when $E(\mathbf{w}) = \sum_{n=1}^N e(\mathbf{w}; \mathbf{x}_n)$, i.e., the total error is the sum of the error at each data point. We update \mathbf{w} as soon as we process \mathbf{x}_n : repeatedly iterate $\mathbf{w} \leftarrow \mathbf{w} + \Delta \mathbf{w}$ with an *update* $\Delta \mathbf{w} = -\eta \nabla e(\mathbf{w}; \mathbf{x}_n)$.
 - *Epoch* = one pass over the whole dataset $\{\mathbf{x}_1, \dots, \mathbf{x}_N\}$. It corresponds to one “noisy” iteration of gradient descent: it need not always decrease $E(\mathbf{w})$ because it doesn't use the correct gradient $\nabla E(\mathbf{w}) = \sum_{n=1}^N \nabla e(\mathbf{w}; \mathbf{x}_n)$.
 - Much faster than (batch) gradient descent to get an approximate solution for \mathbf{w} if N is large (so there is redundancy in the dataset), or if data points come one at a time and we don't store them (*online learning*). However, very slow convergence thereafter.
 - The step size has to decrease slowly over epochs for convergence to occur.
One typically takes $\eta^{(t)} = \frac{\alpha}{\beta + t}$ at epoch t for suitable $\alpha, \beta > 0$.
 - *Shuffling*: in each epoch, process the N points in a random order. Better than a fixed order.
 - *Minibatches*: computing the update $\Delta \mathbf{w} = -\eta \sum_{n \in \mathcal{B}} \nabla e(\mathbf{w}; \mathbf{x}_n)$ based on subsets \mathcal{B} of $1 < |\mathcal{B}| < N$ points works better than with $|\mathcal{B}| = 1$ (pure online learning) or $|\mathcal{B}| = N$ (pure batch learning). In practice $|\mathcal{B}| = 10$ to 1000 points typically.
- Ex: $E(\mathbf{w}) = \frac{1}{2} \sum_{n=1}^N (y_n - \mathbf{w}^T \mathbf{x}_n)^2$ (least-squares error for linear regression). The updates:
GD: $\Delta \mathbf{w} = \eta \sum_{n=1}^N (y_n - \mathbf{w}^T \mathbf{x}_n) \mathbf{x}_n$ SGD: $\Delta \mathbf{w} = \eta (y_n - \mathbf{w}^T \mathbf{x}_n) \mathbf{x}_n$
have the form (ignoring η): “Error \times Input” where Error = DesiredOutput – ActualOutput. This has the following effect:
 - If Error = 0, don't update \mathbf{w} , otherwise update \mathbf{w} proportionally to Error.
 - The update increases \mathbf{w} if Error and Input have the same sign, else it decreases \mathbf{w} .

This results in correcting \mathbf{w} so the error becomes smaller.

10.7 Logistic regression


👁️ a classification method, not a regression method!

Two classes

- The *logistic regression* or *logistic classifier*: given $\mathbf{x} \in \mathbb{R}^D$, calculate $\theta = \sigma(\mathbf{w}^T \mathbf{x} + w_0) \in (0, 1)$ and choose C_1 if $\theta > \frac{1}{2}$, or equivalently calculate $\mathbf{w}^T \mathbf{x} + w_0 \in (-\infty, \infty)$ and choose C_1 if it is positive. It transforms a discriminant value $g(\mathbf{x}) = \mathbf{w}^T \mathbf{x} + w_0$ into a posterior probability $p(C_1|\mathbf{x}) = \theta = \sigma(\mathbf{w}^T \mathbf{x} + w_0) = \frac{1}{1 + \exp(-(\mathbf{w}^T \mathbf{x} + w_0))}$.

Logistic (sigmoid) function

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



- It defines a parametric estimator for $p(C_1|\mathbf{x})$ directly, without having a model for $p(\mathbf{x}|C_1)$, $p(C_1)$ and $p(\mathbf{x}|C_2)$, $p(C_2)$. We learn its parameters (\mathbf{w}, w_0) from a training set $\{(\mathbf{x}_n, y_n)\}_{n=1}^N$.
- Another way to derive the logistic classifier: let us model the log-ratio of the class-conditional densities as a linear function: $\log \frac{p(\mathbf{x}|C_1)}{p(\mathbf{x}|C_2)} = \mathbf{w}^T \mathbf{x} + w_0$. This implies? $p(C_1|\mathbf{x}) = \sigma(\mathbf{w}^T \mathbf{x} + w_0)$ (where $w_0 = w_0^0 + \log \frac{p(C_1)}{p(C_2)}$). Note that, unlike in generative models, we don't have a model of the class-conditional densities $p(\mathbf{x}|C_k)$ themselves.


A linear log-ratio does hold for Gaussian classes with shared covariance ($\mathbf{x}|C_k \sim \mathcal{N}(\boldsymbol{\mu}_k, \boldsymbol{\Sigma})$) 🌀, but can also hold in other cases.

- When the data are Gaussian distributed, logistic regression has a comparable error rate to a parametric Gaussian linear discriminant. Logistic regression can still be applied when the class-conditional densities are non-Gaussian, as long as they are linearly separable.
- Geometry of the logistic classifier: like the geometry of the linear discriminant (\mathbf{w} determines the orientation of the decision boundary and w_0 its location wrt the origin), and the magnitude of \mathbf{w} and w_0 determines how steep the logistic function becomes.
- *Logit function* or *log odds* of θ : $\text{logit}(\theta) = \log\left(\frac{\theta}{1-\theta}\right) \in (-\infty, \infty)$ for $\theta \in (0, 1)$. It is the inverse of the logistic function.


$K > 2$ classes

- We use a *softmax function* to model $p(C_k|\mathbf{x}) = \frac{\exp(\mathbf{w}_k^T \mathbf{x} + w_{k0})}{\sum_{j=1}^K \exp(\mathbf{w}_j^T \mathbf{x} + w_{j0})}$ for $k = 1, \dots, K$.
 - It satisfies $p(C_k|\mathbf{x}) \in (0, 1)$ and $\sum_{k=1}^K p(C_k|\mathbf{x}) = 1$. For $K = 2$, softmax = logistic function.
 - “Soft” $\max(\cdot)$: if $\mathbf{w}_k^T \mathbf{x} + w_{k0} \gg \mathbf{w}_j^T \mathbf{x} + w_{j0} \forall j \neq k$ then $p(C_k|\mathbf{x}) \approx 1$, $p(C_j|\mathbf{x}) \approx 0 \forall j \neq k$. It is differentiable (unlike the $\max(\cdot)$ function).
In this sense, the logistic function is a soft step function.
- *Testing*: given $\mathbf{x} \in \mathbb{R}^D$, calculate $(\theta_1, \dots, \theta_K) = \text{softmax}(\mathbf{w}_1^T \mathbf{x} + w_{10}, \dots, \mathbf{w}_K^T \mathbf{x} + w_{K0})$ and choose C_k if $k = \arg \max_{i=1, \dots, K} \theta_i$.


Decision boundaries for $K = 3$ classes



Linear discriminants



Posterior probabilities (softmax of linear discriminants)



10.8 Learning logistic regression

Given a sample $\{(\mathbf{x}_n, y_n)\}_{n=1}^N$, we want to learn an estimator:

- $K = 2$: $p(C_1|\mathbf{x}) = \sigma(\mathbf{w}^T \mathbf{x} + w_0) = \frac{1}{1 + \exp(-(\mathbf{w}^T \mathbf{x} + w_0))}$ (logistic) with $y_n \in \{0, 1\}$.
- $K > 2$: $p(C_k|\mathbf{x}) = \frac{\exp(\mathbf{w}_k^T \mathbf{x} + w_{k0})}{\sum_{j=1}^K \exp(\mathbf{w}_j^T \mathbf{x} + w_{j0})}$ (softmax) for each class with $y_n \in \{1, \dots, K\}$.

We can learn $\{\mathbf{w}, w_0\}$ or $\{\mathbf{w}_k, w_{k0}\}_{k=1}^K$ by maximum likelihood, or by least-squares regression.

10.8.1 By maximum likelihood


Two classes: $y_n \in \{0, 1\}$

- We model $y_n|\mathbf{x}_n$ as a Bernoulli distribution with parameter $\theta_n = p(C_1|\mathbf{x}_n) = \sigma(\mathbf{w}^T \mathbf{x}_n + w_0)$:

$$\begin{aligned} \max_{\mathbf{w}, w_0} \mathcal{L}(\mathbf{w}, w_0; \{(\mathbf{x}_n, y_n)\}_{n=1}^N) &= \sum_{n=1}^N \log p(y_n|\mathbf{x}_n; \mathbf{w}, w_0) = \sum_{n=1}^N \log (\theta_n^{y_n} (1 - \theta_n)^{1-y_n}) \Leftrightarrow \text{change of sign} \\ \min_{\mathbf{w}, w_0} E(\mathbf{w}, w_0; \{(\mathbf{x}_n, y_n)\}_{n=1}^N) &= - \sum_{n=1}^N (y_n \log \theta_n + (1 - y_n) \log (1 - \theta_n)) = \text{cross-entropy} \\ &= - \sum_{n \in C_1} \log (1 - \theta_n) - \sum_{n \in C_2} \log \theta_n. \end{aligned}$$

- No closed-form solution. We apply gradient descent and obtain:

$$\Delta \mathbf{w} = -\eta \frac{\partial E}{\partial \mathbf{w}} = \eta \sum_{n=1}^N (y_n - \theta_n) \mathbf{x}_n, \quad \Delta w_0 = -\eta \frac{\partial E}{\partial w_0} = \eta \sum_{n=1}^N (y_n - \theta_n)$$

where $\theta_n = \sigma(\mathbf{w}^T \mathbf{x}_n + w_0)$.  Pf. Use the chain rule and $\frac{d\sigma(t)}{dt} = \sigma(t)(1 - \sigma(t))$.

If the classes are linearly separable, $\|\mathbf{w}\| \rightarrow \infty$ and $\theta_n \rightarrow y_n \in \{0, 1\}$ as training proceeds.

Instead, we stop early, when the number of misclassifications is zero.

Newton's method (suitably modified) is much more effective than gradient descent for this problem.




Fig. 10.6 pseudocode

$K > 2$ classes: $y_n \in \{1, \dots, K\}$

- As before, but we model $y_n|\mathbf{x}_n$ as a multinomial distribution with parameter $\theta_{nk} = p(C_k|\mathbf{x}_n) = \frac{\exp(\mathbf{w}_k^T \mathbf{x}_n + w_{k0})}{\sum_{j=1}^K \exp(\mathbf{w}_j^T \mathbf{x}_n + w_{j0})}$. The error function is again the *cross-entropy* (maximum likelihood with a change of sign), where $\mathbf{y}_n \in \{0, 1\}^K$ uses a 1-of- K encoding: p. 254

$$\min_{\{\mathbf{w}_k, w_{k0}\}_{k=1}^K} E(\{\mathbf{w}_k, w_{k0}\}_{k=1}^K; \{(\mathbf{x}_n, \mathbf{y}_n)\}_{n=1}^N) = - \sum_{n=1}^N \sum_{k=1}^K y_{nk} \log \theta_{nk}.$$

Multinomial distribution:

- A die with K faces, each with probability $\theta_k \in [0, 1]$ for $k = 1, \dots, K$ with $\sum_{k=1}^K \theta_k = 1$.
- $p(\mathbf{y}; \boldsymbol{\theta}) = \theta_1^{y_1} \dots \theta_K^{y_K} = \begin{cases} \theta_1, & y_1 = 1 \\ \dots & \\ \theta_K, & y_K = 1 \end{cases}$ where $\mathbf{y} \in \{0, 1\}^K$ has exactly one $y_k = 1$ and the rest are 0.
- For $K = 2$ it becomes the Bernoulli distribution.

p. 600

10.8.2 By least-squares regression

Two classes: $y_n \in \{0, 1\}$

- We define a least-squares regression problem $\min_{\Theta} E(\Theta) = \frac{1}{2} \sum_{n=1}^N (y_n - f(\mathbf{x}_n; \Theta))^2$ where:
 - The labels $\{y_n\}$ are considered as real values (which happen to be either 0 or 1), and can be seen as the desired posterior probabilities for the points $\{\mathbf{x}_n\}$.
 - The parametric function to be learned is $f(\mathbf{x}; \mathbf{w}, w_0) = \sigma(\mathbf{w}^T \mathbf{x} + w_0) = \frac{1}{1 + \exp(-(\mathbf{w}^T \mathbf{x} + w_0))}$.
 - Hence, learning the classifier by regression means trying to estimate the desired posterior probabilities with the function f (whose outputs are in $[0, 1]$).
- No closed-form solution. We apply gradient descent and obtain[?]:

$$\Delta \mathbf{w} = -\eta \frac{\partial E}{\partial \mathbf{w}} = \eta \sum_{n=1}^N (y_n - \theta_n) \theta_n (1 - \theta_n) \mathbf{x}_n, \quad \Delta w_0 = -\eta \frac{\partial E}{\partial w_0} = \eta \sum_{n=1}^N (y_n - \theta_n) \theta_n (1 - \theta_n)$$

where $\theta_n = \sigma(\mathbf{w}^T \mathbf{x}_n + w_0)$.

- As before, if the classes are linearly separable, this would drive $\|\mathbf{w}\| \rightarrow \infty$ and $\theta_n \rightarrow y_n \in \{0, 1\}$. Instead, we stop iterating when the right output dimension is larger than the others for each instance.

$K > 2$ classes: $y_n \in \{1, \dots, K\}$

- Same thing but:
 - We represent the labels y_n as a one-of- K binary vector $\mathbf{y}_n \in \{0, 1\}^K$ (containing a single 1 in position k if \mathbf{y}_n corresponds to class k), and consider the y_{kn} values as real values (which happen to be either 0 or 1).
 - We learn a separate function $f_k(\mathbf{x}; \mathbf{w}_k, w_{k0}) = \sigma(\mathbf{w}_k^T \mathbf{x} + w_{k0}) = \frac{1}{1 + \exp(-(\mathbf{w}_k^T \mathbf{x} + w_{k0}))}$ per output dimension $k = 1, \dots, K$.
- Hence, we solve a least-squares regression problem

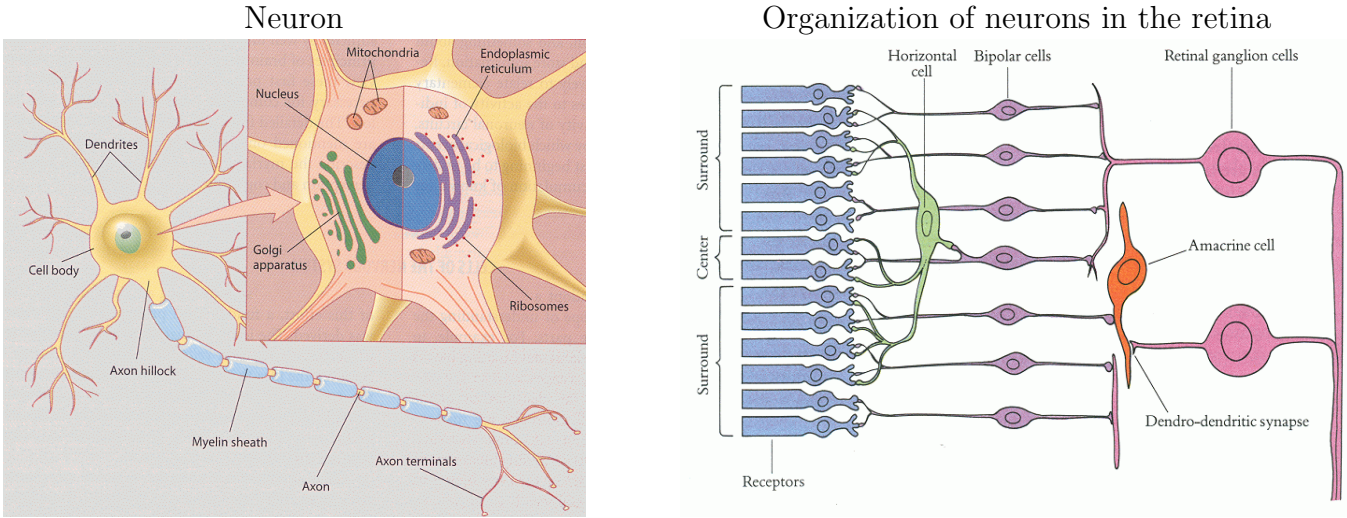
$$\min_{\Theta} E(\Theta) = \frac{1}{2} \sum_{n=1}^N \|\mathbf{y}_n - \mathbf{f}(\mathbf{x}_n; \Theta)\|^2 = \frac{1}{2} \sum_{k=1}^K \sum_{n=1}^N (y_{kn} - f_k(\mathbf{x}_n; \mathbf{w}_k, w_{k0}))^2$$

which is equivalent to K independent regression problems, each on an f_k .

- This formulation ignores the fact that the components of $\mathbf{f}(\mathbf{x})$ should sum to 1. A correct formulation, but more complicated to optimize, would use the softmax function for \mathbf{f} .

11 Multilayer perceptrons (artificial neural nets)

- Parametric nonlinear function approximators $f(\mathbf{x}; \Theta)$ for classification or regression.
- Originally inspired by neuronal circuits in the brain (McCullough-Pitts neurons, Rosenblatt's perceptron, etc.), and by the brain's ability to solve intelligent problems (visual and speech perception, navigation, planning, etc.).
Synapses between neurons = MLP weight parameters (which are modified during learning). Neuron firing = MLP unit nonlinearity.
- They are usually implemented in software in serial computers and more recently in parallel or distributed computers. There also exist VLSI implementations.





11.2 The perceptron

- *Linear perceptron*: computes a linear function $y = \sum_{d=1}^D w_d x_d + w_0 \in \mathbb{R}$ of an input vector $\mathbf{x} \in \mathbb{R}^D$ with weight vector $\mathbf{w} \in \mathbb{R}^D$ (or $y = \mathbf{w}^T \mathbf{x}$ with $\mathbf{w} \in \mathbb{R}^{D+1}$ and we augment \mathbf{x} with a 0th component of value 1). To have K outputs: $\mathbf{y} = \mathbf{W}\mathbf{x}$ with $\mathbf{W} \in \mathbb{R}^{K \times (D+1)}$.
- For classification, we can use:

– Two classes: $y = \sigma(\mathbf{w}^T \mathbf{x}) = \frac{1}{1 + \exp(-\mathbf{w}^T \mathbf{x})} \in (0, 1)$ (logistic).

– $K > 2$ classes: $y_k = \frac{\exp(\mathbf{w}_k^T \mathbf{x})}{\sum_{j=1}^K \exp(\mathbf{w}_j^T \mathbf{x})} \in (0, 1)$, $k = 1, \dots, K$ with $\sum_{k=1}^K y_k = 1$ (softmax).



11.3 Training a perceptron


- Apply *stochastic gradient descent*: to minimize the error $E(\mathbf{w}) = \sum_{n=1}^N e(\mathbf{w}; \mathbf{x}_n, y_n)$, repeatedly update the weights $\mathbf{w} \leftarrow \mathbf{w} + \Delta \mathbf{w}$ with $\Delta \mathbf{w} = -\eta \nabla e(\mathbf{w}; \mathbf{x}_n, y_n)$ and $n = 1, \dots, N$ (one epoch).
 - Regression by least-squares error: $e(\mathbf{w}; \mathbf{x}_n, y_n) = \frac{1}{2}(y_n - \mathbf{w}^T \mathbf{x}_n)^2 \Rightarrow \Delta \mathbf{w} = \eta(y_n - \mathbf{w}^T \mathbf{x}_n) \mathbf{x}_n$.
 - Classification by maximum likelihood, or equivalently cross-entropy:
 - * For two classes: $y_n \in \{0, 1\}$ and $e(\mathbf{w}; \mathbf{x}_n, y_n) = -y_n \log \theta_n - (1 - y_n) \log (1 - \theta_n)$ where $\theta_n = \sigma(\mathbf{w}^T \mathbf{x}_n) \Rightarrow \Delta \mathbf{w} = \eta(y_n - \theta_n) \mathbf{x}_n$.
 - * For $K > 2$ classes: $\mathbf{y}_n \in \{0, 1\}^K$ coded as 1-of- K and $e(\mathbf{w}; \mathbf{x}_n, \mathbf{y}_n) = -\sum_{k=1}^K y_{kn} \log \theta_{kn}$ where $\theta_{kn} = \frac{\exp(\mathbf{w}_k^T \mathbf{x}_n)}{\sum_{j=1}^K \exp(\mathbf{w}_j^T \mathbf{x}_n)} \Rightarrow \Delta w_{kd} = \eta(y_{kn} - \theta_{kn}) x_{dn}$ for $d = 0, \dots, D$ and $k = 1, \dots, K$.

Fig. 11.3 pseudocode

The original *perceptron algorithm* was a variation of stochastic gradient descent. For linearly separable problems, it converges in a finite (possibly large) number of iterations. For problems that are not linearly separable problems, it never converges.


11.4 Learning Boolean functions

- Boolean function: $\{0, 1\}^D \rightarrow \{0, 1\}$. Maps a vector of D bits to a single bit (truth value).
- Can be seen as a binary classification problem where the input instances are binary vectors.
- Since a perceptron can learn linearly separable problems, it can learn AND, OR but not XOR.



11.5 Multilayer perceptrons (MLPs)

- *Multilayer perceptron* (or *feedforward neural net*): nested sequence of perceptrons, with an *input layer*, an *output layer* and zero or more *hidden layers*: $\mathbf{x} \xrightarrow{\text{perceptrons}} \cdot \xrightarrow{\text{perceptrons}} \dots \xrightarrow{\text{perceptrons}} \mathbf{y}$.
- It can represent nonlinear discriminants (for classification) or functions (for regression).
- *Architecture of the MLP*: each layer has several units. Each unit h takes as input the output \mathbf{z} of the previous layer's units and applies to it a linear function $\mathbf{w}_h^T \mathbf{z}$ (using a weight vector \mathbf{w}_h , including a bias term) followed by a nonlinearity $s(t)$: output of unit $h = s(\mathbf{w}_h^T \mathbf{z})$.
- Typical nonlinearity functions $s(t)$ used:
 - *Logistic function*: $s(t) = \frac{1}{1+e^{-t}}$ (or softmax).
 - *Hyperbolic tangent*: $s(t) = \tanh t = \frac{e^t - e^{-t}}{e^t + e^{-t}}$.
 - *Rectified linear unit (ReLU)*: $s(t) = \max(0, t)$.
 - *Step function*: $s(t) = 0$ if $t < 0$, else 1.
 - *Identity function*: $s(t) = t$ (no nonlinearity).



The output layer uses as nonlinearity:


- For regression: the identity (so the outputs can take any real value).
- For classification: the sigmoid and a single unit ($K = 2$ classes), or the softmax and K units ($K > 2$ classes).

All nonlinearities (sigmoid, etc.) have a similar universal approximation ability, but some (e.g. ReLU) are easier to optimize than others.


- If all the layers are linear ($s(t) = t$ for all units in each layer) the MLP is overall a linear function, which is not useful, so we must have some nonlinear layers.
- Ex: an MLP with one hidden layer having H units, with inputs $\mathbf{x} \in \mathbb{R}^D$ and outputs $\mathbf{y} \in \mathbb{R}^{D'}$, where the hidden units are sigmoidal and the output units are linear:

$$\left. \begin{aligned} z_h(\mathbf{x}) &= \sigma(\mathbf{w}_h^T \mathbf{x}), \quad h = 1, \dots, H \\ f_i(\mathbf{x}) &= \mathbf{v}_i^T \mathbf{z}(\mathbf{x}), \quad i = 1, \dots, D' \end{aligned} \right\} \Rightarrow f_i(\mathbf{x}) = \sum_{h=1}^H v_{ih} \sigma \left(\sum_{d=1}^D w_{hd} x_d + w_{h0} \right) + v_{i0}.$$

An MLP with a single nonlinear hidden layer...



... solves the XOR problem



11.6 MLP as a universal approximator

- A Boolean function can always be written as a disjunction of conjunctions, and this can be implemented by an MLP with one hidden layer. Each conjunction (AND) is implemented by one hidden unit. Ex: $x_1 \text{ XOR } x_2 = (x_1 \text{ AND } \bar{x}_2) \text{ OR } (\bar{x}_1 \text{ AND } x_2)$. The disjunction (OR) is implemented by the output unit.

This existence proof generates very large MLPs (up to 2^D hidden units with D inputs). Practically, MLPs are much smaller.

- *Universal approximation*: any continuous function (satisfying mild assumptions) from \mathbb{R}^D to $\mathbb{R}^{D'}$ can be approximated by an MLP with a single hidden layer with an error as small as desired (by using sufficiently many hidden units).

Simple proof for the case of two hidden layers: we can enclose every input instance (or region) with a set of hyperplanes using hidden units in the first layer. A hidden unit in the second layer ANDs them together to bound the region. We then set the weight of the connection from that hidden unit to the output unit equal to the desired function value. This gives a piecewise constant approximation to the function. Its accuracy may be increased by using more hidden units that define finer regions in input space.

11.7 Learning MLPs: the backpropagation algorithm

- The *backpropagation algorithm* is (stochastic) gradient descent applied to an MLP using a least-squares or cross-entropy error function. Since the MLP consists of a nested sequence of functions (layers) $\mathbf{x} \xrightarrow{\mathbf{W}_1} \mathbf{z}_1 \xrightarrow{\mathbf{W}_2} \dots \xrightarrow{\mathbf{W}_K} \mathbf{y} = \mathbf{f}(\mathbf{x}; \Theta)$ with $\Theta = \{\mathbf{W}_1, \dots, \mathbf{W}_K\}$, each being a nonlinear perceptron, the gradient is computed with the *chain rule*, and can be interpreted as backpropagating the error at the output layer through the hidden layers back to the input.
- Ex: an MLP with one hidden layer having H units, with inputs $\mathbf{x} \in \mathbb{R}^D$ and outputs $\mathbf{y} \in \mathbb{R}^{D'}$, where the hidden units are sigmoidal and the output units are linear. The gradient is as follows:

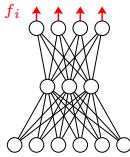
$$\left. \begin{array}{l} \text{– wrt second-layer weight } v_{ih}: \frac{\partial E}{\partial v_{ih}} = \frac{\partial E}{\partial f_i} \frac{\partial f_i}{\partial v_{ih}} = \delta_i \frac{\partial f_i}{\partial v_{ih}} \text{ (as with a perceptron} \\ \text{given fixed inputs } z_h) \\ \text{– wrt first-layer weight } w_{hd}: \frac{\partial E}{\partial w_{hd}} = \sum_{i=1}^{D'} \frac{\partial E}{\partial f_i} \frac{\partial f_i}{\partial z_h} \frac{\partial z_h}{\partial w_{hd}} = \sum_{i=1}^{D'} \delta_i \frac{\partial f_i}{\partial z_h} \frac{\partial z_h}{\partial w_{hd}} \end{array} \right\} \text{with “error” } \delta_i = \frac{\partial E}{\partial f_i}.$$

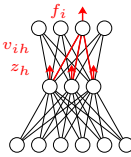
- With more hidden layers, the gradient wrt each layer’s weights is computed recursively.
- The updates using stochastic gradient descent with a minibatch \mathcal{B} to minimize an error function $E(\Theta; \{(\mathbf{x}_n, \mathbf{y}_n)\}_{n=1}^N) = \sum_{n=1}^N e(\Theta; \mathbf{x}_n, \mathbf{y}_n)$ are of the form $\Theta \leftarrow \Theta + \Delta\Theta$ with $\Delta\Theta = -\eta \sum_{n \in \mathcal{B}} \nabla_{\Theta} e(\Theta; \mathbf{x}_n, \mathbf{y}_n)$, where the gradients are given below. $\mathcal{B} = \{1, \dots, N\}$ gives batch gradient descent.

Regression We assume outputs \mathbf{y} of dimension D' . $\Theta = \{\mathbf{W}, \mathbf{V}\}$. *Least-squares error:*

Fig. 11.11 pseudocode

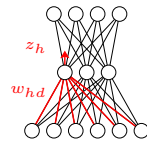
$$e(\mathbf{W}, \mathbf{V}; \mathbf{x}_n, \mathbf{y}_n) = e_n(\mathbf{W}, \mathbf{V}) = \frac{1}{2} \|\mathbf{y}_n - \mathbf{f}(\mathbf{x}_n; \mathbf{W}, \mathbf{V})\|^2 = \frac{1}{2} \sum_{i=1}^{D'} (y_{in} - f_i(\mathbf{x}_n; \mathbf{W}, \mathbf{V}))^2$$



$$e = e(\{f_i\}_{i=1}^{D'}) = \frac{1}{2} \sum_{i=1}^{D'} \underbrace{(y_i - f_i)^2}_{\delta_i = \text{error of output unit } i}$$


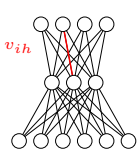
$$f_i = f_i(\{v_{ih}, z_h\}_{h=0}^H) = \sum_{h=0}^H v_{ih} z_h,$$

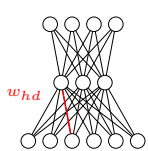
output of output unit i



$$z_h = z_h(\{w_{hd}\}_{d=0}^D) = \sigma\left(\sum_{d=0}^D w_{hd} x_d\right)$$

output of hidden unit h



$$\frac{\partial e_n}{\partial v_{ih}} = \underbrace{-(y_{in} - f_i(\mathbf{x}_n))}_{\partial e_n / \partial f_i} \underbrace{z_{hn}}_{\partial f_i / \partial v_{ih}}$$


$$\frac{\partial e_n}{\partial w_{hd}} = \left(\sum_{i=1}^{D'} \underbrace{-(y_{in} - f_i(\mathbf{x}_n))}_{\partial e_n / \partial f_i} \underbrace{v_{ih}}_{\partial f_i / \partial z_h} \right) \underbrace{z_{hn} (1 - z_{hn})}_{\partial z_h / \partial w_{hd}} x_{dn}.$$

Classification, two classes We need a single output unit $f(\mathbf{x}_n)$. $\Theta = \{\mathbf{W}, \mathbf{v}\}$. *Cross-entropy error:*


$$e(\mathbf{W}, \mathbf{v}; \mathbf{x}_n, \mathbf{y}_n) = e_n(\mathbf{W}, \mathbf{v}) = -y_n \log(f(\mathbf{x}_n; \mathbf{W}, \mathbf{v})) - (1 - y_n) \log(1 - f(\mathbf{x}_n; \mathbf{W}, \mathbf{v}))$$

$$\frac{\partial e_n}{\partial v_h} = -(y_n - f(\mathbf{x}_n)) z_{hn} \quad \frac{\partial e_n}{\partial w_{hd}} = -(y_n - f(\mathbf{x}_n)) v_h z_{hn} (1 - z_{hn}) x_{dn}.$$


Classification, $K > 2$ classes We need one output unit $f_i(\mathbf{x}_n)$ per class $i = 1, \dots, K$, given by the softmax function. $\Theta = \{\mathbf{W}, \mathbf{V}\}$. *Cross-entropy error:*

$$e(\mathbf{W}, \mathbf{V}; \mathbf{x}_n, \mathbf{y}_n) = e_n(\mathbf{W}, \mathbf{V}) = -\sum_{i=1}^K y_{in} \log(f_i(\mathbf{x}_n; \mathbf{W}, \mathbf{V})), \quad f_i(\mathbf{x}_n) = \frac{\exp(o_{in})}{\sum_{k=1}^K \exp(o_{kn})}, \quad o_{in} = \sum_{h=0}^H v_{ih} z_{hn}$$


$$\frac{\partial e_n}{\partial v_{ih}} = -(y_{in} - f_i(\mathbf{x}_n)) z_{hn} \quad \frac{\partial e_n}{\partial w_{hd}} = \left(\sum_{i=1}^{D'} -(y_{in} - f_i(\mathbf{x}_n)) v_{ih} \right) z_{hn} (1 - z_{hn}) x_{dn}.$$



MLP with $H = 2$ hidden units after 100, 200 and 300 epochs



Error on training & validation sets



a) hyperplanes of hidden units in L1;
b) hidden unit outputs;
c) hidden unit outputs \times weights in L2

11.8 Training procedures

Training neural nets in practice is tricky and requires some expertise and trial-and-error.

Improving convergence Training a neural net can be computationally very costly:

- The dataset $\{(\mathbf{x}_n, \mathbf{y}_n)\}_{n=1}^N$ and the number of weights can both be very large in practical problems.
- Gradient descent converges very slowly in optimization problems with many parameters.
- *Vanishing gradients problem*: there is one product $\sigma'(z) = \sigma(z)(1 - \sigma(z))$ per layer. If the value of z is large, which will happen if the weights or inputs at that layer are large, then the sigmoid saturates and $\sigma'(z) \approx 0$, so the gradient becomes tiny, and it takes many iterations to make progress. This is particularly problematic with deep networks (having several layers of hidden units). Other nonlinearities have less of a problem, e.g. ReLU.

The following techniques are typically used to speed up the convergence:

- *Initial weights*: small random values (e.g. uniform in $[-0.01, 0.01]$) so as not to saturate the sigmoids.
- *Normalizing the inputs* so they have zero mean and unit variance speeds up the optimization (since we use a single step size η for all parameters).
- *Momentum*: we use an update (for any given weight w) $\Delta w^{\text{new}} = -\eta \frac{\partial E}{\partial w} + \alpha \Delta w^{\text{old}}$ where α is generally taken between 0.5 and 1. This tends to smooth the trajectory of the iterates and reduce oscillations. It is a limited form of conjugate gradients.
- *Rescaling the learning rate η_w for each weight w* by using a diagonal approximation to the Hessian of the objective function $E(\mathbf{w})$. This helps to correct for the fact that weights in different layers, or subject to weight sharing, may have different effects on the output.
- It is also possible to use other optimization methods (modified Newton's method, conjugate gradients, L-BFGS, etc.) but, for problems with large N and redundant samples, they don't seem to improve significantly over SGD (with properly tuned learning rate, minibatch size and momentum term).

Overtraining



- *Model selection* for the number of hidden units (hence weights): the bias-variance tradeoff applies as usual:
 - Neural nets with many hidden units can achieve a very low training error but memorize the noise as well as the signal, hence overfit.
 - Neural nets with few hidden units can have a large bias, hence underfit.

The number of hidden units can be estimated by cross-validation.

- More generally, one needs to select the overall architecture of the neural net (number of layers and of hidden units in each layer, connectivity between them, choice of nonlinearity, etc.). This is done by trial and error and can be very time-consuming.
- *Early stopping*: during the (stochastic) gradient descent optimization (with a given number of hidden units), we monitor the error on a validation set and stop training when the validation error doesn't keep decreasing. This helps to prevent overfitting as well.
- *Weight decay*: we discourage weights from taking large values by adding to the objective function, or equivalently to the update, a penalty term:

$$E'(\mathbf{w}) = E(\mathbf{w}) + \frac{\lambda}{2} \sum_i w_i^2 \iff \Delta w_i = -\eta \left(\frac{\partial E}{\partial w_i} + \lambda w_i \right).$$

This has the effect of choosing networks with small weights as long as they have a low training error (depending on λ). Such networks are smoother and generalize better to unseen data. The value of λ is set by cross-validation.



Structuring the network


- Depending on the application and data, some types of connectivity may be more effective than having all layers be fully-connected.
- *Convolutional neural nets*: useful with images or time series that have local structure (around a location in space or time, respectively). Each hidden unit receives input from a subset of the input layer's units, corresponding to a spatial or temporal window in the input image or signal.
Ex: in handwritten digit images, nearby pixels are correlated and give rise to local features such as edges, corners or T-junctions. A stroke or a digit can be seen as a combination of such primitive features.
- *Weight sharing*: the values of certain weights are constrained to be equal. For example, with convolutional neural nets, the weights in the window are the same for every hidden unit. This corresponds to using a filter that is homogenous in space or time, and helps to detect features regardless of their location in the input. Ex: edges at different locations and of different orientations.

- Convolutional neural nets with weight sharing have a much smaller number of weights, can be trained faster, and usually give better results than fully-connected networks.
- This process may be repeated in successive layers and make it possible for the network to learn a hierarchical representation of the data with each layer learning progressively more complex and abstract concepts.

Ex: pixels \rightarrow edges in different orientations \rightarrow corners, T-junctions, contour segments \rightarrow parts \rightarrow objects \rightarrow classes of objects.

- With temporal data, as in speech recognition or machine translation, one may use *time-delay neural nets* or *recurrent neural nets*.
- *Deep learning* refers to neural networks having multiple layers that can learn hierarchical representations of complex data such as images. Properly trained on large, labeled datasets (using GPUs), they have achieved impressive results in recent years in difficult tasks such as object recognition, speech recognition, machine translation or game learning. This success is due to the ability of the network to learn nonlinear mappings in high-dimensional spaces, and to the fact that the network is learned end-to-end, i.e., all its weights are optimized with little human contribution (rather than e.g. having a preprocessing layer that is fixed by hand by an expert).

Ex: instead of extracting preselected features from an image (e.g. SIFT features) or from the speech waveform (e.g. MFCCs), the first layer(s) of the net learn those features optimally.



Hints

- *Hints* are properties of the target function that are known to us independent of the training data. They should be built into the neural net if possible, because they help to learn a good network, particularly when the training data is limited.
- Ex: *invariance*. The output of a classifier is often invariant to certain known transformations of the input:
 - Object recognition in images: translation, rotation, scaling.
 - Speech recognition: loudness, reverberation.
- Hints may be incorporated in different ways, such as:
 - *Virtual examples*: generate multiple copies of every object to be recognized at different locations, rotations and scales, and add them as training examples with the same label. This doesn't change the optimization algorithm but it increases the training set size considerably.
 - *Preprocessing stage*: the image could be centered, aligned and rescaled to a standard position/orientation/scale before being fed to the network.
 - *Putting hints into the network structure*, as with convolutional nets and weight sharing.

11.11 Dimensionality reduction: autoencoders

- *Autoencoder*: an MLP that, given a training set $\{\mathbf{x}_n\}_{n=1}^N \in \mathbb{R}^D$ without labels, tries to reconstruct its input (“labels” $\mathbf{y}_n = \mathbf{x}_n$), but has a *bottleneck layer*, whose number of hidden units H is smaller than the dimension D of the input \mathbf{x} . Its output layer is linear and it is trained by least-squares error:


$$\min_{\mathbf{W}_1, \mathbf{W}_2} E(\mathbf{W}_1, \mathbf{W}_2) = \frac{1}{2} \sum_{n=1}^N \|\mathbf{y}_n - \mathbf{f}(\mathbf{F}(\mathbf{x}_n; \mathbf{W}_1); \mathbf{W}_2)\|^2$$

where $\mathbf{F}: \mathbb{R}^D \rightarrow \mathbb{R}^H$ is the *encoder* network and $\mathbf{f}: \mathbb{R}^H \rightarrow \mathbb{R}^D$ the *decoder* network.

- This forces the overall network $\mathbf{f}(\mathbf{F}(\mathbf{x}))$ to learn an optimal low-dimensional representation $\mathbf{z}_n = \mathbf{F}(\mathbf{x}_n) \in \mathbb{R}^H$ of each instance $\mathbf{x}_n \in \mathbb{R}^D$ in the bottleneck (or “code”) layer. The codes (= outputs of the bottleneck hidden layer) are optimal for reconstructing the input instances.
- If the hidden layers are all linear then the network becomes equivalent to PCA?

The hidden unit weights need not be the H leading eigenvectors of the data covariance matrix, but they span the same subspace.

If they are nonlinear (e.g. sigmoidal) then it learns a nonlinear dimensionality reduction.



- If a neural network is trained for classification and has a bottleneck (a hidden layer with dimension smaller than the input), then the network will learn a low-dimensional representation in that hidden layer that is optimal for classification (similar to LDA if using linear layers).

12 Local models: radial basis function networks

12.3 Radial basis function networks



- Assume a dataset $\{\mathbf{x}_n, \mathbf{y}_n\}_{n=1}^N$ with $\mathbf{x}_n \in \mathbb{R}^D$ and $\mathbf{y}_n \in \mathbb{R}^{D'}$ for regression or $y_n \in \{1, \dots, K\}$ for classification.
- *Radial basis function (RBF) network*:

$$\mathbf{f}(\mathbf{x}) = \sum_{h=1}^H \mathbf{w}_h \phi_h(\mathbf{x}) \quad \phi_h(\mathbf{x}) = \exp\left(-\frac{1}{2} \|\mathbf{x} - \boldsymbol{\mu}_h\|^2 / \sigma^2\right)$$


where $\{\phi_h(\cdot)\}_{h=1}^H$ are radial basis functions, typically (proportional to) Gaussians with *centroids* $\{\boldsymbol{\mu}_h\}_{h=1}^H \subset \mathbb{R}^D$ and common *width* σ , and $\{\mathbf{w}_h\}_{h=1}^H \subset \mathbb{R}^{D'}$ are *weights*. We take $\phi_1(\mathbf{x}) \equiv 1$ if we want to use a *bias*.

It is also possible to use a separate width σ_h per RBF.

- They can be seen as a feedforward neural net with a single hidden layer of Gaussian units and an output layer of linear units. They are not usually constructed with more hidden layers, unlike MLPs.
- Two types of representation of information in neural nets: assume for simplicity that units output either 0 (not activated) or 1 (activated):
 - *Distributed representation*, as in sigmoidal MLPs: an input $\mathbf{x} \in \mathbb{R}^D$ is encoded by the simultaneous activation of many hidden units.
Every hidden unit splits the space into two half-spaces, so \mathbf{x} will activate half of the units on average.
 - *Local representation*, as in RBF nets: an input $\mathbf{x} \in \mathbb{R}^D$ is encoded by the simultaneous activation of few hidden units.
Every unit splits the space into inside/outside a hypersphere, so \mathbf{x} activates only a few units, unless the width σ is large.
Each unit is locally tuned to a small area of the input space called its *receptive field*, and the input space is paved with such units.
Neurons in the primary visual cortex respond only to highly localized stimuli in retinal position and angle of visual orientation.



Local representation in the space of (p_1, p_2, p_3)
 $\mathbf{x}^a : (1.0, 0.0, 0.0)$
 $\mathbf{x}^b : (0.0, 0.0, 1.0)$
 $\mathbf{x}^c : (1.0, 1.0, 0.0)$



Distributed representation in the space of (h_1, h_2)
 $\mathbf{x}^a : (1.0, 1.0)$
 $\mathbf{x}^b : (0.0, 1.0)$
 $\mathbf{x}^c : (1.0, 0.0)$

- *Training for regression*: we minimize the least-squares error

$$E(\{\mathbf{w}_h, \boldsymbol{\mu}_h\}_{h=1}^H, \sigma) = \frac{1}{2} \sum_{n=1}^N \|\mathbf{y}_n - \mathbf{f}(\mathbf{x}_n)\|^2 + \lambda \sum_{h=1}^H \|\mathbf{w}_h\|^2$$

where $\lambda \geq 0$ is a regularization user parameter which controls the smoothness of \mathbf{f} . We can use gradient descent, where the gradient is computed using the chain rule and is similar to that of an MLP. However, RBF nets can be trained in an approximate but much simpler and faster way as follows:

1. Set the centroids $\{\boldsymbol{\mu}_h\}_{h=1}^H$ in an unsupervised way using only the input points $\{\mathbf{x}_n\}_{n=1}^N$, usually by running K -means with H centroids, or by selecting H points at random.
2. Set the width σ by cross-validation (see below).
It is also possible to set σ to a rough value as follows. Having run k -means, we compute the distance from each centroid $\boldsymbol{\mu}_h$ to the farthest point \mathbf{x}_n in its cluster. We then set σ to the average such distance over the H centroids.
3. Given the centroids and width, the values $\phi_h(\mathbf{x}_n)$ are fixed, and the weights $\{\mathbf{w}_h\}_{h=1}^H$ are determined by optimizing E , which reduces to a simple linear regression. We solve the linear system[Ⓔ] (recall the normal equations of chapter 5):

$$(\boldsymbol{\Phi}\boldsymbol{\Phi}^T + \lambda\mathbf{I})\mathbf{W} = \boldsymbol{\Phi}\mathbf{Y}^T \quad \boldsymbol{\Phi}_{H \times N} = (\phi_h(\mathbf{x}_n))_{hn}, \quad \mathbf{W}_{H \times D'} = (\mathbf{w}_1, \dots, \mathbf{w}_H)^T, \quad \mathbf{Y}_{D' \times N} = (\mathbf{y}_1, \dots, \mathbf{y}_N).$$

- *Training for classification*: we minimize the cross-entropy error

$$E(\{\mathbf{w}_h, \boldsymbol{\mu}_h\}_{h=1}^H, \sigma) = - \sum_{n=1}^N \sum_{i=1}^K y_{in} \log \theta_{in} + \lambda \sum_{h=1}^H \|\mathbf{w}_h\|^2$$

$$\theta_{in} = \frac{\exp(f_i(\mathbf{x}_n))}{\sum_{k=1}^K \exp(f_k(\mathbf{x}_n))} \quad f_i(\mathbf{x}) = \sum_{h=1}^H w_{ih} \phi_h(\mathbf{x}).$$

That is, we use an RBF net with H BF's and K outputs, which we pass through a softmax. This can either be optimized by gradient descent, or trained approximately as above (but the optimization over the weights cannot be solved by a linear system and requires gradient descent itself).

- *Model selection*: the complexity of a RBF net is controlled by:
 - The number H of basis functions: $H \downarrow$ high bias, low variance; $H \uparrow$ low bias, high variance.
 - The regularization parameter $\lambda \geq 0$: the larger, the smoother the RBF net.
Numerically, we typically need $\lambda \gtrsim 10^{-7}$ to make the linear system sufficiently well conditioned.
 - The width σ (if not optimized over or set by hand): the larger, the smoother the RBF net.

They are all set by cross-validation. We usually search over a range of values of H , $\log \lambda$ and $\log \sigma$.

- Because they are localized, RBF nets may need many basis functions to achieve a desired accuracy.
- *Universal approximation*: any continuous function (satisfying mild assumptions) from \mathbb{R}^D to $\mathbb{R}^{D'}$ can be approximated by a RBF net with an error as small as desired (by using sufficiently many basis functions).

Intuitive proof: we can enclose every input instance (or region) with a basis function having as centroid that instance and a small enough width so it has little overlap with other BF's. This way we “grid” the input space \mathbb{R}^D with sufficiently many BF's, so that only one BF is active for a given instance \mathbf{x}_n . Finally, we set the weight of each BF to the desired function value. This gives an approximately piecewise constant approximation to the function (essentially equal to a nonparametric kernel smoother). Its accuracy may be increased by using more hidden units and placing a finer grid in the input.

12.5 Normalized basis functions

- With Gaussian BFs, it is possible that $\phi_h(\mathbf{x}) = 0$ for all $h = 1, \dots, H$ for some $\mathbf{x} \in \mathbb{R}^D$. This happens if \mathbf{x} is far enough (wrt σ) from each centroid.
- We can define the basis functions as normalized Gaussians:

$$\phi_h(\mathbf{x}) = \frac{\exp\left(-\frac{1}{2}\|\mathbf{x} - \boldsymbol{\mu}_h\|^2/\sigma^2\right)}{\sum_{j=1}^H \exp\left(-\frac{1}{2}\|\mathbf{x} - \boldsymbol{\mu}_j\|^2/\sigma^2\right)} \implies \phi_h(\mathbf{x}) \in (0, 1) \text{ and } \sum_{h=1}^H \phi_h(\mathbf{x}) = 1.$$

Now, even if \mathbf{x} is far (wrt σ) from each centroid, at least one BF will have a nonzero value.

- $\phi_h(\mathbf{x})$ can be seen as the posterior probability $p(h|\mathbf{x})$ of BF h given input \mathbf{x} assuming a Gaussian mixture model $p(\mathbf{x})$, where component h has mean $\boldsymbol{\mu}_h$, and all components have the same proportion $\frac{1}{H}$ and the same, isotropic covariance matrix $\boldsymbol{\Sigma}_h = \sigma^2\mathbf{I}$.
- The nonparametric kernel smoother of chapter 8 using a Gaussian kernel of bandwidth $h > 0$

$$\mathbf{g}(\mathbf{x}) = \sum_{n=1}^N \frac{K(\|(\mathbf{x} - \mathbf{x}_n)/h\|)}{\sum_{n'=1}^N K(\|(\mathbf{x} - \mathbf{x}_{n'})/h\|)} \mathbf{y}_n$$

is identical to a RBF network with the following special choices:

- it uses normalized BFs;
- it has $H = N$ BFs, one per input data point;
- the centroids are the input points: $\boldsymbol{\mu}_n = \mathbf{x}_n$;
- the weights equal the output vectors: $\mathbf{w}_n = \mathbf{y}_n$.

13 Kernel machines (support vector machines, SVMs)

- Discriminant-based method: it models the classification boundary, not the classes themselves.
- It can produce linear classifiers (*linear SVMs*) or nonlinear classifiers (*kernel SVMs*).
- Very effective in practice with certain problems:
 - It gives good generalization to test cases.
 - The optimization problem is convex and has a unique solution (no local optima).
 - It can be trained on reasonably large datasets.
 - Special kernels can be defined for many applications.
- It can be extended beyond classification to solve regression, dimensionality reduction, outlier detection and other problems.

The basic approach is the same in all cases: maximize the margin and penalize deviations, resulting in a convex quadratic program.
- We have seen several linear classifiers: logistic regression, Gaussian classes with common covariance, now linear SVMs... They give different results because they have different inductive bias, i.e., make different assumptions (the objective function, etc.).

13.2 Binary classification, linearly separable case: optimal separating hyperplane

- Consider a dataset $\{\mathbf{x}_n, y_n\}_{n=1}^N$ where $\mathbf{x}_n \in \mathbb{R}^D$ and $y_n \in \{-1, +1\}$, and a linear discriminant function $g(\mathbf{x}) = \mathbf{w}^T \mathbf{x} + w_0$, i.e., a hyperplane, where $\mathbf{w} \in \mathbb{R}^D$ and $w_0 \in \mathbb{R}$. The classification rule induced by the discriminant is given by $\text{sgn}(\mathbf{w}^T \mathbf{x} + w_0) \in \{-1, +1\}$.
- Assume the two classes are linearly separable, so there exist an infinite number of separating hyperplanes. Although they are all equally good on the training set, they differ with test data. Which one is best?
- We define the *margin* of a separating hyperplane as the distance to the hyperplane of the closest instance to it. *We want to find the hyperplane having the largest margin.*
- This has two advantages:
 - It provides a unique solution to the separating hyperplane problem.
 - It gives better classification performance on test data.

If noise perturbs slightly an instance \mathbf{x}_n , it will still be correctly classified if it is not too close to the boundary.
- The signed distance of a point \mathbf{x}_n to the hyperplane is $(\mathbf{w}^T \mathbf{x}_n + w_0) / \|\mathbf{w}\| \in \mathbb{R}$, hence its absolute value is $y_n(\mathbf{w}^T \mathbf{x}_n + w_0) / \|\mathbf{w}\| \geq 0$ (since $y_n \in \{-1, +1\}$). Let \mathbf{x}_{n^*} be the closest instance to the hyperplane and $\rho = y_{n^*}(\mathbf{w}^T \mathbf{x}_{n^*} + w_0) / \|\mathbf{w}\| > 0$ its distance (the margin), i.e., $y_n(\mathbf{w}^T \mathbf{x}_n + w_0) / \|\mathbf{w}\| \geq \rho \forall n = 1, \dots, N$. To make this distance as large as possible, we can maximize ρ subject to $y_n(\mathbf{w}^T \mathbf{x}_n + w_0) / \|\mathbf{w}\| \geq \rho \forall n = 1, \dots, N$. But $y_n(\mathbf{w}^T \mathbf{x}_n + w_0) / \|\mathbf{w}\|$ is invariant to rescaling both \mathbf{w} and w_0 by a constant factor, so we arbitrarily fix this factor by requiring $\|\mathbf{w}\| = 1/\rho$ (or, equivalently, by requiring $y_n(\mathbf{w}^T \mathbf{x}_n + w_0) = 1$ for the instance \mathbf{x}_n closest to the hyperplane). Then, since maximizing ρ is the same as minimizing $1/2\rho^2 = \frac{1}{2} \|\mathbf{w}\|^2$, we finally reach the following *maximum margin* optimization problem:

$$\text{Primal QP: } \min_{\mathbf{w} \in \mathbb{R}^D, w_0 \in \mathbb{R}} \frac{1}{2} \|\mathbf{w}\|^2 \quad \text{s.t.} \quad y_n(\mathbf{w}^T \mathbf{x}_n + w_0) \geq 1, \quad n = 1, \dots, N.$$

- This is a constrained optimization problem, specifically a *convex quadratic program (QP)*, defined on $D + 1$ variables (\mathbf{w}, w_0) with N linear constraints and a quadratic objective function. It has a unique minimizer (\mathbf{w}, w_0) , which can be found with different QP algorithms. The closest instances to each side of the hyperplane are at a distance? $1 / \|\mathbf{w}\|$ from the hyperplane.
- Solving the *primal problem*, a QP on $D + 1$ variables, is equivalent to solving the following *dual problem*, another convex QP but on N variables $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_N)^T$ and $N + 1$ constraints ($\alpha_1, \dots, \alpha_N$ are the *Lagrange multipliers* of the N constraints in the primal QP):

$$\text{Dual QP: } \min_{\boldsymbol{\alpha} \in \mathbb{R}^N} \frac{1}{2} \sum_{n,m}^N (y_n y_m (\mathbf{x}_n^T \mathbf{x}_m)) \alpha_n \alpha_m - \sum_{n=1}^N \alpha_n \quad \text{s.t.} \quad \sum_{n=1}^N y_n \alpha_n = 0, \quad \alpha_n \geq 0, \quad n = 1, \dots, N.$$

This can be solved in $\mathcal{O}(N^3 + N^2 D)$ times and $\mathcal{O}(N^2)$ space (because we need the dot products $\mathbf{x}_n^T \mathbf{x}_m$ for $n, m = 1, \dots, N$). The optimal $\boldsymbol{\alpha} \in \mathbb{R}^N$ has the property that $\alpha_n > 0$ only for those instances that lie on the margin (the closest ones to the hyperplane), i.e., $y_n(\mathbf{w}^T \mathbf{x}_n + w_0) = 1$, called the *support vectors (SVs)*. For all other instances, which lie beyond the margin, $\alpha_n = 0$.

- From the optimal $\boldsymbol{\alpha} \in \mathbb{R}^N$, we recover the solution (\mathbf{w}, w_0) as follows:

$$\mathbf{w} = \sum_{n=1}^N \alpha_n y_n \mathbf{x}_n = \sum_{n \in \text{SVs}} \alpha_n y_n \mathbf{x}_n \quad w_0 = y_n - \mathbf{w}^T \mathbf{x}_n \text{ for any support vector } \mathbf{x}_n?$$

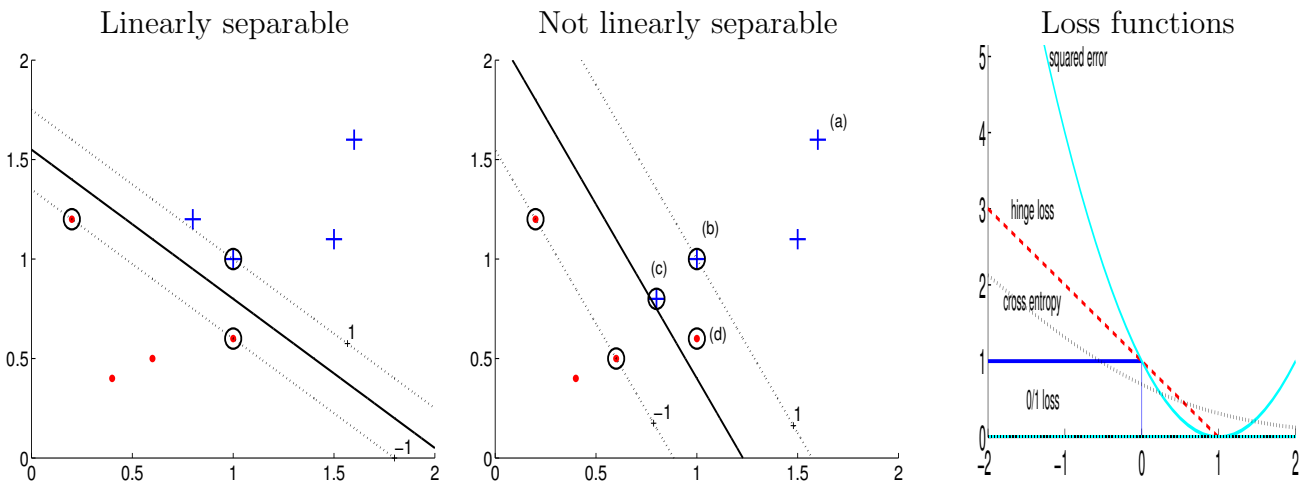
so the optimal weight vector can be written as a l.c. of the SVs. This allows kernelization later.

- The resulting, linear discriminant (the *linear SVM*) is

$$g(\mathbf{x}) = \mathbf{w}^T \mathbf{x} + w_0 = \sum_{n=1}^N \alpha_n y_n \mathbf{x}_n^T \mathbf{x} + w_0 = \sum_{n \in \text{SVs}} \alpha_n y_n \mathbf{x}_n^T \mathbf{x} + w_0.$$

- Properties of the support vectors:

- They are the instances that are closest to the boundary and thus the more difficult ones to classify.
- The number of SVs is usually much smaller than N , though this depends on the dimension D and the problem.
- Solving the problem using as training set just the SVs gives the same result.
In practice, we don't know which instances are SVs, so we have to solve the QP using the entire dataset.



In terms of ξ_n : (a) $\xi_n = 0$, (b) $\xi_n = 0$, (c) $0 < \xi_n < 1$, (d) $\xi_n > 1$;

In terms of α_n : (a) $\alpha_n = 0$, (b) $0 < \alpha_n < C$, (c)–(d) $\alpha_n = C$;

(b)–(d) are SVs (circled instances $\oplus \odot$, on or beyond the margin)

13.3 Binary classification, not linearly separable case: soft margin hyperplane

- If the dataset is not linearly separable, we look for the hyperplane that incurs least classification error while having the largest margin. For each data point \mathbf{x}_n , we allow a deviation $\xi_n \geq 0$ (a *slack variable*) from the margin, but penalize such deviations proportionally to $C > 0$:

$$\text{Primal QP: } \min_{\mathbf{w} \in \mathbb{R}^D, w_0 \in \mathbb{R}, \boldsymbol{\xi} \in \mathbb{R}^N} \frac{1}{2} \|\mathbf{w}\|^2 + C \sum_{n=1}^N \xi_n \quad \text{s.t.} \quad \begin{cases} y_n(\mathbf{w}^T \mathbf{x}_n + w_0) \geq 1 - \xi_n, \\ \xi_n \geq 0, \quad n = 1, \dots, N. \end{cases}$$

$C > 0$ is a user parameter that controls the tradeoff between maximizing the margin ($2/\|\mathbf{w}\|^2$) and minimizing the deviations or “soft error” ($\sum_{n=1}^N \xi_n$):

- $C \uparrow$: tries to avoid every single error but reduces the margin so it can overfit;
- $C \downarrow$: enlarges the margin but ignores most errors so it can underfit.

C is set by cross-validation. Typically, one chooses C on a log scale, e.g. $C \in \{10^{-6}, 10^{-5}, \dots, 10^{+5}, 10^{+6}\}$.

- This is again a convex QP on $D + 1 + N$ variables and its unique minimizer $(\mathbf{w}, w_0, \boldsymbol{\xi})$ can be found with different QP algorithms. It is equivalent to solving the dual QP on N variables:

$$\text{Dual QP: } \min_{\boldsymbol{\alpha} \in \mathbb{R}^N} \frac{1}{2} \sum_{n,m} (y_n y_m (\mathbf{x}_n^T \mathbf{x}_m)) \alpha_n \alpha_m - \sum_{n=1}^N \alpha_n \quad \text{s.t.} \quad \sum_{n=1}^N y_n \alpha_n = 0, \quad C \geq \alpha_n \geq 0, \quad n = 1, \dots, N$$

which differs from the dual QP of the linearly separable case in the extra “ $\alpha_n \leq C$ ” constraints. The optimal $\boldsymbol{\alpha} \in \mathbb{R}^N$ has the property that $\alpha_n > 0$ only for those instances that lie on or within the margin or are misclassified, i.e., $y_n(\mathbf{w}^T \mathbf{x}_n + w_0) \leq 1$, called the *support vectors (SVs)*. For all other instances, which lie beyond the margin, $\alpha_n = 0$. As a result, for each training point $n = 1, \dots, N$ we can have:

- $\xi_n = 0$: \mathbf{x}_n is correctly classified and beyond the margin (here, $\alpha_n = 0$).
- $0 < \xi_n < 1$: \mathbf{x}_n is correctly classified but within the margin (here, $0 < \alpha_n \leq C$).
- $\xi_n \geq 1$: \mathbf{x}_n is misclassified (here, $\alpha_n = C$).

- From the optimal $\boldsymbol{\alpha} \in \mathbb{R}^N$, we recover the solution (\mathbf{w}, w_0) as before:

$$\mathbf{w} = \sum_{n=1}^N \alpha_n y_n \mathbf{x}_n = \sum_{n \in \text{SVs}} \alpha_n y_n \mathbf{x}_n \quad w_0 = y_n - \mathbf{w}^T \mathbf{x}_n \text{ for any } \mathbf{x}_n \text{ with } 0 < \alpha_n < C.$$

- The resulting, linear discriminant (the *linear SVM*) is

$$g(\mathbf{x}) = \mathbf{w}^T \mathbf{x} + w_0 = \sum_{n=1}^N \alpha_n y_n \mathbf{x}_n^T \mathbf{x} + w_0 = \sum_{n \in \text{SVs}} \alpha_n y_n \mathbf{x}_n^T \mathbf{x} + w_0.$$

- The primal problem can be written equivalently but without constraints by eliminating ξ_n as:

$$\min_{\mathbf{w} \in \mathbb{R}^D, w_0 \in \mathbb{R}} \frac{1}{2} \|\mathbf{w}\|^2 + C \sum_{n=1}^N [1 - y_n g(\mathbf{x}_n)]_+ \quad \text{where} \quad g(\mathbf{x}) = \mathbf{w}^T \mathbf{x} + w_0.$$

This defines the *hinge loss*, which penalizes errors linearly:

$$[1 - y_n g(\mathbf{x}_n)]_+ = \max(0, 1 - y_n g(\mathbf{x}_n)) = \begin{cases} 0 & \text{if } y_n g(\mathbf{x}_n) \geq 1 \\ 1 - y_n g(\mathbf{x}_n) & \text{otherwise} \end{cases}$$

and is more robust against outliers compared to the quadratic loss $(1 - y_n g(\mathbf{x}_n))^2$.

13.5 Kernelization: kernel SVMs for nonlinear classification

- We map the data points \mathbf{x} to a higher-dimensional space. Learning a linear model in the new space corresponds to learning a nonlinear model in the original space.
- Assume we map $\mathbf{x} \in \mathbb{R}^D \rightarrow \mathbf{z} = \phi(\mathbf{x}) = (\phi_1, \dots, \phi_K)^T \in \mathbb{R}^K$ where $K > D$ (usually, K is much larger or even infinite), and $z_1 = \phi_1(\mathbf{x}) \equiv 1$ (so we don't have to write a bias term explicitly). The discriminant is now $g(\mathbf{x}) = \mathbf{w}^T \phi(\mathbf{x}) = \sum_{i=1}^K w_i \phi_i(\mathbf{x})$.
- The primal and dual soft-margin QPs are as before but replacing \mathbf{x}_n with $\phi(\mathbf{x}_n)$:

$$\text{Primal QP: } \min_{\mathbf{w} \in \mathbb{R}^D, \xi \in \mathbb{R}^N} \frac{1}{2} \|\mathbf{w}\|^2 + C \sum_{n=1}^N \xi_n \quad \text{s.t.} \quad y_n(\mathbf{w}^T \phi(\mathbf{x}_n)) \geq 1 - \xi_n, \quad \xi_n \geq 0, \quad n = 1, \dots, N$$

$$\text{Dual QP: } \min_{\alpha \in \mathbb{R}^N} \frac{1}{2} \sum_{n,m} (y_n y_m \underbrace{\phi(\mathbf{x}_n)^T \phi(\mathbf{x}_m)}_{K(\mathbf{x}_n, \mathbf{x}_m)}) \alpha_n \alpha_m - \sum_{n=1}^N \alpha_n \quad \text{s.t.} \quad \sum_{n=1}^N y_n \alpha_n = 0, \quad C \geq \alpha_n \geq 0, \quad n = 1, \dots, N$$

$$\text{Optimal solution: } \mathbf{w} = \sum_{n=1}^N \alpha_n y_n \phi(\mathbf{x}_n) = \sum_{n \in \text{SVs}} \alpha_n y_n \phi(\mathbf{x}_n)$$

$$\text{Discriminant: } g(\mathbf{x}) = \mathbf{w}^T \phi(\mathbf{x}) = \sum_{n=1}^N \alpha_n y_n \underbrace{\phi(\mathbf{x}_n)^T \phi(\mathbf{x})}_{K(\mathbf{x}_n, \mathbf{x})}$$

- *Kernelization*: we replace the inner product $\phi(\mathbf{x}_n)^T \phi(\mathbf{x}_m)$ between basis functions with a *kernel function* $K(\mathbf{x}_n, \mathbf{x}_m)$ that operates on a pair of instances in the original input space. This saves the work of having to construct $\phi(\cdot)$, whose dimension can be very high or infinite, and taking the dot product. All we need in practice is the kernel $K(\cdot, \cdot)$, not the basis functions $\phi(\cdot)$.
- The resulting, nonlinear discriminant (the *kernel SVM*) is

$$g(\mathbf{x}) = \mathbf{w}^T \phi(\mathbf{x}) = \sum_{n=1}^N \alpha_n y_n \phi(\mathbf{x}_n)^T \phi(\mathbf{x}) = \sum_{n=1}^N \alpha_n y_n K(\mathbf{x}_n, \mathbf{x}) = \sum_{n \in \text{SVs}} \alpha_n y_n K(\mathbf{x}_n, \mathbf{x})$$

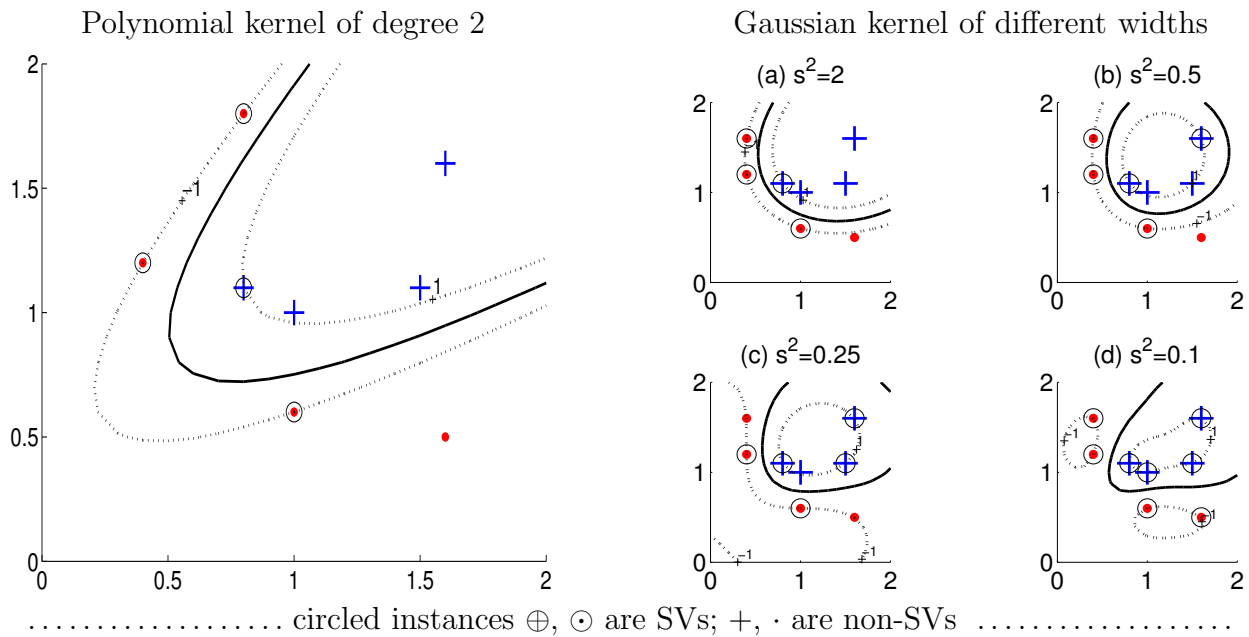
Again, we can use the kernel $K(\cdot, \cdot)$ directly and don't need the basis functions.

- Many other algorithms that depend on dot products of input instances have been kernelized (hence made nonlinear) in the same way: PCA, LDA, etc.
- The fundamental object in a kernel SVM, which determines the resulting discriminant, is the kernel, and the $N \times N$ *Gram matrix* of dot products it defines on a training set of N points: $\mathbf{K} = (K(\mathbf{x}_n, \mathbf{x}_m))_{nm} = (\phi(\mathbf{x}_n)^T \phi(\mathbf{x}_m))_{nm}$.

For this matrix and the dual QP to be well defined, the kernel $K(\cdot, \cdot)$ must be a symmetric positive definite function, i.e., it must satisfy $\mathbf{u}^T \mathbf{K} \mathbf{u} \geq 0 \quad \forall \mathbf{u} \in \mathbb{R}^N, \mathbf{u} \neq \mathbf{0}$.

- A kernel SVM can be seen as a basis function expansion (as in RBF networks), but the number of BF's is not set by the user (possibly by cross-validation); it is determined automatically during training, and each BF corresponds to one input instance that is a support vector.
- At test time, a kernel SVM works like a *template classifier*, by “comparing” the test pattern \mathbf{x} with the SVs (templates) by means of the kernel, and combining these comparisons into $g(\mathbf{x})$ to compute the final discriminant. Ex: for MNIST handwritten digits with SVs **0**, **4**...

$$g(\mathbf{x}) = \sum_{n \in \text{SVs}} \alpha_n y_n K(\mathbf{x}_n, \mathbf{x}) = \alpha_1 y_1 K(\mathbf{0}, \mathbf{x}) + \alpha_2 y_2 \underbrace{K(\mathbf{4}, \mathbf{x})}_{\text{how similar } \mathbf{x} \text{ is to } \mathbf{4}} + \dots$$



13.6 Vectorial kernels

- Intuitively, the kernel $K(\mathbf{x}, \mathbf{y})$ measures the similarity between two input points \mathbf{x}, \mathbf{y} . Different kernels may be useful in different applications.
- The most practically used kernels (and their hyperparameters) are, for $\mathbf{x}, \mathbf{y} \in \mathbb{R}^D$:
 - *Polynomial kernel* of degree q : $K(\mathbf{x}, \mathbf{y}) = (\mathbf{x}^T \mathbf{y} + 1)^q$.
 Ex. for $D = 2$ and $q = 2$: $K(\mathbf{x}, \mathbf{y}) = (x_1 y_1 + x_2 y_2 + 1)^2 = 1 + 2x_1 y_1 + 2x_2 y_2 + 2x_1 x_2 y_1 y_2 + x_1^2 y_1^2 + x_2^2 y_2^2$, which corresponds to the inner product of the basis function $\phi(\mathbf{x}) = (1, \sqrt{2}x_1, \sqrt{2}x_2, \sqrt{2}x_1 x_2, x_1^2, x_2^2)^T \in \mathbb{R}^6$.
 - For $q = 1$ we recover the *linear kernel* of the linear SVM.
 - *Gaussian or RBF kernel* of width σ : $K(\mathbf{x}, \mathbf{y}) = \exp(-\frac{1}{2} \|\mathbf{x} - \mathbf{y}\|^2 / \sigma^2)$.
 The Gaussian kernel implicitly defines a certain infinite-dimensional basis function $\phi(\mathbf{x})$.
 - *Sigmoidal (neural net) kernel* of hyperparameters $a, b \in \mathbb{R}$: $K(\mathbf{x}, \mathbf{y}) = \tanh(a\mathbf{x}^T \mathbf{y} + b)$.
- In addition to the kernel and its hyperparameters (q, σ , etc.), the user also has to select the C hyperparameter from the SVM. With kernel SVMs (nonlinear discriminants) this has the following effect:
 - $C \uparrow$ tries to avoid any error by discouraging any positive ξ_n , which leads to a wiggly boundary in the original input space with a small margin, so it can overfit.
 - $C \downarrow$ encourages a small $\|\mathbf{w}\|$ and hence a large margin, which leads to a smoother boundary in the original input space, but ignores most errors, so it can underfit.
- Likewise, a large q or small σ result in a more flexible boundary.
- The best values of the hyperparameters are found by cross-validation. Typically, we do a grid search, i.e., we select a set of values for each hyperparameter (say, C and σ), train an SVM for each combination of hyperparameter values (C, σ), and pick the SVM with the lowest error on the validation set.
- It is possible to define kernels for data where the instances are not represented in vector form.
 Ex: if \mathbf{x}, \mathbf{y} are documents using an unknown dictionary, we could define $K(\mathbf{x}, \mathbf{y})$ as the number of words they have in common.

13.9 Multiclass kernel machines

With $K > 2$ classes, there are two common ways to define the decision for a point $\mathbf{x} \in \mathbb{R}^D$:

- *One-vs-all*: we use K SVMs $g_k(\mathbf{x})$, $k = 1, \dots, K$.
 - *Training*: SVM g_k is trained to classify the training points of class C_k (label +1) vs the training points of all other classes (label -1).
 - *Testing*: choose C_k if $k = \arg \max_{i=1, \dots, K} g_i(\mathbf{x})$.
- *One-vs-one*: we use $K(K-1)/2$ SVMs $g_{ij}(\mathbf{x})$, $i, j = 1, \dots, K$, $i \neq j$.
 - *Training*: SVM g_{ij} is trained to classify the training points of class C_i (label +1) vs the training points of class C_j (label -1); training points from other classes are not used.
 - *Testing*: choose C_k if $k = \arg \max_{i=1, \dots, K} \sum_{j \neq i}^K g_{ij}(\mathbf{x})$. Also possible: for each class k , count the number of times $g_{kj}(\mathbf{x}) > 0$ for $j \neq k$, and pick the class with most votes.

13.10 Kernel machines for regression: support vector regression

- We are given a sample $\{(\mathbf{x}_n, y_n)\}_{n=1}^N$ with $\mathbf{x}_n \in \mathbb{R}^D$ and $y_n \in \mathbb{R}$.
The generalization to the case where \mathbf{y}_n has dimension $D' > 1$ is straightforward.
- We consider first linear regression: $f(\mathbf{x}) = \mathbf{w}^T \mathbf{x} + w_0$.
- Instead of the least-squares error $(y_n - f(\mathbf{x}_n))^2$, in *support vector regression* we use the ϵ -sensitive loss function:

$$[|y_n - f(\mathbf{x}_n)| - \epsilon]_+ = \max(0, |y_n - f(\mathbf{x}_n)| - \epsilon) = \begin{cases} 0, & \text{if } |y_n - f(\mathbf{x}_n)| < \epsilon \\ |y_n - f(\mathbf{x}_n)| - \epsilon, & \text{otherwise} \end{cases}$$

which means we tolerate errors up to $\epsilon > 0$ and that errors beyond ϵ have a linear penalty and not a quadratic one. This error function is more robust to noise and outliers; the estimated f will be less affected by them.

- As with the soft-margin hyperplane, we introduce slack variables ξ_n^+ , ξ_n^- to account for deviations (positive and negative) out of the ϵ -zone. We get the following, primal QP:

$$\min_{\mathbf{w} \in \mathbb{R}^D, w_0 \in \mathbb{R}, \xi^+, \xi^- \in \mathbb{R}^N} \frac{1}{2} \|\mathbf{w}\|^2 + C \sum_{n=1}^N (\xi_n^+ + \xi_n^-) \quad \text{s.t.} \quad \begin{cases} -\epsilon - \xi_n^+ \leq y_n - (\mathbf{w}^T \mathbf{x}_n + w_0) \leq \epsilon + \xi_n^+ \\ \xi_n^+, \xi_n^- \geq 0, \quad n = 1, \dots, N. \end{cases}$$

- This is a convex QP on $D + 1 + 2N$ variables and its unique minimizer $(\mathbf{w}, w_0, \xi^+, \xi^-)$ can be found with different QP algorithms. It is equivalent to solving the dual QP on $2N$ variables:

$$\min_{\alpha^+, \alpha^- \in \mathbb{R}^N} \frac{1}{2} \sum_{n,m} ((\alpha_n^+ - \alpha_n^-)(\alpha_m^+ - \alpha_m^-)(\mathbf{x}_n^T \mathbf{x}_m)) + \epsilon \sum_{n=1}^N (\alpha_n^+ + \alpha_n^-) - \sum_{n=1}^N y_n (\alpha_n^+ - \alpha_n^-)$$

$$\text{s.t.} \quad \sum_{n=1}^N (\alpha_n^+ - \alpha_n^-) = 0, \quad C \geq \alpha_n^+, \alpha_n^- \geq 0, \quad n = 1, \dots, N.$$

The optimal $\alpha \in \mathbb{R}^N$ has the property that $\alpha_n^+ = \alpha_n^- = 0$ only for those instances that lie within the ϵ -tube; these are the instances that are fitted with enough precision. The support vectors satisfy either $\alpha_n^+ > 0$ or $\alpha_n^- > 0$. As a result, for each training point $n = 1, \dots, N$ we can have:

- $\alpha_n^+ = \alpha_n^- = 0$: \mathbf{x}_n is fitted within the ϵ -tube.
- Either α_n^+ or α_n^- is in $(0, C)$: \mathbf{x}_n is fitted on the boundary of the ϵ -tube.
We use these instances to calculate w_0 , since they satisfy $y_n = \mathbf{w}^T \mathbf{x}_n + w_0 \pm \epsilon$ if $\alpha_n^\mp > 0$.
- $\alpha_n^+ = C$ or $\alpha_n^- = C$: \mathbf{x}_n is fitted outside the ϵ -tube.

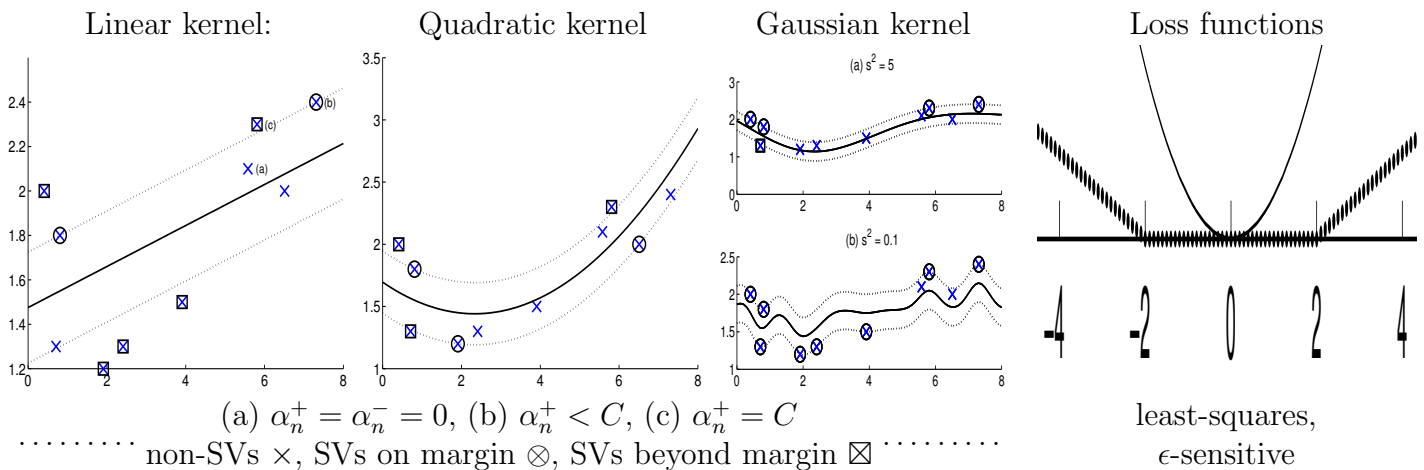
- From the optimal $\boldsymbol{\alpha}^+, \boldsymbol{\alpha}^- \in \mathbb{R}^N$, we recover the solution (\mathbf{w}, w_0) :

$$\mathbf{w} = \sum_{n=1}^N (\alpha_n^+ - \alpha_n^-) \mathbf{x}_n \quad w_0 = y_n \mp \epsilon - \mathbf{w}^T \mathbf{x}_n \text{ for any } n \text{ with } \alpha_n^\pm > 0.$$

- Hence, the fitted line can be written as a weighted sum of the support vectors:

$$f(\mathbf{x}) = \mathbf{w}^T \mathbf{x} + w_0 = \sum_{n=1}^N (\alpha_n^+ - \alpha_n^-) (\mathbf{x}_n^T \mathbf{x}) + w_0.$$

- Kernelization: we replace the dot products $\mathbf{x}_n^T \mathbf{x}_m$ (in the dual QP) or $\mathbf{x}_n^T \mathbf{x}$ (in the regression line) with $K(\mathbf{x}_n, \mathbf{x}_m)$ or $K(\mathbf{x}_n, \mathbf{x})$, respectively, where $K(\cdot, \cdot)$ is a kernel (polynomial, Gaussian, etc.). This results in a nonlinear regression function.



14 Graphical models

Joint distribution:	$p(X = x, Y = y).$
Conditioning (product rule):	$p(Y = y X = x) = \frac{p(X = x, Y = y)}{p(X = x)}.$
Marginalizing (sum rule):	$p(X = x) = \sum_y p(X = x, Y = y).$
Bayes' theorem: (inverse probability)	$p(X = x Y = y) = \frac{p(Y = y X = x) p(X = x)}{p(Y = y)}.$

Directed graphical models

- Also called *Bayesian networks*, *belief networks*, *probabilistic networks*.
- They represent the interaction between random variables visually, through a *directed acyclic graph* (DAG):
 - Each node represents one random variable X .
 - A directed arc $Y \rightarrow X$ indicates that Y has a *direct influence* on X (possibly but not necessarily indicating causality).
 - At node X , its probability distribution is conditional only on the nodes that point to it, and has tunable parameters (e.g. a table of probabilities for discrete variables, a mean/covariance for Gaussian variables).
Ex: if $X_1 \rightarrow X_3$ and $X_2 \rightarrow X_3$ over (X_1, \dots, X_d) , then we write $p(X_3 | \text{all variables}) = p(X_3 | X_1, \dots, X_d)$ as $p(X_3 | X_1, X_2)$.
- Types of independence of random variables:
 - X and Y are *independent* iff $p(X, Y) = p(X) p(Y)$.
Equivalently? $p(X|Y) = p(X)$ and $p(Y|X) = p(Y)$. Knowing Y tells me nothing about X and vice versa.
 - X and Y are *conditionally independent* given Z iff $p(X, Y | Z) = p(X | Z) p(Y | Z)$.
Equivalently? $p(X|Y, Z) = p(X|Z)$ and $p(Y|X, Z) = p(Y|Z)$.
- A graphical model represents a joint distribution by making conditional independence assumptions. In a graphical model, not all nodes are connected; typically, each node is connected to a few others only. The subgraphs containing these connections imply assumptions that the model makes about conditional independence among groups of variables. This allows inference over a set of variables to be decomposed into smaller sets of variables, each requiring local calculations.
- In this chapter, we focus mostly on discrete, binary random variables, and illustrate how to do inference in simple cases of graphical models.

Undirected graphical models

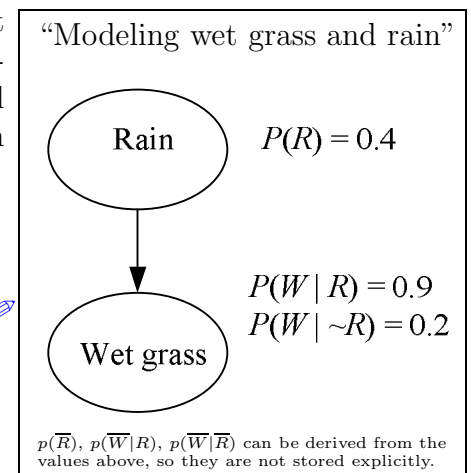
- Also called *Markov random fields* (MRFs).
- Convenient to represent symmetric interactions between random variables.

Example Notation: $p(R)$ means $p(R = 1)$, $p(\bar{W} | R)$ means $p(W = 0 | R = 1)$, etc.

A graphical model with completely specified conditional distributions at each node (in the example, tables giving $p(R)$ and $p(W | R)$ for all combinations of values of R and W), specifies the *joint distribution* over all variables (in the example, $p(R, W) = p(R) p(W | R)$), from which we can compute any specific distribution over groups of variables:

- $p(\bar{R}) = \text{⚡}$, $p(\bar{W} | R) = \text{⚡}$, $p(\bar{W} | \bar{R}) = \text{⚡}$
- Marginal: $p(W) = \sum_R p(R, W) = p(W | R) p(R) + p(W | \bar{R}) p(\bar{R}) = \text{⚡}$
- Bayes' rule: $p(R | W) = \frac{p(W | R) p(R)}{p(W)} = \text{⚡}$

Inverts the dependency to give a diagnosis.



14.2 Canonical cases for conditional independence

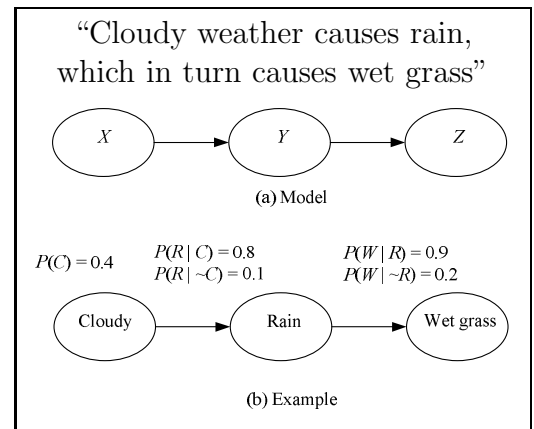
- By repeated application of the rule of conditional probability, we can write any joint distribution over d random variables without assumptions as follows (What graphical model corresponds to this?)

$$p(X_1, X_2, \dots, X_d) \stackrel{?}{=} p(X_1) p(X_2|X_1) p(X_3|X_1, X_2) \cdots p(X_d|X_1, X_2, \dots, X_{d-1}).$$

- In particular, consider 3 random variables X , Y and Z . Their joint distribution can always be written, without assumption, as $p(X, Y, Z) = p(X) p(Y|X) p(Z|X, Y)$. This can be simplified with assumptions given by conditional independencies, i.e., by removing arrows from the full graphical model $(X \rightarrow Y \rightarrow Z)$. There are 3 canonical cases.

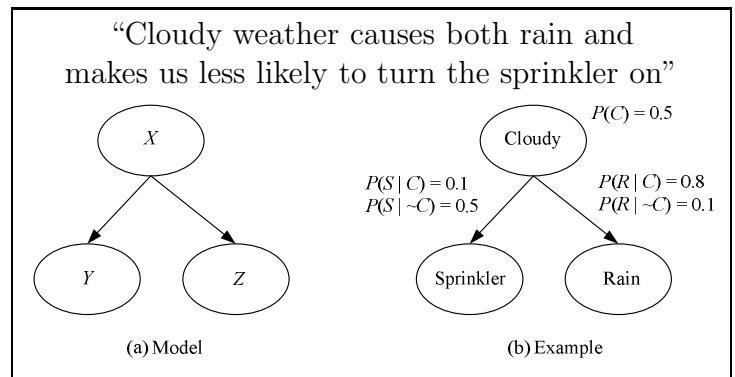
Case 1: head-to-tail connection

- $(X \rightarrow Y \rightarrow Z)$ means $p(X, Y, Z) = p(X) p(Y|X) p(Z|Y)$.
Note $p(Z|Y)$, not $p(Z|X, Y)$; we removed $X \rightarrow Z$.
- Knowing Y tells Z everything, knowing also X adds nothing, so X and Z are independent given Y : $p(Z|X, Y) = p(Z|Y)$
- Typically, X is the cause of Y and Y is the cause of Z .
- $p(R) = \text{blue}$, $p(W) = \text{blue}$, $p(W|C) = \text{blue}$, $p(C|W) = \text{blue}$




Case 2: tail-to-tail connection

- $(Z \leftarrow X \rightarrow Y)$ means $p(X, Y, Z) = p(X) p(Y|X) p(Z|X)$.
Note $p(Z|X)$, not $p(Z|X, Y)$; we removed $Y \rightarrow Z$.
- If we don't know X : Y and Z are dependent. If we know X : Y and Z become independent, $p(Y, Z|X) = p(Y|X) p(Z|X)$
- Typically, X is the cause of Y and Z .
- $p(C|R) = \text{blue}$, $p(R|S) = \text{blue}$, $p(R|\bar{S}) = \text{blue}$




Case 3: head-to-head connection

- $(X \rightarrow Z \leftarrow Y)$ means $p(X, Y, Z) = p(X) p(Y) p(Z|X, Y)$.
Note $p(Y)$, not $p(Y|X)$; we removed $X \rightarrow Y$.
- If we don't know Z : X and Y are independent, $p(X, Y) = p(X) p(Y)$. If we know Z : X and Y become dependent.
- Typically, Z has two independent causes X and Y .
- $p(W) = \text{blue}$, $p(W|S) = \text{blue}$, $p(S|W) = \text{blue}$, $p(S|R, W) < p(S|W)$ (explaining away), $p(S|\bar{R}, W) > p(S|W)$



A more general example

- Joint distribution without assumptions (requires 15 parameters):
 $p(C, S, R, W) = p(C) p(S|C) p(R|C, S) p(W|C, S, R)$.
- Joint distribution with assumptions (requires 9 parameters):
 $p(C, S, R, W) = p(C) p(S|C) p(R|C) p(W|S, R)$.
- $p(W|C) = \text{✎}$, $p(C|W) = \text{✎}$



A real-world example: QMR-DT, a Bayesian network for medical diagnosis, with binary variables for diseases (flu, asthma...) and symptoms (fatigue, dyspnea, wheezing, fever, chest-pain...), and directed arcs from each disease to the symptoms it causes.

Summary For d variables X_1, X_2, \dots, X_d (binary, but this generalizes to discrete or continuous variables):

- The general expression for a joint distribution without assumptions (requires 2^{d-1} parameters):

$$p(X_1, X_2, \dots, X_d) = p(X_1) p(X_2|X_1) p(X_3|X_1, X_2) \cdots p(X_d|X_1, X_2, \dots, X_{d-1})$$

simplifies under the conditional independence assumptions of the graphical model to:

$$p(X_1, X_2, \dots, X_d) = \prod_{i=1}^d p(X_i|\text{parents}(X_i))$$

which requires $2^{|\text{parents}(X_i)|}$ parameters Θ_i at node X_i : $p(X_i|\text{parents}(X_i)) = p(X_i|\text{parents}(X_i); \Theta_i)$.

- In turn, this simplifies the computations required for:
 - *Inference* (testing): answering questions of the form “ $p(X_1, X_7|X_3, X_5) = ?$ ”.
 - *Learning* (training): learning the parameters at each node (tables of probability values).
- In graphical models, we need not specify explicitly some variables as “inputs” and others as “outputs” as in a supervised problem (e.g. classification). Having the trained graphical model (i.e., with known values for the conditional distribution table at each node), we can set the values of any subset of the random variables $\mathcal{S}_{\text{outputs}} \subset \{X_1, \dots, X_d\}$ (e.g. based on observed data) and do *inference* over any other subset $\mathcal{S}_{\text{inputs}} \subset \{X_1, \dots, X_d\}$ (unobserved variables), i.e., compute $p(\mathcal{S}_{\text{outputs}}|\mathcal{S}_{\text{inputs}})$, where $\mathcal{S}_{\text{outputs}} \cap \mathcal{S}_{\text{inputs}} = \emptyset$. The graphical model is a “probabilistic database”, a machine that can answer queries regarding the values of random variables.
- We can also have *hidden variables*, which are never observed, but which help in defining the dependency structure.

Ex: basket analysis. We want to find dependencies between items sold for “baby food”, “diapers” and “milk” (a customer buying one of these is likely to buy the other two). Instead of putting arcs among these, we may designate a hidden node “baby at home” as the hidden cause of the consumption of these 3 items.

- Training, i.e., estimating the parameters of the graphical model given a dataset, is typically done by maximum likelihood.
- The computational complexity of training and inference depends on properties of the graphical model:
 - If the graph is a tree: exact inference takes polynomial time (junction-tree algorithm).
 - Otherwise, exact inference is NP-complete and becomes intractable for large models. In that case one uses approximate algorithms (variational methods, Markov chain Monte Carlo, etc.).


14.3 Generative models

- *Generative model*: a graphical model that represents the process we believe created the data.
- Convenient to design specific graphical models for supervised and unsupervised problems in machine learning (classification, regression, clustering, etc.).
- Ex: classification with inputs $\mathbf{x} = (x_1, \dots, x_D)$ and class labels C :

- Generative model: $C \rightarrow \mathbf{x}$ means $p(C, \mathbf{x}) = p(C)p(\mathbf{x}|C)$, or “first pick a class C by sampling from $p(C)$, then (given C) pick an instance \mathbf{x} by sampling from $p(\mathbf{x}|C)$ ”. This allows us to sample pairs (C, \mathbf{x}) if we know the model for $p(\mathbf{x}|C)$ (the class-conditional probability), and the model parameters ($p(C)$ and $p(\mathbf{x}|C)$).

Ex: $p(C) = \pi_C \in (0, 1)$ and $p(\mathbf{x}|C) = \mathcal{N}(\boldsymbol{\mu}_C, \boldsymbol{\Sigma}_C)$, for continuous variables.

- Bayes’ rule inverts the generative process to classify an instance \mathbf{x} : $p(C|\mathbf{x}) = \frac{p(\mathbf{x}|C)p(C)}{p(\mathbf{x})}$.
- If, in addition, we assume that $\mathbf{x} = (x_1, \dots, x_D)$ are conditionally independent given C (i.e., we ignore possible correlations between features for instances within the same class), then we achieve a simpler model, the *naive Bayes classifier*: $p(\mathbf{x}|C) = p(x_1|C) \cdots p(x_D|C)$.



14.6 Undirected graphical models (Markov random fields)

- A different way to express graphically a probability distribution over d random variables, more convenient than directed graphical models if the influences between variables are symmetric.

Ex: pixels on an image. Two neighboring pixels tend to have the same color, and the correlation goes both ways.

- The joint distribution is not defined using parents (which imply a directed graph), but using:
 - *Cliques* (clique = set of nodes such that there is a link between every pair of nodes in the clique).
 - *Potential functions* $\psi_C(X_C)$, where X_C is the set of variables in clique C .

$$X = (X_1, \dots, X_d): \quad p(X) = \frac{1}{Z} \prod_{\text{all cliques } C} \psi_C(X_C) \quad Z = \sum_X \prod_{\text{all cliques } C} \psi_C(X_C).$$

The *partition function* Z is a normalization constant that ensures that $\sum_X p(X) = 1$.

- Ex: consider an image X with d black or white pixels ($X_i \in \{0, 1\}$, so $X = (X_1, \dots, X_d)$ can take 2^d different values). Connecting each pixel to its 4 neighboring pixels defines two types of cliques: single pixels X_i , and pairs of neighboring pixels (X_i, X_j) . Then:

$$p(X) = \frac{1}{Z} \left(\prod_{i=1}^d \psi_1(X_i) \right) \left(\prod_{i \sim j} \psi_2(X_i, X_j) \right) \quad \psi_1(X_i) = e^{-|Y_i - X_i|} \quad \psi_2(X_i, X_j) = e^{-\alpha |X_i - X_j|}$$

where Y_1, \dots, Y_d are the pixel values for an observed, noisy image Y and $\alpha > 0$. Then, $p(X)$ is high if X is close to Y (single-pixel cliques), and if X is smooth (neighboring pixels tend to have the same value), as controlled by α . We can compute a denoised image as $\arg \max_X p(X)$.

In statistical mechanics, the *Ising model* for magnetic materials is similar to an MRF: pixel = atom with spin $\in \{\uparrow, \downarrow\}$.

15 Discrete Markov models and hidden Markov models

- Up to now, we have regarded data points $\mathbf{x}_1, \dots, \mathbf{x}_N$ as *identically independently distributed* (*iid*): each \mathbf{x}_n is an independent sample from a (possibly unknown) distribution $p(\mathbf{x}; \Theta)$ with parameters Θ . Hence, the log-likelihood is simply a sum of the individual log-likelihood for each data point: $\log P(\mathbf{x}_1, \dots, \mathbf{x}_N) = \sum_{n=1}^N \log p(\mathbf{x}_n; \Theta)$. As a graphical model: $\bigcirc(\mathbf{x}_1) \bigcirc(\mathbf{x}_2) \cdots \bigcirc(\mathbf{x}_N)$.
- Now, we relax that assumption: $\mathbf{x}_1, \dots, \mathbf{x}_N$ depend on each other.
- Ex: temporal or spatial series (discrete or continuous):
 - Temporal, continuous: tides, trajectory of a moving body, etc.
 - Temporal, discrete:
 - * Letter or word sequences in English: $t \rightarrow 'h'$ more likely than $'x'$. Because of grammar or phonological rules, only certain sequences are allowed, and some are more likely than others.
 - * Speech: sequence of sounds, each corresponding to a phoneme. Physical constraints of the vocal tract articulators (tongue, lips, etc.) mean their motion is smooth, and so is the acoustic speech they produce.
 - Spatial, discrete: base pairs in a DNA sequence.

15.2 Discrete Markov processes (Markov chains)

- Consider a system that at any time is in one of a set of N distinct states S_1, S_2, \dots, S_N . Write the state at time $t = 1, 2, \dots$ as a discrete random variable $q_t \in \{S_1, S_2, \dots, S_N\}$.
- In general, we can always write the probability of a sequence of states of length T as:

$$p(q_1, q_2, \dots, q_T) = p(q_1) p(q_2|q_1) p(q_3|q_1, q_2) \cdots p(q_T|q_1, q_2, \dots, q_{T-1}),$$

that is, the state at time $t + 1$ depends on all the previous states since $t = 1$:

$$p(q_{t+1} = S_j \mid q_t = S_i, q_{t-1} = S_k, \dots) \quad \leftarrow \text{requires a table of } N^T \text{ parameters.}$$

- *First-order Markov model*: assumes the state at time $t + 1$ depends only on the state at time t (it is conditionally independent of all other times given time t , or “the future is independent of the past given the present”):

$$p(q_{t+1} = S_j \mid q_t = S_i, q_{t-1} = S_k, \dots) = p(q_{t+1} = S_j \mid q_t = S_i).$$

Markov model of order k : the state at time $t + 1$ depends on the state at the previous k times $t, t - 1, \dots, t - k + 1$.

- *Homogeneous Markov model*: assumes the *transition probability* from S_i to S_j (for any pair of states S_i, S_j) is independent of the time (so going from S_i to S_j has the same probability no matter when it happens):

$$a_{ij} \equiv p(q_{t+1} = S_j \mid q_t = S_i) \quad \sum_{j=1}^N a_{ij} = 1, \quad i = 1, \dots, N, \quad a_{ij} \in [0, 1], \quad i, j = 1, \dots, N.$$


- *Transition probability matrix* $\mathbf{A} = (a_{ij})$ of $N \times N$: contains nonnegative elements, rows sum 1.
- The first state in a sequence has its own *initial distribution* $\boldsymbol{\pi}$ (a vector of $N \times 1$):

$$\pi_i \equiv p(q_1 = S_i), \quad \sum_{i=1}^N \pi_i = 1, \quad \pi_i \in [0, 1], \quad i = 1, \dots, N.$$

- The probability of a given sequence q_1, q_2, \dots, q_T of length T is:

$$p(q_1, \dots, q_T; \mathbf{A}, \boldsymbol{\pi}) = p(q_1) p(q_2|q_1) \cdots p(q_T|q_{T-1}) = \pi_{q_1} a_{q_1 q_2} \cdots a_{q_{T-1} q_T} = \pi_{q_1} \left(\prod_{t=2}^T a_{q_{t-1} q_t} \right).$$

- A discrete Markov process can be seen as a *stochastic automaton* with states S_1, \dots, S_N , transition matrix $\mathbf{A} = (a_{ij})$ and initial distribution $\boldsymbol{\pi} = (\pi_i)$.



- The sequence of states (each a random variable) can be seen as a *graphical model*: $(q_1) \rightarrow (q_2) \rightarrow \dots \rightarrow (q_T)$.

- *Sampling*: given \mathbf{A} and $\boldsymbol{\pi}$, we can generate a sequence of states of length T as follows: sample q_1 from $\boldsymbol{\pi}$; sample q_2 from $\mathbf{A}_{q_1 \bullet}$; ...; sample q_T from $\mathbf{A}_{q_{T-1} \bullet}$.
- *Training*: given K observed sequences of length T , we want to estimate the parameters $(\mathbf{A}, \boldsymbol{\pi})$. This can be done by maximum likelihood:

$$\max_{\boldsymbol{\pi}, \mathbf{A}} \sum_{k=1}^K \log P(\text{sequence } k) \quad \text{where} \quad P(\text{sequence } k) = \pi_{q_{k1}} \left(\prod_{t=2}^T a_{q_{k,t-1} q_{kt}} \right).$$

The solution is:

$$\pi_i = \frac{\# \text{ sequences starting with } S_i}{\# \text{ sequences}} \quad a_{ij} = \frac{\# \text{ transitions from } S_i \text{ to } S_j}{\# \text{ transitions from } S_i}.$$


15.3 Hidden Markov models (HMMs)

- At each time $t = 1, 2, \dots$ there are two random variables:
 - The state q_t , which is unobserved (a hidden variable). It depends only on the previous time state, q_{t-1} : *transition probability* $p(q_t|q_{t-1})$.
 - The observation O_t , which we do observe. It depends only on the current state, q_t : *observation (or emission) probability* $p(O_t|q_t)$. The observation O_t can be:
 - * Discrete, i.e., symbols from an alphabet (= a set of symbols). Then, $p(O_t|q_t)$ is a matrix \mathbf{B} of probabilities for each pair (state, symbol).
 - * Continuous and possibly multidimensional. Then, $p(O_t|q_t)$ is a continuous distribution, e.g. a Gaussian $\mathcal{N}(\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i)$ for state S_i .

Hence, an observed sequence O_1, \dots, O_T results from two sources of randomness: randomly moving from state to state, and randomly emitting an observation at each state.

- Again, we assume these probabilities don't change over time (homogenous HMM).
- The first state in a sequence has its own *initial distribution* $p(q_1)$.
- Altogether, the HMM is characterized by 3 groups of parameters: the transition matrix \mathbf{A} , the observation matrix \mathbf{B} (assuming discrete observations), and the initial distribution $\boldsymbol{\pi}$.

- An HMM for an observation sequence O_1, \dots, O_T (and unobserved state sequence q_1, \dots, q_T) is a graphical model with one node per random variable (O_t or q_t). Each node contains the parameters needed to generate its random variable:




node q_t for $t > 1$ contains \mathbf{A} (necessary to compute $p(q_t|q_{t-1})$), node O_t contains \mathbf{B} (necessary to compute $p(O_t|q_t)$), and node q_1 contains $\boldsymbol{\pi}$ (necessary to compute $p(q_1)$).

What can we do with an HMM?

- *Sampling*: given $(\mathbf{A}, \mathbf{B}, \boldsymbol{\pi})$ we can generate a sequence of observations as follows: sample q_1 from $\boldsymbol{\pi}$ and O_1 from $\mathbf{B}_{q_1\bullet}$; sample q_2 from $\mathbf{A}_{q_1\bullet}$ and O_2 from $\mathbf{B}_{q_2\bullet}$; ...; sample q_T from $\mathbf{A}_{q_{T-1}\bullet}$ and O_T from $\mathbf{B}_{q_T\bullet}$.
- *Evaluating* the probability $p(O; \mathbf{A}, \mathbf{B}, \boldsymbol{\pi})$ of a given observation sequence $O = (O_1, \dots, O_T)$. The joint probability of a given sequence of both observations and states (if they were observed) is:

$$p(O, Q; \mathbf{A}, \mathbf{B}, \boldsymbol{\pi}) = p(q_1) \left(\prod_{t=2}^T p(q_t|q_{t-1}) \right) \left(\prod_{t=1}^T p(O_t|q_t) \right) = \pi_{q_1} b_{q_1 O_1} a_{q_1 q_2} \cdots a_{q_{T-1} q_T} b_{q_T O_T}.$$

We can compute $p(O|\mathbf{A}, \mathbf{B}, \boldsymbol{\pi})$ by marginalizing this joint distribution over all possible state sequences Q : $p(O|\mathbf{A}, \mathbf{B}, \boldsymbol{\pi}) = \sum_Q p(O, Q; \mathbf{A}, \mathbf{B}, \boldsymbol{\pi})$. Although there are N^T such sequences, this marginalization can be computed exactly in $\mathcal{O}(N^2T)$ using the *forward-backward algorithm* (a dynamic programming algorithm).



- *Decoding*: given a sequence of observations $O = (O_1, \dots, O_T)$, find the sequence of states $Q = (q_1, \dots, q_T)$ that has highest probability $p(Q|O; \mathbf{A}, \mathbf{B}, \boldsymbol{\pi})$ (many state sequences are possible, but some are likelier than others to have generated the observation sequence O). Again, this can be done exactly in $\mathcal{O}(N^2T)$ using the *Viterbi algorithm* (a dynamic programming algorithm).
Speech recognition: states = phonemes, observations = acoustic measurements (e.g. mel-frequency cepstral coefficients (MFCCs)).
- *Training*: given a training set of observation sequences, learn the HMM parameters $(\mathbf{A}, \mathbf{B}, \boldsymbol{\pi})$ that maximize the probability of generating those sequences. This can be solved with a form of the EM algorithm (the *Baum-Welch algorithm*). E step: compute “ $p(Q|O; \mathbf{A}, \mathbf{B}, \boldsymbol{\pi})$ ” for the current parameters $(\mathbf{A}, \mathbf{B}, \boldsymbol{\pi})$. M step: reestimate the parameters $(\mathbf{A}, \mathbf{B}, \boldsymbol{\pi})$, as in a discrete Markov process.

In an HMM, the state variables follow a first-order Markov process: $p(q_t|q_1, \dots, q_{t-1}) = p(q_t|q_{t-1})$. But the observed variables do not: $p(O_t|O_{t-1}, \dots, O_1) \neq p(O_t|O_{t-1})$. In fact, O_t depends (implicitly) on all previous observations. An HMM achieves this long-range dependency with fewer parameters than using a Markov process of order k . The latter needs N^{k+1} transition probabilities, and a huge training set where every possible sequence of $k + 1$ symbols (bigrams, trigrams, ..., $k + 1$ -grams) should be observed.

Continuous states and observations: tracking (dynamical models)

- Both the states and the observations are continuous. We want to predict the state at time $t + 1$ given the state q_t and the observation O_t at time t .
- Many applications in control, computer vision, robotics, etc. Ex: predicting the trajectory of a guided missile, mobile robot, pedestrian, etc. State $q_t = (x_t, y_t, z_t)$ spatial coordinates of robot at time t , observation O_t = sensor information at time t (camera image, depth sensor, etc.).
- A *Kalman filter* (or *linear dynamical system*) is like an HMM, but:
 - Transition prob. $p(q_{t+1}|q_t)$: q_{t+1} is a linear function of q_t plus zero-mean Gaussian noise.
 - Emission prob. $p(O_t|q_t)$: O_t is another linear function of q_t plus zero-mean Gaussian noise.

Training can also be done with an EM algorithm.

- Extensions to nonlinear functions and/or nongaussian distributions: *extended Kalman filter*, *particle filters*.

17 Combining multiple learners

- By suitably combining multiple learners the accuracy can be improved (but it need not to).
 - How to generate *base learners* that complement each other?
 - How to *combine* their outputs for maximum accuracy?
- Ex: train an ensemble of L decision trees on L different subsets of the training set and define the ensemble output for a test instance as the majority vote (for classification) or the average (for regression) of the L trees.
- Ensembles of decision trees (random forest, boosted decision trees) are practically among the most accurate models in machine learning.
- Disadvantages:
 - An ensemble of learners is computationally more costly in time and space than a single learner, both at training and test time.
 - A decision tree is interpretable (as rules), but an ensemble of trees is hard to interpret.

17.1 Generating diverse learners

- If the learners behave identically, i.e., the outputs of the learners are the same for any given input, their combination will be identical to any individual learner. The accuracy doesn't improve and the computation is slower.
 - ☞ *Diversity*: we need learners whose decisions differ and complement each other.
- If each learner is extremely accurate, or extremely inaccurate, the ensemble will barely improve the individual learners (if at all).
 - ☞ *Accuracy*: the learners have to be sufficiently accurate, but not very accurate.
- Good ensembles need a careful interplay of accuracy and diversity. Having somewhat inaccurate base learners can be compensated by making them diverse and independent from each other.
- Although one can use any kind of models to construct ensembles, practically it is best to use base learners that are simple and unstable.

Mechanisms to generate diversity

- *Different models*: each model (linear, neural net, decision tree. . .) makes different assumptions about the data and lead to different classifiers.
- *Different hyperparameters*: within the same model (polynomials, neural nets, RBF networks. . .), using different hyperparameters leads to different trained models. Ex: number of hidden units in multilayer perceptrons or of basis functions in RBF networks, k in k -nearest-neighbor classifiers, error threshold in decision trees, etc.
- *Different optimization algorithm or initialization*: for nonconvex problems (e.g. neural nets), each local optimum of the objective function corresponds to a different trained model. The local optimum found depends on the optimization algorithm used (gradient descent, alternating optimization, etc.) and on the initialization given to it.
- *Different features*: each learner can use a different (possibly random) subset of features from the whole feature vector. This also makes each learner faster, since it uses fewer features.

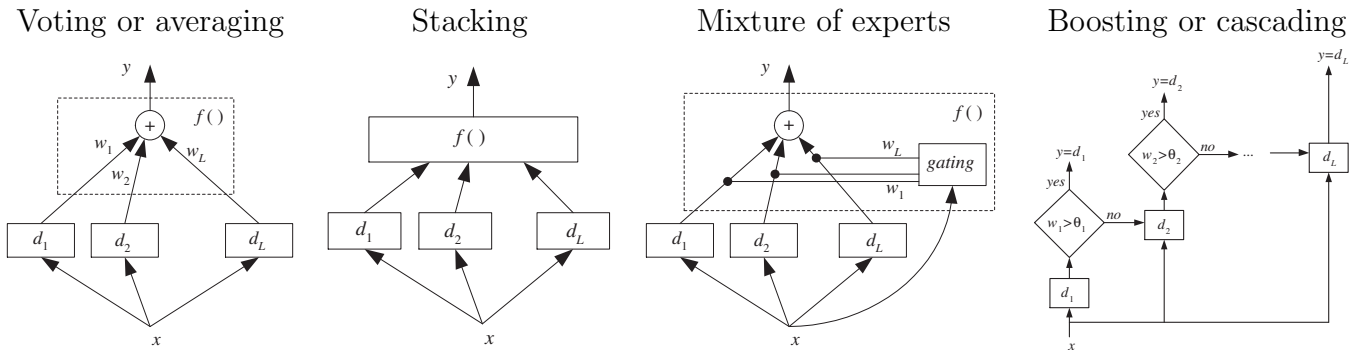
- *Different training sets*: each learner is trained on a different subset of the data. We can do this:
 - In parallel, by drawing independent random subsets from the training set, as in *bagging*.
 - Sequentially, by giving more emphasis to instances on which the preceding learners are not accurate, as in *boosting* or *cascading*.
 - By making the subsets local, e.g. obtained by clustering.
 - By defining the main task in terms of several subtasks to be implemented by the learners, as in *error-correcting output codes*.

17.2 Model combination schemes

Given L trained models (learners), their outputs $\mathbf{y}_1, \dots, \mathbf{y}_L$ for an input \mathbf{x} can be combined:

- *In parallel or multiexpert*:
 - *Global approach (learner fusion)*: all learners generate an output and all these outputs are combined. Ex: a fixed combination (*voting*, *averaging*) or a learned one (*stacking*).
 - *Local approach (learner selection)*: a “gating” model selects one learner as responsible to generate the final output. Ex: *mixture of experts*.
- *In sequence or multistage*: the learners are sorted (usually in increasing complexity) and we apply them in sequence to the input until one of them is confident. Ex: *boosting*, *cascading*.

In the case of K -class classification, each learner may have K outputs (for the discriminant of each class, or from a softmax). We can have each learner output a single class (the one with largest discriminant or softmax), or have the combination use all K outputs of all L learners.



17.3 Voting and averaging

- For discrete outputs (classification):
 - *Majority vote*: the class with most votes wins.
 - *Weighted vote* with the posterior prob. $p_l(C_1|\mathbf{x}), \dots, p_l(C_K|\mathbf{x})$ from each learner $l = 1, \dots, L$.
- For continuous outputs (regression):
 - *Average* $\mathbf{y} = \frac{1}{L} \sum_{l=1}^L \mathbf{y}_l$. Possibly weighted: $\mathbf{y} = \frac{1}{L} \sum_{l=1}^L w_l \mathbf{y}_l$ with $\sum_{l=1}^L w_l = 1$ and $w_1, \dots, w_L \in (0, 1)$.
 - *Median*: more robust to outlying outputs.
- *Bayesian model combination*: in classification, the posterior prob. marginalized over all models is $p(C_k|\mathbf{x}) = \sum_{\text{all models } \mathcal{M}_i} p(C_k|\mathbf{x}, \mathcal{M}_i) p(\mathcal{M}_i)$, which can be seen as weighted averaging using as weights the model prior probabilities. Simple voting corresponds to a uniform prior (all models equally likely).

Bias and variance Why (and when) does diversity help?

- Consider L *independent* binary classifiers with success probability $> \frac{1}{2}$ (i.e., better than random guessing) combined by taking a majority vote. One can prove the accuracy increases with L . Not necessarily true if they are not independent (e.g. if the L classifiers are equal the accuracy will not change).
- Consider L *iid* random variables y_1, \dots, y_L with expected value $E\{y_l\} = \mu$ and variance $\text{var}\{y_l\} = \sigma^2$. Then, the average $y = \frac{1}{L} \sum_{l=1}^L y_l$ has the following moments:

$$E\{y\} = E\left\{\frac{1}{L} \sum_{l=1}^L y_l\right\} = \frac{1}{L} \sum_{l=1}^L E\{y_l\} = \mu$$
$$\text{var}\{y\} = \text{var}\left\{\frac{1}{L} \sum_{l=1}^L y_l\right\} = \frac{1}{L^2} \text{var}\left\{\sum_{l=1}^L y_l\right\} = \frac{1}{L^2} \sum_{l=1}^L \text{var}\{y_l\} = \frac{1}{L} \sigma^2.$$

So the expected value (hence the bias) doesn't change, but the variance (hence the mean squared error) decreases as L increases. If y_1, \dots, y_L are identically but *not independently* distributed:

$$\text{var}\{y\} = \frac{1}{L^2} \text{var}\left\{\sum_{l=1}^L y_l\right\} = \frac{1}{L^2} \left(\sum_{l=1}^L \text{var}\{y_l\} + 2 \sum_{i,j=1, i \neq j}^L \text{cov}\{y_i, y_j\} \right) = \frac{1}{L} \sigma^2 + \frac{2}{L^2} \sum_{i,j=1, i \neq j}^L \sigma_{ij}.$$

So, if the learners are positively correlated ($\sigma_{ij} > 0$), the variance (and error) is larger than if they are independent. Hence, a well-designed ensemble combination will aim at reducing, if not completely eliminating, positive correlation between learners.

Further decrease in variance is possible with negatively correlated learners. However, it is impossible to have many learners that are both accurate and negatively correlated.

- *Ensemble averaging and noise*: if we view each learner as a random noise function added to the true discriminant/regression function, and if these noise functions are independent with zero mean, then the ensemble average is smoother than each ensemble member.

Voting or averaging over models with low bias and high variance produces an ensemble with low bias but lower variance.

fig. 4.5

17.5 Bagging

- *Bootstrap*: given a training set $\mathcal{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ containing N samples, the *bootstrap* generates L subsets $\mathcal{X}_1, \dots, \mathcal{X}_L$ of \mathcal{X} , each of size N , as follows: subset \mathcal{X}_l is obtained by sampling at random *with replacement* N points from \mathcal{X} .
 - This means that some points \mathbf{x}_n will appear repeated multiple times in \mathcal{X}_l and some other points \mathbf{x}_n will not appear in \mathcal{X}_l .
 - The L subsets are partly different from each other. On average, each subset contains $\approx 63\%$ of the training set. Proof: the probability we don't pick instance \mathbf{x}_n after N draws is $(1 - \frac{1}{N})^N \approx e^{-1} \approx 0.37$.
- *Stable vs unstable* learning algorithms:
 - A learning algorithm is *unstable* if small changes in the training set cause a large difference in the trained model, i.e., the training algorithm has large variance.
 - Running a stable algorithm on resampled versions of the training set lead to learners with high positive correlation, so little diversity. Using an unstable algorithm reduces this correlation and thus the ensemble has lower variance.
 - Ex. of unstable algorithms: decision trees (the bigger the more unstable), neural nets.
 - Ex. of stable algorithms: linear classifiers or regressors, nearest-neighbor classifier.

- *Bagging (bootstrap aggregating)*: applicable to classification and regression.
 - We generate L (partly different) subsets of the training set with the bootstrap.
 - We train L learners, each on a different subset, using an unstable learning algorithm. Since the training sets are partly different, the resulting learners are diverse.
 - The ensemble output is defined as the vote or average (or median) of the learners' outputs.
- *Random forest*: a variation of bagging using *decorrelated* decision trees as base learners.
 - As in bagging, each tree is trained on a bootstrap sample of the training set, and the ensemble output is defined as the vote or average (or median) of the learners' outputs.
 - In addition, the l th tree is constructed with a randomized algorithm, independently of the other trees. Typically, we do random feature selection at each node of the tree: we select the feature to split not among all possible features, but among a random subset of them. It is recommended to grow a full tree (no pruning), because it has lower bias and higher variance; and to use as size of the random subset $1 + \lceil \log_2 D \rceil$ (where $D = \dim \mathbf{x}$ is the dimensionality of the instances).
- Random forests are among the best classifiers in practice, and simpler to train than boosting.

17.6 Boosting

- Bagging generates complementary learners through random sampling and unstable learning algorithms. *Boosting actively generates complementary learners by training the next learner on the mistakes of the previous learners.*
- *Weak learner*: a learner that has probability of error $< \frac{1}{2}$ (i.e., better than random guessing on binary classification). Ex: decision trees, *decision stumps* (tree grown to only 1 or 2 levels). *Strong learner*: a learner that can have arbitrarily small probability of error. Ex: neural net.
- There are many versions of boosting, we focus on *AdaBoost.M1*, for classification (it can be applied to regression with some modifications):
 - It combines L weak learners. They have high bias, but the decrease in variance in the ensemble compensates for that.
 - Each learner $l = 1, \dots, L$ is trained on the entire training set $\mathcal{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$, but each point \mathbf{x}_n has an associated probability $p_n^{(l)}$ indicating how “important” it is for that learner. The training algorithm must be able to use these probabilities $\mathbf{p}^{(l)}$ together with the training set \mathcal{X} . Otherwise, this can be simulated by sampling a training set of size N from \mathcal{X} according to $\mathbf{p}^{(l)}$.
 - The first learner uses $p_n^1 = \frac{1}{N}$ (all points equally important).
 - Assume that after training learner l on $(\mathcal{X}, \mathbf{p}^{(l)})$ its error rate is ϵ_l . If $\epsilon_l \geq \frac{1}{2}$ we stop, since the learner is not weak. Otherwise, we update the probabilities as follows: for each point \mathbf{x}_n , $n = 1, \dots, N$:

$$p_n^{(l+1)} = \begin{cases} \beta_l p_n^{(l)} & \text{if learner } l \text{ correctly classifies } \mathbf{x}_n \\ p_n^{(l)} & \text{otherwise} \end{cases} \quad \text{where } \beta_l = \frac{\epsilon_l}{1 - \epsilon_l} \in [0, \frac{1}{2}).$$

Then we renormalize the probabilities so they sum 1: $p_n^{(l+1)} = p_n^{(l+1)} / \sum_{n=1}^N p_n^{(l+1)}$.

- This decreases the probability of a point if it is correctly classified, so the next learner focuses on the misclassified points. That is why learners are chosen to be weak. If they are strong (= very accurate), the next learner's training set will emphasize a few points, many of which could be very noisy or outliers.
- After training, the ensemble output for a given instance is a weighted vote, where the weight of learner l is $w_l \propto \log(1/\beta_l)$, so the weaker learners have a lower weight.

Other variations of boosting make the overall decision by applying learners in sequence, as in cascading.

fig. 17.2

- Boosting usually achieves very good classification performance, although it does depend on the dataset and the type of learner used (it should be weak but not too weak). It is sensitive to noise and outliers.

A very successful application in computer vision: the Viola-Jones face detector. This is a cascade of AdaBoost ensembles of decision stumps, each trained on a large set of Haar features. It is robust and (with clever image-based operations) very fast for test images.

Cascading is similar to boosting, but the learners are trained sequentially. For example, in boosting the next learner could be trained on the residual error of the previous learner. Another way:

- When applied to an instance \mathbf{x} , each learner gives an output (e.g. class label) and a confidence (which can be defined as the largest posterior probability $p(C_k|\mathbf{x})$).
- Learner l is trained on instances for which the previous learner is not confident enough.
- When applying the ensemble to a text instance, we apply the learners in sequence until one is confident enough. We use learner l only if the previous learners $1, \dots, l-1$ are not confident enough on their outputs.
- The goal is to order the learners in increasing complexity, so the early learners are simple and classify the easy instances quickly.

17.4 Error-correcting output codes (ECOC)

- An ensemble learning method for K -class classification.
- Idea: instead of solving the main classification task directly with one classifier, which may be difficult, create simpler classification subtasks which can be combined to get the main classifier.
- *Base learners*: L binary classifiers with outputs in $\{-1, +1\}$.
- *Code matrix* \mathbf{W} of $K \times L$ with elements in $\{-1, +1\}$: if $w_{kl} = -1$ (+1) then class k should be on the negative (positive) side of learner l . Hence, each learner tries to classify one subset of classes vs the rest (it partitions the K classes into two groups). Ex.:

$$\text{one-vs-all: } \begin{pmatrix} +1 & -1 & -1 & -1 \\ -1 & +1 & -1 & -1 \\ -1 & -1 & +1 & -1 \\ -1 & -1 & -1 & +1 \end{pmatrix} \quad \text{one-vs-one: } \begin{pmatrix} +1 & +1 & +1 & 0 & 0 & 0 \\ -1 & 0 & 0 & +1 & +1 & 0 \\ 0 & -1 & 0 & -1 & 0 & +1 \\ 0 & 0 & -1 & 0 & -1 & -1 \end{pmatrix} \quad \begin{matrix} \text{all possible} \\ \text{learners:} \end{matrix} \begin{pmatrix} -1 & -1 & -1 & -1 & -1 & -1 \\ -1 & -1 & -1 & +1 & +1 & +1 \\ -1 & +1 & +1 & -1 & -1 & +1 \\ +1 & -1 & +1 & -1 & +1 & +1 \end{pmatrix}$$

- Particular cases of the code matrix:
 - One-vs-all: $L = K$ and \mathbf{W} has +1 along the diagonal and -1 in elsewhere.
 - One-vs-one: $L = K(K-1)/2$ and each column of \mathbf{W} has one -1 , one $+1$ and 0 elsewhere (“0” means don’t care).

An ECOC code matrix has L between K (one-vs-all, fewest learners) and $2^{K-1} - 1$ (all possible learners). Because negating a column gives the same learner; an a column of all -1 s (or all $+1$ s) is useless.

- The code matrix allows us to define a K -class classification problem in terms of several binary classification problems. We can use any binary classifier for the latter (decision trees, etc.).
- To classify a test input \mathbf{x} , we apply the L classifiers to it and obtain a row vector \mathbf{y} with L entries in $\{-1, +1\}$. Ideally, if the classifiers were perfect, this would equal the row of \mathbf{W} corresponding to \mathbf{x} ’s class, but in practice some classifiers will classify \mathbf{x} incorrectly. Then, we find the row $1 \leq k \leq K$ in \mathbf{W} that is closest to this \mathbf{y} in Hamming distance, and output k as label (as in error-correcting codes). Equivalently⁵, we can compute a weighted vote $\sum_{l=1}^L w_{kl}y_l$ for class k (where the weights w_{kl} are the elements of \mathbf{W}) and then pick the class with most votes.

- The point of ECOC is to introduce robustness against errors of the learners by making their codewords be farther from each other in Hamming distance (number of mismatching bits). This requires sufficiently many learners, and introduces redundancy.
- Given a value of L (the number of classifiers, with $L > K$) selected by the user, we generate \mathbf{W} so its rows are as different as possible in Hamming distance (redundant codes so protection against errors), and its columns are as different as possible (so the learners are diverse).
- An ECOC can also be seen as an ensemble of classifiers where, to obtain each classifier, we manipulate the output targets (rather than the input features or the training set, which in ECOC are equal to the original training set for each classifier).
- Practically. the main benefit of ECOC seems to be in variance reduction.

17.9 Stacked generalization (stacking)


- *The combination of the learners' outputs is itself learned, but on a validation set.*
 1. We train L learners using the training set (in whatever way that introduces diversity).
 2. We apply these learners to a validation set. For each validation point \mathbf{x}_n , we obtain the outputs of the L learners ($\mathbf{y}_1, \dots, \mathbf{y}_L$). For classification, the output of learner l could be the class label or the posterior probabilities (softmax). For regression, the output of learner l is its real-valued output vector.
 3. We train a model f to predict the ground-truth output for a given instance \mathbf{x}_n from the learners' outputs for that instance.
- If we choose f to be a linear function, this produces a weighted average or vote, where the weights are learned on the validation set. But we can choose f to be nonlinear.
- By training f on the validation set (rather than the training set where the learners were trained), it learns to correct for mistakes that the learners make.
- Whether it works better than a fixed, simple combination rule (such as majority vote or average) depends on the problem.

17.10 Fine-tuning an ensemble

- Constructing an ensemble that does decrease the error requires some art in selecting the right number of learners and in making them be diverse and of the right accuracy.
- Typically, the resulting ensemble contains learners that are somewhat correlated with each other, or that are useless.
- It often helps to postprocess the ensemble by removing some learners (*ensemble pruning*), so we obtain a smaller ensemble (hence faster at test time) having about the same error.
- This is similar to feature selection and we can indeed apply feature selection algorithms. The most usual is *forward selection*: starting with an empty ensemble, we sequentially add one learner at a time, the one that gives highest accuracy when added to the previous ones.

18 Reinforcement learning

- How an autonomous agent that senses and acts in its environment can learn to choose optimal actions to achieve a goal. Ex.: board games (e.g. chess, backgammon), robot navigation (e.g. looking for the exit of a maze).
 - Each time the agent performs an action, it may receive a reward (or penalty) to indicate how good the resulting state of the environment is.
 - Often the reward is delayed.
At the end of the game (positive reward if we win, negative if we lose); or when (if) the robot reaches the exit.
 - The task of the agent is to learn from this reward to choose sequences of actions that produce the greatest cumulative reward.
- As in other machine learning settings, the task is to learn a function, here a control policy $\pi: \mathcal{S} \rightarrow \mathcal{A}$ that maps states $s \in \mathcal{S}$ to actions $a \in \mathcal{A}$, but with the following differences:
 - *Delayed reward*: we are not given pairs (state,action) to train the function, as in supervised learning. Instead, we (may) have a reward when the agent executes an action. There is no such thing as the best move at any given time, what matters is the sequence of moves. Supervised learning is “learning with a teacher”. Reinforcement learning is “learning with a critic”. The critic doesn’t tell us what to do but only how well we have done in the past. Its feedback is scarce and when it comes, it comes late.
 - *Credit assignment problem*: after taking several actions and getting the reward, which actions helped? So that we can record and recall them later on.
A reinforcement learning program learns to generate an internal value for the intermediate states and actions in terms of how good they are in leading us to the goal. Having learned this internal reward mechanism, the agent can just take local actions to maximize it.
 - *Exploration vs exploitation*: the agent influences the distribution of training examples by the actions it chooses. It should trade off exploration of unknown states and actions (to gain information) vs exploitation of states and actions that it has already learned will yield high reward (to maximize its cumulative reward).
 - *Partially observed states*: practically, sensors provide only partial information about the environment state (e.g. a camera doesn’t provide the robot location).
- Basic elements of a reinforcement learning problem:
 - *Agent*: the decision maker (e.g. game player, robot). It has *sensors* to observe the environment (e.g. robot camera).
 - *Environment* (e.g. board, maze). At any time t , the environment is in a certain *state* s_t that is one of a set of possible states \mathcal{S} (e.g. board state, robot position). Often, there is an initial state and a goal state.
 - A set \mathcal{A} of possible *actions* a_t (e.g. legal chess movements, possible robot steps). The state changes after an action: $s_{t+1} = \delta(s_t, a_t)$. The solution requires a sequence of actions.
 - *Reward* $r_t = r(s_t, a_t) \in \mathbb{R}$: the feedback we receive, usually at the end of the game. It helps to learn the policy.
 - *Policy* $\pi: \mathcal{S} \rightarrow \mathcal{A}$: a control strategy for choosing actions that achieve a goal.



- The task of the agent: *perform sequences of actions, observe their consequences and learn a control policy that, from any initial state, chooses actions that maximize the reward accumulated over time.*

The learning task

- There are several settings of the reinforcement learning problem. We may assume:
 - that the agent’s actions are deterministic, or that they are nondeterministic;
 - that the agent can predict the next state that will result from each action, or that it cannot.
 - that the agent is trained by giving it examples of optimal action sequences, or that it must train itself by performing actions of its own choice.

We consider a simple setting: deterministic actions that satisfy the Markov property. At each discrete time step t :

- The agent senses the current state s_t and performs an action a_t .
- The environment gives a reward $r_t = r(s_t, a_t)$ and a next state $s_{t+1} = \delta(s_t, a_t)$.

The functions r, δ are part of the environment and are not necessarily known to the agent; they depend only on the current state and action and not on the previous ones (Markov assumption).


- The task of the agent: learn a policy $\pi: \mathcal{S} \rightarrow \mathcal{A}$ for selecting its next action given the current observed state s_t : $\pi(s_t) = a_t$. We seek a policy that produces the greatest *cumulative reward* over time:

$$V^\pi(s_t) = r_t + \gamma r_{t+1} + \gamma^2 r_{t+2} + \dots = \sum_{i=0}^{\infty} \gamma^i r_{t+i}$$

i.e., the weighted sum of rewards obtained by starting at s_t and following the policy π : $a_i = \pi(s_i)$ for $i \geq t$. The constant $0 \leq \gamma < 1$ (*discount rate*) determines the relative value of delayed vs immediate rewards ($\gamma = 0$ considers only the most immediate reward). We take $\gamma < 1$ because, practically, we prefer rewards sooner rather than later (e.g. robot running on a battery).

- Hence: learn a policy that maximizes $V^\pi(s)$ for all states $s \in \mathcal{S}$: $\pi^* = \arg \max_{\pi} V^\pi(s) \forall s \in \mathcal{S}$, and $V^*(s)$ is the value function of such an optimal policy π^* .

Ex.: grid world with GOAL state, squares = states, arrows = actions ($\gamma = 0.9$):



Q-learning


- Learning $\pi^*: \mathcal{S} \rightarrow \mathcal{A}$ directly is difficult because we don't have examples (s, a) . Instead, we learn a numerical evaluation function $Q(s, a)$ of states and actions, and derive π from it.
- Define the *Q function*: $Q(s, a) = r(s, a) + \gamma V^*(\delta(s, a))$.
- We want to find $\pi^*(s) = \arg \max_a Q(s, a)$. This can be done with the following iterative algorithm. Noting that $V^*(s) = \max_{a'} Q(s, a')$, we can write the following recursive expression for Q : $Q(s, a) = r(s, a) + \gamma \max_{a'} Q(\delta(s, a), a')$. The algorithm starts with a table $\hat{Q}(s, a) = 0$ for every (s, a) , and repeatedly does the following:
 - observe its current state s
 - execute some action a
 - observe the new state $s' = \delta(s, a)$ and reward $r = r(s, a)$
 - update the table entry (s, a) like this: $\hat{Q}(s, a) \leftarrow r + \gamma \max_{a'} \hat{Q}(s', a')$.
- Each update of $\hat{Q}(s, a)$ affects the old state s (not the new one s').
- Assuming $r(s, a)$ is bounded and every possible (state,action) pair is visited infinitely often, one can prove: 1) the \hat{Q} values never decrease, and are between 0 and their optimal values Q ; 2) Q -learning converges to the correct Q function.
- Note *the agent need not know the functions r and δ ahead of time* (which is often not practical, e.g. for a mobile robot). It just needs to know their particular values as it takes actions, moves to new states and receives the corresponding rewards (i.e., it *samples* those functions).

In choosing new actions to apply, the agent has an exploration-exploitation tradeoff (repeat actions that seem currently good vs trying new ones). Commonly, one moves from pure exploration at the beginning towards exploitation as training proceeds.

If we have perfect knowledge of the functions r and δ (i.e., we can evaluate them for any (s, a)), we can find the optimal policy more efficiently with a *dynamic programming* algorithm to solve *Bellman's equation*: $V^*(s) = E \{r(s, \pi(s)) + \gamma V^*(\delta(s, \pi(s)))\} \forall s \in \mathcal{S}$. This is convenient in some applications, e.g. factory automation and job scheduling problems.

Ex.: grid world from a particular current \hat{Q} table, at state s_1 , taking action a_R and moving to state s_2 . The update is:

$$\begin{aligned} \hat{Q}(s_1, a_R) &\leftarrow r + \gamma \max_{a'} \hat{Q}(s_2, a') \\ &= 0 + 0.9 \max(66, 81, 100) \\ &= 90 \end{aligned}$$



Training consists of a series of *episodes*, each beginning at a random state and executing actions until reaching the goal state. The first update occurs when reaching the goal state and receiving a nonzero reward; in subsequent episodes, updates propagate backward, eventually filling the entire $\hat{Q}(\cdot, \cdot)$ table.

Nondeterministic rewards and actions

- Instead of deterministic functions $r(s, a)$ and $\delta(s, a)$, we have probability distributions $p(r|s, a)$ and $P(s'|s, a)$ (Markov assumption).

Ex.: roll of dice in backgammon, noisy sensors and effectors in robot navigation.

- Q -learning generalizes by using expectations wrt these distributions instead of deterministic functions:

- Value of a policy: $V^\pi(s_t) = \mathbb{E} \{ \sum_{i=0}^{\infty} \gamma^i r_{t+i} \}$.
- Optimal policy (as before): $\pi^* = \arg \max_{\pi} V^\pi(s) \forall s \in \mathcal{S}$.
- Q function: $Q(s, a) = \mathbb{E} \{ r(s, a) + \gamma V^*(\delta(s, a)) \} = \mathbb{E} \{ r(s, a) \} + \gamma \mathbb{E} \{ V^*(\delta(s, a)) \} = \mathbb{E} \{ r(s, a) \} + \gamma \sum_{s'} P(s'|s, a) V^*(s')$.
- Recursive expression for Q : $Q(s, a) = \mathbb{E} \{ r(s, a) \} + \gamma \sum_{s'} P(s'|s, a) \max_{a'} Q(s', a')$.
- Update for \hat{Q} at iteration k : $\hat{Q}_k(s, a) \leftarrow (1 - \alpha_k) \hat{Q}_{k-1}(s, a) + \alpha_k (r + \gamma \max_{a'} \hat{Q}_{k-1}(s', a'))$, where the step size α_k decreases towards zero as $k \rightarrow \infty$ for convergence to occur.

This is similar to SGD compared to GD. Convergence is slow and may require many thousands of iterations.

Generalizing from examples

- Using a table of states \times actions to represent Q has two limitations:
 - In practice, there may be a large or infinite number of states and actions (e.g. turning the steering wheel by a continuous angle).
 - Even if we could store the table, the previous Q -learning algorithm performs a kind of rote learning: it simply stores seen (state,action) pairs and doesn't generalize to unseen ones. Indeed, the convergence proof assumes every possible (state,action) pair is visited (infinitely often).

- Instead, we can represent $Q(s, a)$ with a function approximator, such as a neural net $Q(s, a; \mathbf{W})$ with inputs (s, a) , weights \mathbf{W} , and a single output corresponding to the value of $Q(s, a)$. Each $\hat{Q}(s, a)$ update in the Q -learning algorithm provides one training example, which is used to update the weights via SGD.

A classic example: the 1995 TD-gammon program for playing backgammon (which has $\approx 10^{20}$ states and randomness because of the dice roll) used a neural net trained on 1.5 million self-generated games (playing against itself) and achieved master-level play.

Partially observable states

- With a Markov decision process, we assume the agent can observe the state directly. In some applications, this is not possible (e.g. a mobile robot with a camera has a certain view of its environment but does not observe its location directly).
- This can be modeled using a *Partially Observable Markov Decision Process (POMDP)*, which uses $P(o|s, a)$ to model an observation o given the current state s (unobserved) and action a , rather than $P(s'|s, a)$ (e.g. o = camera image, s = 3D location of the robot).

This is similar to the difference between a discrete (observable) Markov process and a hidden Markov model (HMM).