

IB031 Úvod do strojového učení

Tomáš Brázdil

Course Info

Resources:

- ▶ Lectures & tutorials (the **main** source)
- ▶ Many books, few perfect for introductory level
One relatively good, especially the first part:
A. Géron. Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems. O'Reilly Media; 3rd edition, 2022
- ▶ (Almost) infinitely many online courses, tutorials, materials, etc.

Evaluation

The evaluation is composed of three parts:

- ▶ Mid-term exam: Written exam from the material of the first half of the semester.
- ▶ End-term exam: The "big" one containing everything from the semester (with possibly more stress in the second half).
- ▶ Projects: During tutorials, you will work on larger projects (in pairs or triples).

Each part contributes the following number of points:

- ▶ Mid-term exam: 25
- ▶ End-term exam: 50
- ▶ Project: 25

To pass, you need to obtain at least 60 points.

Distinguishing Properties of the Course

- ▶ Introductory, prerequisites are held to a minimum
- ▶ Formal and precise: Be prepared for a complete and “mathematical” description of presented methods.

Distinguishing Properties of the Course

- ▶ Introductory, prerequisites are held to a minimum
- ▶ Formal and precise: Be prepared for a complete and “mathematical” description of presented methods.

I assume that you have basic knowledge of

- ▶ Elementary mathematical notation (operations on sets, logic, etc.)
- ▶ Linear algebra: Vectors in \mathbb{R}^n , operations on vectors (including the dot product). Geometric interpretation!
- ▶ Calculus: Functions of multiple real variables, partial derivatives, basic differential calculus.
- ▶ Probability: Notion of probability distribution, random variables/vectors, expectation.

What Is Machine Learning?

Machine learning is the science (and art) of programming computers so they can **learn from data**.

What Is Machine Learning?

Machine learning is the science (and art) of programming computers so they can **learn from data**.

Here is a slightly more general definition:

Arthur Samuel, 1959

Machine learning is the field of study that allows computers to learn without being explicitly programmed.

What Is Machine Learning?

Machine learning is the science (and art) of programming computers so they can **learn from data**.

Here is a slightly more general definition:

Arthur Samuel, 1959

Machine learning is the field of study that allows computers to learn without being explicitly programmed.

And a more engineering-oriented one:

Tom Mitchell, 1997

A computer program is said to learn from experience E concerning some task T and some performance measure P if its performance on T , as measured by P , improves with experience E .

Example

In the context of spam filtering:

- ▶ The task T is to flag spam in new emails.
- ▶ The experience E is represented by a set of emails labeled either spam or ham by hand (the training data).
- ▶ The performance measure P could be the accuracy, which is the ratio of the number of correctly classified emails and all emails.

There are many more performance measures; we will study the basic ones later.

Example

In the context of spam filtering:

- ▶ The task T is to flag spam in new emails.
- ▶ The experience E is represented by a set of emails labeled either spam or ham by hand (the training data).
- ▶ The performance measure P could be the accuracy, which is the ratio of the number of correctly classified emails and all emails.

There are many more performance measures; we will study the basic ones later.

In the context of housing price prediction:

- ▶ The task T is to predict prices of new houses based on their basic parameters (size, number of bathrooms, etc.)
- ▶ The experience E is represented by information about existing houses.
- ▶ The performance measure P could be, e.g., an absolute difference between the predicted and the real price.

Examples (cont.)

In the context of game playing:

- ▶ The task T is to play chess.
- ▶ The experience E is represented by a series of self-plays where the computer plays against itself.
- ▶ The performance measure P is winning/losing the game.

Here, the trick is to spread the delayed and limited feedback about the result of the game throughout the individual decisions in the game.

Examples (cont.)

In the context of game playing:

- ▶ The task T is to play chess.
- ▶ The experience E is represented by a series of self-plays where the computer plays against itself.
- ▶ The performance measure P is winning/losing the game.
Here, the trick is to spread the delayed and limited feedback about the result of the game throughout the individual decisions in the game.

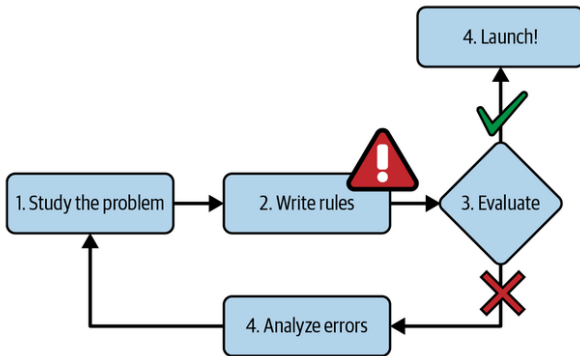
In the context of customer behavior:

- ▶ The task T is to group customers with similar shopping habits in an e-shop.
- ▶ The experience E consists of lists of items individual customers bought in the shop.
- ▶ The performance measure P ?
Measure how "nicely" the customers are grouped.
(whether people with similar habits, as seen by humans, fall into the same group).

Comparison of Programming and Learning

How to code the spam filter?

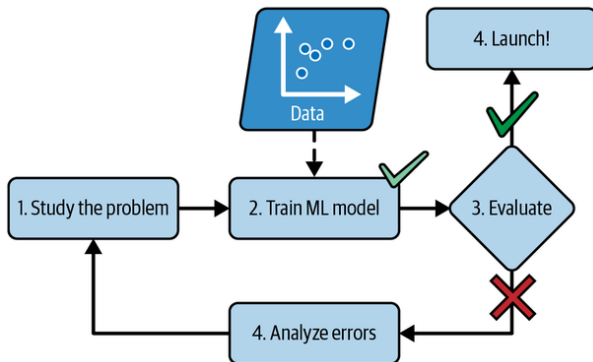
- ▶ Examine what spam mails typically contain: Specific words ("Viagra"), sender's address, etc.
- ▶ Write down a rule-based system that detects specific features.
- ▶ Test the program on new emails and (most probably) go back to look for more spam features.



Comparison of Programming and Learning

The machine learning way:

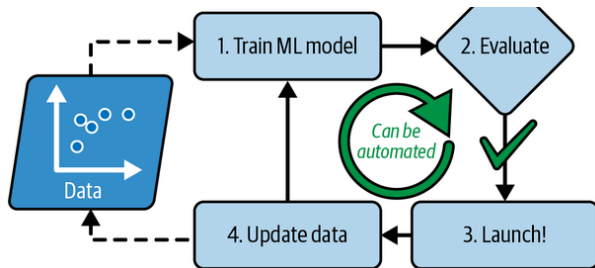
- ▶ Study the problem and collect lots of emails, labeling them spam or ham.
- ▶ Train a machine learning model that reads an email and decides whether it's spam or ham.
- ▶ Test the model and (most probably) go back to collect more data and adjust the model.



ML Solutions are Adaptive

Spam filter: Authors of spam might and will adapt to your spam filter (possibly change the wording to pass through).

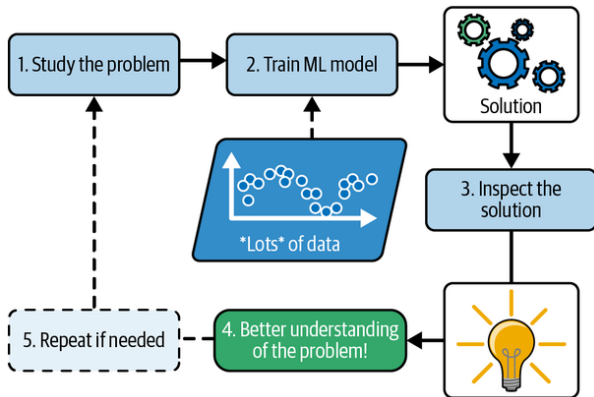
ML systems can be adjusted to new situations by retraining on new data (unless the data becomes ugly).



ML for Human Understanding

Spam filter: A trained system can be inspected for notorious spam features.

Some models allow direct inspection, such as decision trees or linear/logistic regression models.



Usage of Machine Learning

Machine learning suits various applications, especially where traditional methods fall short. Here are some areas where it excels:

- ▶ Solving complex problems where fine-tuning and rule-based solutions are inadequate.
- ▶ Tackling complex issues that resist traditional problem-solving approaches.
- ▶ Adapting to fluctuating environments through retraining on new data.
- ▶ Gaining insights from large and complex datasets.

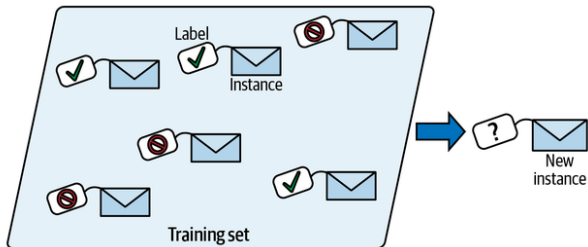
In summary, machine learning offers innovative solutions and adaptability for today's complex and ever-changing problems, (sometimes) providing insights beyond the reach of traditional approaches.

Types of Learning

There are main categories based on information available during the training:

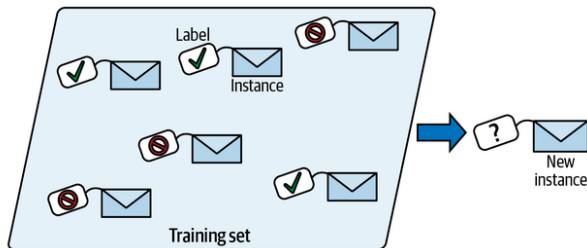
- ▶ Supervised learning
- ▶ Unsupervised learning
- ▶ Semi-supervised learning
- ▶ Self-supervised learning
- ▶ Reinforcement learning

Supervised Learning



Labels are available for all input data.

Supervised Learning

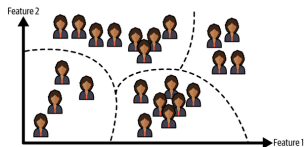
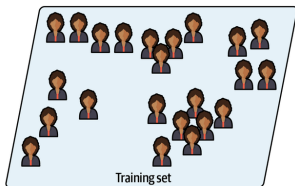


Labels are available for all input data.

Typical supervised learning tasks are

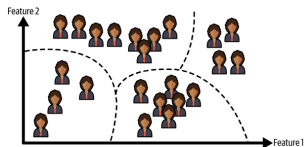
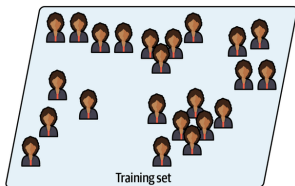
- ▶ *Classification* where the aim is to classify inputs into (typically few) classes
(e.g., the spam filter where the classes are spam/ham)
- ▶ *Regression* where a numerical value is output for a given input
(e.g., housing prices)

Unsupervised Learning



No labels are available for input data.

Unsupervised Learning

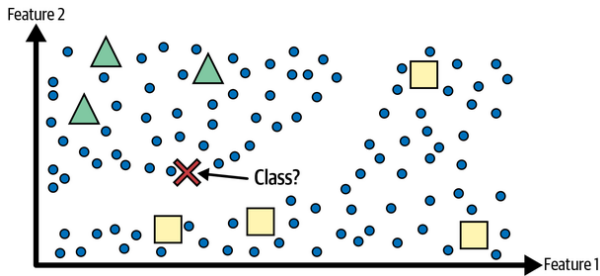


No labels are available for input data.

Typical unsupervised learning tasks are

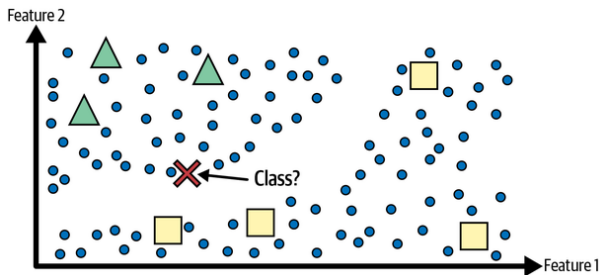
- ▶ *Clustering* where inputs are grouped according to their features
(e.g., clients of a bank grouped according to their age, wealth, etc.)
- ▶ *Association* where interesting relations and rules are discovered among the features of inputs
(e.g., market basket mining where associations between various types of goods are being learned from the behavior of customers)
- ▶ *Dimensionality reduction* reduce high-dimensional data to few dimensions (e.g., images to few image features)

Semi-Supervised Learning



Labels for some data.

Semi-Supervised Learning

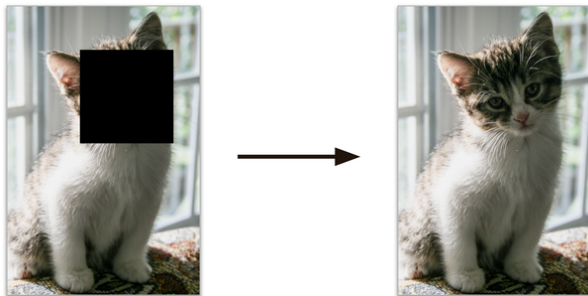


Labels for some data.

For example, Medical data, where elaborate diagnosis is available only for some patients.

Combines supervised and unsupervised learning: e.g., clusters all data and labels the unlabeled inputs with the most common labels in their clusters.

Self-Supervised Learning

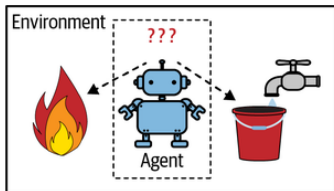


Generate labels from (unlabeled) inputs.

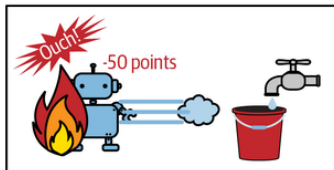
The goal is to learn typical features of the data.

It can be later modified to generate images, classify, etc.

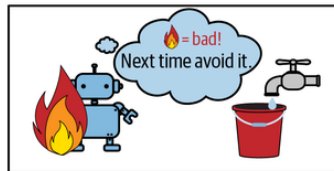
Reinforcement Learning



- 1 Observe
- 2 Select action using policy



- 3 Action!
- 4 Get reward or penalty



- 5 Update policy (learning step)
- 6 Iterate until an optimal policy is found

Learn from performing *actions* and getting feedback from *environment*.

ML Applications Highlights

- ▶ ChatGPT (and similar generative models)
 - ▶ The basis forms a generative language model, i.e., a text-generating model trained on texts in a self-supervised way
 - ▶ Recently extended to multimodal versions (text, image, sound)

ML Applications Highlights

- ▶ ChatGPT (and similar generative models)
 - ▶ The basis forms a generative language model, i.e., a text-generating model trained on texts in a self-supervised way
 - ▶ Recently extended to multimodal versions (text, image, sound)
- ▶ Machine translation, image captioning
 - ▶ Google translate, etc.
 - ▶ Typically (semi)-supervised learning,

ML Applications Highlights

- ▶ ChatGPT (and similar generative models)
 - ▶ The basis forms a generative language model, i.e., a text-generating model trained on texts in a self-supervised way
 - ▶ Recently extended to multimodal versions (text, image, sound)
- ▶ Machine translation, image captioning
 - ▶ Google translate, etc.
 - ▶ Typically (semi)-supervised learning,
- ▶ Various image recognition and processing tasks
 - ▶ In medicine, such systems are slowly making their way into hospitals as assistance tools
 - ▶ Automotive, quality control, advertising, etc., etc., etc.

ML Applications Highlights

- ▶ ChatGPT (and similar generative models)
 - ▶ The basis forms a generative language model, i.e., a text-generating model trained on texts in a self-supervised way
 - ▶ Recently extended to multimodal versions (text, image, sound)
- ▶ Machine translation, image captioning
 - ▶ Google translate, etc.
 - ▶ Typically (semi)-supervised learning,
- ▶ Various image recognition and processing tasks
 - ▶ In medicine, such systems are slowly making their way into hospitals as assistance tools
 - ▶ Automotive, quality control, advertising, etc., etc., etc.
- ▶ Science
 - ▶ Chemistry & biology: E.g., prediction of features of chemical compounds (Alpha-fold)

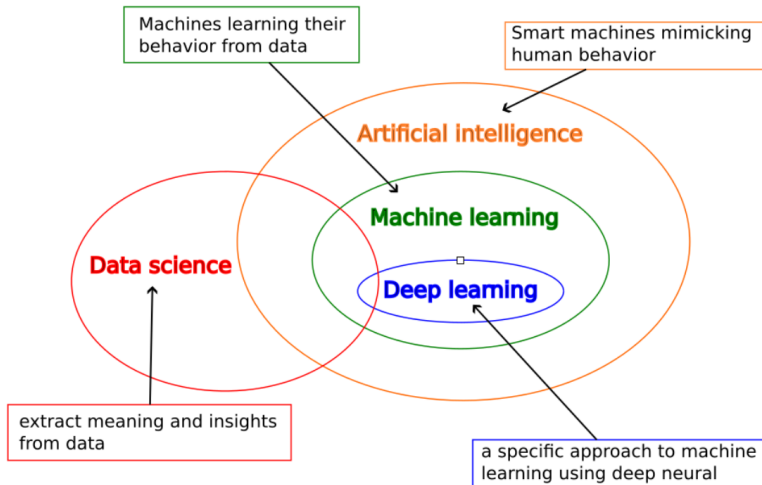
ML Applications Highlights

- ▶ ChatGPT (and similar generative models)
 - ▶ The basis forms a generative language model, i.e., a text-generating model trained on texts in a self-supervised way
 - ▶ Recently extended to multimodal versions (text, image, sound)
- ▶ Machine translation, image captioning
 - ▶ Google translate, etc.
 - ▶ Typically (semi)-supervised learning,
- ▶ Various image recognition and processing tasks
 - ▶ In medicine, such systems are slowly making their way into hospitals as assistance tools
 - ▶ Automotive, quality control, advertising, etc., etc., etc.
- ▶ Science
 - ▶ Chemistry & biology: E.g., prediction of features of chemical compounds (Alpha-fold)
- ▶ Various "table" data processing in finance, management, etc.
 - ▶ Often straightforward methods (linear/logistic regression)
 - ▶ Essential but not fancy

ML Applications Highlights

- ▶ ChatGPT (and similar generative models)
 - ▶ The basis forms a generative language model, i.e., a text-generating model trained on texts in a self-supervised way
 - ▶ Recently extended to multimodal versions (text, image, sound)
- ▶ Machine translation, image captioning
 - ▶ Google translate, etc.
 - ▶ Typically (semi)-supervised learning,
- ▶ Various image recognition and processing tasks
 - ▶ In medicine, such systems are slowly making their way into hospitals as assistance tools
 - ▶ Automotive, quality control, advertising, etc., etc., etc.
- ▶ Science
 - ▶ Chemistry & biology: E.g., prediction of features of chemical compounds (Alpha-fold)
- ▶ Various "table" data processing in finance, management, etc.
 - ▶ Often straightforward methods (linear/logistic regression)
 - ▶ Essential but not fancy
- ▶ Game playing: More fancy than useful, learning models beating humans in several difficult games.

ML in Context



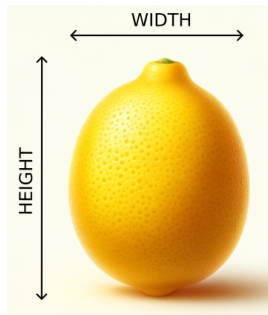
Supervised Learning

Example - Fruit Recognition

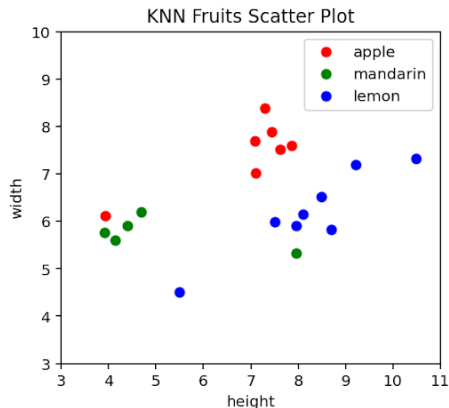
The goal: Create an automatic system for fruit recognition, concretely apple, lemon, and mandarin.

Inputs: Measures of *height* and *width* of each fruit.

Suppose we have a dataset of dimensions of several fruits labeled with the correct class.



Data



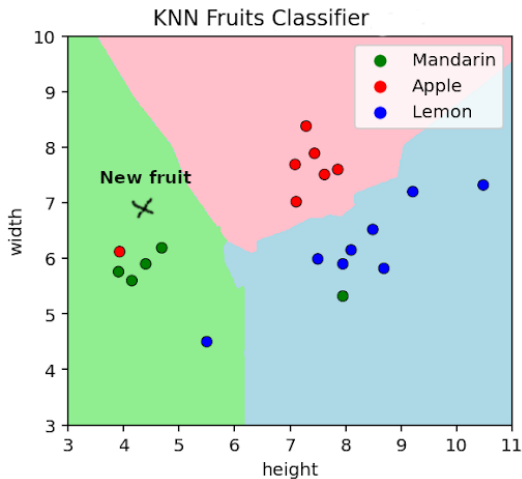
	height	width	fruit
0	3.91	5.76	Mandarin
1	7.09	7.69	Apple
2	10.48	7.32	Lemon
3	9.21	7.20	Lemon
4	7.95	5.90	Lemon
5	7.62	7.51	Apple
6	7.95	5.32	Mandarin
7	4.69	6.19	Mandarin
8	7.50	5.99	Lemon
9	7.11	7.02	Apple
10	4.15	5.60	Mandarin
11	7.29	8.38	Apple
12	8.49	6.52	Lemon
13	7.44	7.89	Apple
14	7.86	7.60	Apple
15	3.93	6.12	Apple
16	4.40	5.90	Mandarin
17	5.50	4.50	Lemon
18	8.10	6.15	Lemon
19	8.69	5.82	Lemon

Use similarity to solve the problem.

KNN Classification

Given a new fruit.
What is it?

Find five closest
examples



Where is the machine learning?

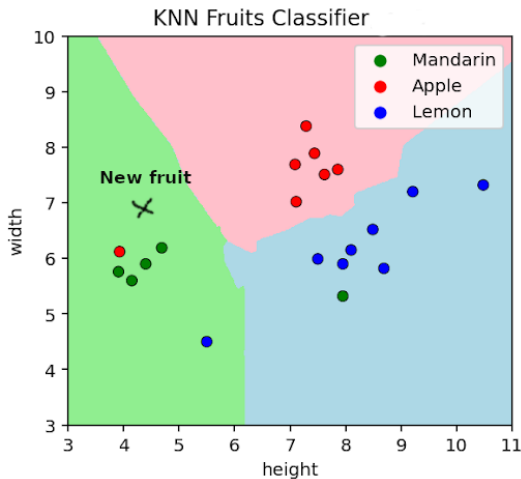
KNN Classification

Given a new fruit.
What is it?

Find five closest
examples

Among the five closest:

- ▶ $M = 4$ mandarins
- ▶ $A = 1$ apples
- ▶ $L = 0$ lemons



Where is the machine learning?

KNN Classification

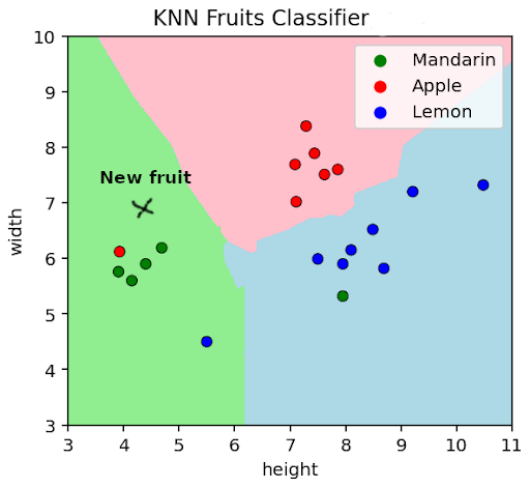
Given a new fruit.
What is it?

Find five closest
examples

Among the five closest:

- ▶ $M = 4$ mandarins
- ▶ $A = 1$ apples
- ▶ $L = 0$ lemons

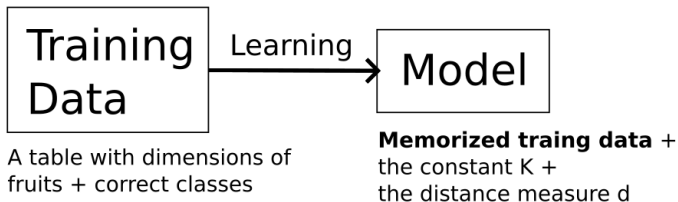
It is a **mandarin**!



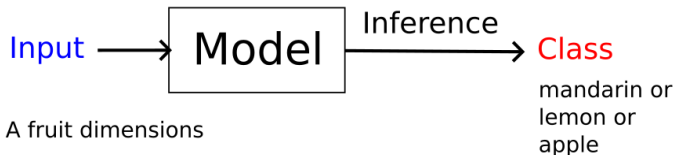
Where is the machine learning?

Learning in Fruit Classification with KNN

Learning:



Inference:



Fruit Classification Algorithm

Input: A fruit F with dimensions *height*, *width*

Output: *mandarin*, *lemon*, *apple*

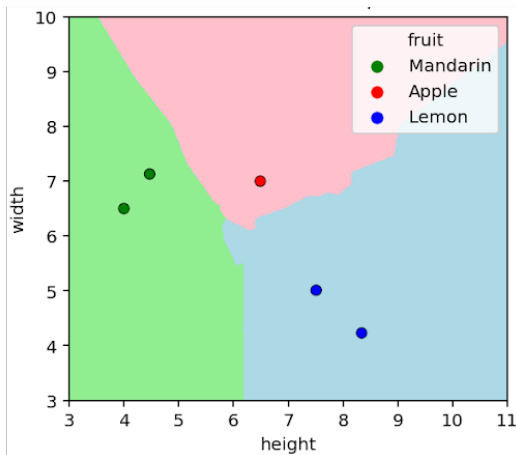
- 1: Find K examples $\{E_1, \dots, E_K\}$ in the dataset whose dimensions are closest to the dimensions of the fruit F
- 2: Count the number of examples of each class in $\{E_1, \dots, E_K\}$
 - M mandarins in $\{E_1, \dots, E_K\}$
 - L lemons in $\{E_1, \dots, E_K\}$
 - A apples in $\{E_1, \dots, E_K\}$
- 3: **if** $M \geq L$ and $M \geq A$ **then return** *mandarin*
- 4: **else if** $L \geq A$ **then return** *lemon*
- 5: **else return** *apple*
- 6: **end if**

Does it work?

Testing the Model for Fruit Classification

Consider a test set of new instances ($K = 5$, d is Euclidean):

height	width	fruit
4.0	6.5	Mandarin
4.47	7.13	Mandarin
6.49	7.0	Apple
7.51	5.01	Lemon
8.34	4.23	Lemon



Perfect classification of new data! Just deploy and sell!!

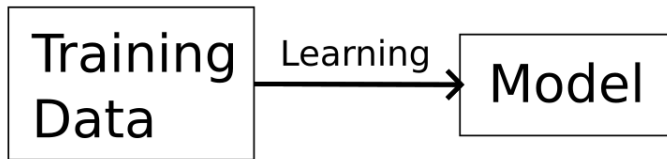
K Nearest Neighbors

... on ideal data

Learning and Inference

Two crucial components of machine learning are the following:

Learning:
Creating model



Inference:
Using model



Training Data

Assume table training data, i.e., of the form

x_{11}	x_{12}	\cdots	x_{1n}	c_1
x_{21}	x_{22}	\cdots	x_{2n}	c_2
\vdots	\vdots	\ddots	\vdots	\vdots
x_{p1}	x_{p2}	\cdots	x_{pn}	c_p

height	width	fruit
4.0	6.5	Mandarin
4.47	7.13	Mandarin
6.49	7.0	Apple
7.51	5.01	Lemon
8.34	4.23	Lemon

Formally, we define **training dataset**

$$\mathcal{T} = \{(\vec{x}_k, c_k) \mid k = 1, \dots, p\}$$

Here each $\vec{x}_k \in \mathbb{R}^n$ is an input vector and $c_k \in C$ is the correct class.

$$\mathcal{T} = \{(4.0, 6.5), M), \\ (4.47, 7.13), M), \\ (6.49, 7.0), A), \\ \dots\}$$

KNN: Learning

Consider the training set:

$$\mathcal{T} = \{(\vec{x}_k, c_k) \mid k = 1, \dots, p\}$$

and *memorize it* exactly as it is.

Store in a table.

Possibly use a clever representation allowing fast computation of nearest neighbors such as KDTrees (out of the scope of this lecture).

Also,

- ▶ determine the number of neighbors $K \in \mathbb{N}$,
- ▶ and the distance measure d .

Inference in KNN

Assume a KNN "trained" by memorizing

$\mathcal{T} = \{(\vec{x}_k, c_k) \in \mathbb{R}^n \times C \mid k = 1, \dots, p\}$, a constant $K \in \mathbb{N}$ and a distance measure d .

For d , consider Euclidean distance, but different norms may also be used to define different distance measures.

Inference in KNN

Assume a KNN "trained" by memorizing

$\mathcal{T} = \{(\vec{x}_k, c_k) \in \mathbb{R}^n \times C \mid k = 1, \dots, p\}$, a constant $K \in \mathbb{N}$ and a distance measure d .

For d , consider Euclidean distance, but different norms may also be used to define different distance measures.

Input: A vector $\vec{z} = (z_1, \dots, z_n) \in \mathbb{R}^n$

Output: A class from C

- 1: Find K indices of examples $X = \{i_1, \dots, i_K\} \subseteq \{1, \dots, p\}$ with minimum distance to \vec{z} , i.e., satisfying

$$\max\{d(\vec{z}, \vec{x}_\ell) \mid \ell \in X\} \leq \min\{d(\vec{z}, \vec{x}_\ell) \mid \ell \in \{1, \dots, p\} \setminus X\}$$

- 2: For every $c \in C$ count the number $\#c$ of elements ℓ in X such that $c_\ell = c$
- 3: Return some

$$c_{\max} \in \arg \max_{c \in C} \#c$$

A class $c_{\max} \in C$ which maximizes $\#c$.

The resulting model

What exactly constitutes the model? The *model* consists of

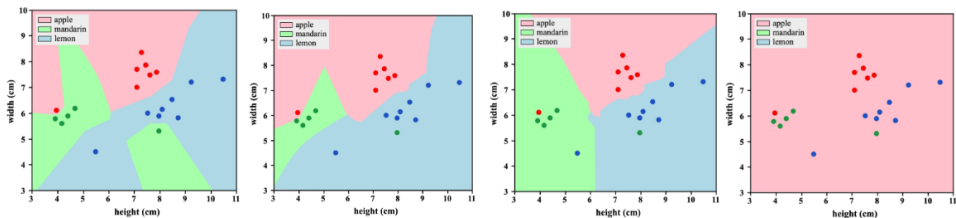
- ▶ The *trained parameters*: In this case the memorized training data.
- ▶ The *hyperparameters* set “from the outside”: In this case, the number of neighbors K and the distance measure d .

The resulting model

What exactly constitutes the model? The *model* consists of

- ▶ The *trained parameters*: In this case the memorized training data.
- ▶ The *hyperparameters* set “from the outside”: In this case, the number of neighbors K and the distance measure d .

Note that different settings of K lead to different classifiers (for the same d):



In Practice

...to get an efficient solution:

In Practice

...to get an efficient solution:

- ▶ Deal with issues in the data
 - ▶ Data almost always comes in weird formats, with inconsistencies, missing values, wrong values, etc.
 - ▶ Data rarely have the ideal form for a given learning model.

We need to ingest, validate, and preprocess the data.

In Practice

... to get an efficient solution:

- ▶ Deal with issues in the data
 - ▶ Data almost always comes in weird formats, with inconsistencies, missing values, wrong values, etc.
 - ▶ Data rarely have the ideal form for a given learning model.

We need to ingest, validate, and preprocess the data.

- ▶ Deal with issues in the model
 - ▶ In KNN, the training memorizes the example, but at least the K can be tuned.

We need to tune the model.

In Practice

... to get an efficient solution:

- ▶ Deal with issues in the data
 - ▶ Data almost always comes in weird formats, with inconsistencies, missing values, wrong values, etc.
 - ▶ Data rarely have the ideal form for a given learning model.

We need to ingest, validate, and preprocess the data.

- ▶ Deal with issues in the model
 - ▶ In KNN, the training memorizes the example, but at least the K can be tuned.

We need to tune the model.

- ▶ Deal with the wrong model by testing and validation in as realistic conditions as possible.

In Practice

... to get an efficient solution:

- ▶ Deal with issues in the data

- ▶ Data almost always comes in weird formats, with inconsistencies, missing values, wrong values, etc.
- ▶ Data rarely have the ideal form for a given learning model.

We need to ingest, validate, and preprocess the data.

- ▶ Deal with issues in the model

- ▶ In KNN, the training memorizes the example, but at least the K can be tuned.

We need to tune the model.

- ▶ Deal with the wrong model by testing and validation in as realistic conditions as possible.
- ▶ Deal with deployment - real-world application issues involving, e.g., implementation in embedded devices with limited resources.

Models Considered in This Course

Throughout this course, we will meet the following models:

- ▶ KNN (already did)
- ▶ Decision trees
- ▶ (Naive) Bayes classifier
- ▶ Clustering: K-means and hierarchical
- ▶ Linear and logistic regression
- ▶ Support Vector Machines (SVM)
- ▶ Kernel linear models
- ▶ Neural networks (light intro to feed-forward networks)
- ▶ Ensemble methods + random forests
- ▶ (maybe some reinforcement learning)

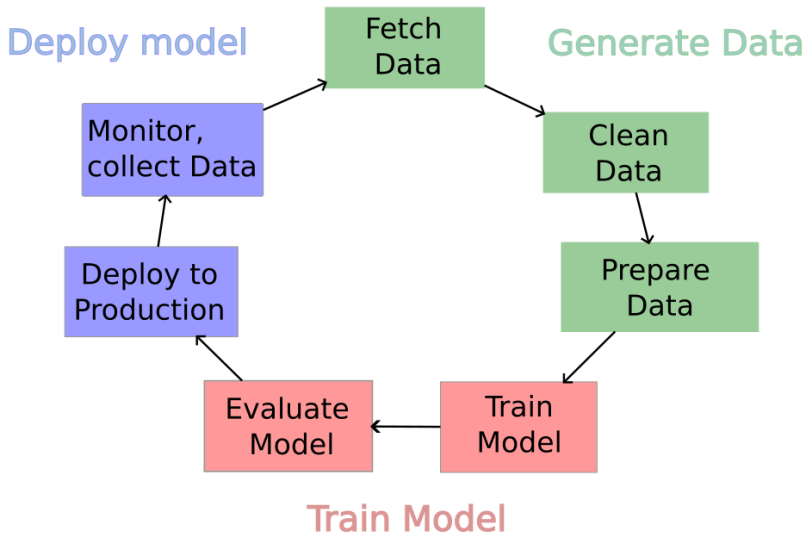
Models Considered in This Course

Throughout this course, we will meet the following models:

- ▶ KNN (already did)
- ▶ Decision trees
- ▶ (Naive) Bayes classifier
- ▶ Clustering: K-means and hierarchical
- ▶ Linear and logistic regression
- ▶ Support Vector Machines (SVM)
- ▶ Kernel linear models
- ▶ Neural networks (light intro to feed-forward networks)
- ▶ Ensemble methods + random forests
- ▶ (maybe some reinforcement learning)

...but first, let us see the whole machine-learning pipeline.

Machine Learning Pipeline



Fetch Data

Always start with

Fetch Data

Always start with

- ▶ The problem formulation & understanding.

For example, diagnosis of diabetes from medical records. What info is (possibly) sufficient for such a diagnosis?

Fetch Data

Always start with

- ▶ The problem formulation & understanding.

For example, diagnosis of diabetes from medical records. What info is (possibly) sufficient for such a diagnosis?

- ▶ Find data sources.

In our example, the sources are hospitals. It would be best to persuade them to give you the data and sign a contract.

Fetch Data

Always start with

- ▶ The problem formulation & understanding.

For example, diagnosis of diabetes from medical records. What info is (possibly) sufficient for such a diagnosis?

- ▶ Find data sources.

In our example, the sources are hospitals. It would be best to persuade them to give you the data and sign a contract.

- ▶ Collect the data.

In our example, the data is possibly small (just tables with results of tests). But for other diagnoses, you may include huge amounts of data from MRI, CT, etc. Then, the collection itself might be a serious technical problem.

Fetch Data

Always start with

- ▶ The problem formulation & understanding.

For example, diagnosis of diabetes from medical records. What info is (possibly) sufficient for such a diagnosis?

- ▶ Find data sources.

In our example, the sources are hospitals. It would be best to persuade them to give you the data and sign a contract.

- ▶ Collect the data.

In our example, the data is possibly small (just tables with results of tests). But for other diagnoses, you may include huge amounts of data from MRI, CT, etc. Then, the collection itself might be a serious technical problem.

- ▶ Integrate data from various sources.

A serious diagnostic system must be trained/tested on data from many hospitals. You must blend the data from various sources (different formats, etc.).

Fetch Data

For simple “toy” machine learning projects, you may fetch prepared datasets from various databases on the internet.

Fetch Data

For simple “toy” machine learning projects, you may fetch prepared datasets from various databases on the internet.

The data should be stored in an identified location and versioned. You will probably keep adding data and training models on the ever-changing datasets. You have to be able to keep track of the changes and map training data to particular models.

Tools such as ML Flow or Weights & Biases might be helpful.

Fetch Data

For simple “toy” machine learning projects, you may fetch prepared datasets from various databases on the internet.

The data should be stored in an identified location and versioned. You will probably keep adding data and training models on the ever-changing datasets. You have to be able to keep track of the changes and map training data to particular models.

Tools such as ML Flow or Weights & Biases might be helpful.

Data Separation

At this point, you should randomize the ordering of the data and select a test set to be used in model evaluation!

The test data are supposed to simulate the actual conditions, i.e., they should be “unseen”.

Fetch Data

For simple “toy” machine learning projects, you may fetch prepared datasets from various databases on the internet.

The data should be stored in an identified location and versioned. You will probably keep adding data and training models on the ever-changing datasets. You have to be able to keep track of the changes and map training data to particular models.

Tools such as ML Flow or Weights & Biases might be helpful.

Data Separation

At this point, you should randomize the ordering of the data and select a test set to be used in model evaluation!

The test data are supposed to simulate the actual conditions, i.e., they should be “unseen”.

Data Exploration

Compute basic statistics to identify missing values, outliers, etc.

Clean Data

The cleaning usually comprises the following steps:

- ▶ Fix or remove incorrect or corrupted values.
- ▶ Identify outliers and decide what to do with them.
Outliers may harm some training methods and are not “representative”. However, sometimes, they naturally belong to the dataset, and expert insight is needed.
- ▶ Fix formatting.
For example, the Date may be expressed in many ways, and a simple Yes/No answer.
- ▶ Resolve missing values (by either removing the whole examples or imputing)
Many methods have been developed for missing values imputation. It is a susceptible issue because new values may strongly bias the model.
- ▶ Remove duplicates.

The above steps often affect the training and need expertise in the application domain.

Later in this course, we will discuss techniques for data cleaning.

ID	Age	Income	Gender	Customer_Satisfaction
1	38	46641.356413713	nan	Unsatisfied
2	42	49129.0615585107	female	Neutral
3	18	119965.049731014	Male	nan
4	18	66828.0762224329	nan	very unsatisfied
5	58	57422.2721106762	female	very unsatisfied
6	28	59502.8174855665	Other	Satisfied
7	18	42659.6675768587	Other	Neutral
8	18	54019.1173206374	Other	Satisfied
9	40	25429.1604541137	female	Unsatisfied
10	21	15595.5862129548	Other	Satisfied
11	18	58094.2328460069	Other	very unsatisfied
12	18	39097.3278583155	female	Very Satisfied
13	30		Other	Satisfied
14	50	30617.3914472273	Female	Very Satisfied
15	18		nan	Neutral
16	34	39902.4430953214	male	nan
17	49	68381.6997683133	Female	Very Satisfied
18	33	44796.0962271524	Other	Very Satisfied
19	47	39218.9560738814	Female	very unsatisfied
20		14544.9226784447	Other	Satisfied

Prepare Data

Unlike cleaning, which is application-dependent, data preparation/transformation is model-dependent. This usually subsumes:

- ▶ **Scaling:** Settings values of inputs to a similar range.

Some models, especially those utilizing distance, are sensitive to large differences between input sizes.

- ▶ **Encoding:** Encode non-numeric data using real-valued vectors.

Many models, especially those based on geometry, work only with numeric data. Non-numeric data such as Yes/No, Short/Medium/Long must be encoded appropriately.

- ▶ **Binning or Discretization** Convert continuous features into discrete bins to capture patterns in ranges.

Comment: Sometimes **Normalization**, that is changing the distribution of inputs to resemble the normal distribution, is mentioned. However, this step is typically not essential for machine learning itself. However, it is important to use statistical inference to test the significance of learned parameters.

Prepare Data

- ▶ **Feature selection** Throw out input features that are too “similar” to other features.

For example, if the temperature is measured both in Celsius and in Kelvin, keep one of them. The relationship can, of course, be a more complex (non-linear) correlation.

Prepare Data

- ▶ **Feature selection** Throw out input features that are too “similar” to other features.

For example, if the temperature is measured both in Celsius and in Kelvin, keep one of them. The relationship can, of course, be a more complex (non-linear) correlation.

- ▶ **Dimensionality reduction** Transforming data from \mathbb{R}^n to \mathbb{R}^m where $m \ll n$.

Growing dimension means growing difficulty of training for all models. Some models cease to work for high-dimensional data. The reduction typically searches for a few important characteristic features of inputs.

Prepare Data

- ▶ **Feature selection** Throw out input features that are too “similar” to other features.

For example, if the temperature is measured both in Celsius and in Kelvin, keep one of them. The relationship can, of course, be a more complex (non-linear) correlation.

- ▶ **Dimensionality reduction** Transforming data from \mathbb{R}^n to \mathbb{R}^m where $m \ll n$.

Growing dimension means growing difficulty of training for all models. Some models cease to work for high-dimensional data. The reduction typically searches for a few important characteristic features of inputs.

- ▶ **Feature aggregation** Introducing new features using operations on the original ones.

We will see kernel transformations later in this course, allowing simple models to solve complex problems.

Train Model

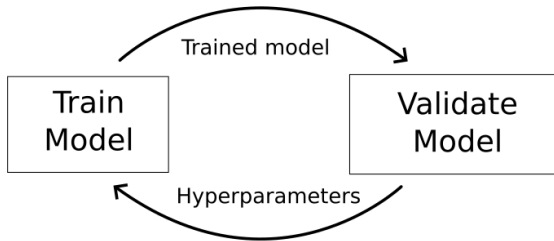
Now the dataset has been cleaned; we may train a model.

Train Model

Now the dataset has been cleaned; we may train a model.

Before training, we should split the dataset into

- ▶ *training* dataset on which the model will learn
- ▶ *validation* dataset on which we fine-tune hyperparameters



The resulting model is obtained after several iterations of the above process.

Evaluate Model

Here, we use the test set that we separated during data fetching.

In some cases, a brand new test set can be generated.

patients are examined regularly, creating new records continuously.

In some cases, it is tough to obtain new data.

For example, new expensive and difficult measurements are needed to obtain new data.

Evaluate Model

Here, we use the test set that we separated during data fetching.

In some cases, a brand new test set can be generated.

patients are examined regularly, creating new records continuously.

In some cases, it is tough to obtain new data.

For example, new expensive and difficult measurements are needed to obtain new data.

Critical issue: Make sure that you are truly testing
exactly the whole inference process.

Often, just a model is tested, and the testing and production inference engines are separated. This leads to truly nasty errors in the production!

We will discuss various generic metrics helpful in measuring the quality of the resulting model.

Deploy to Production

Deployment of machine learning models is a complex question, application dependent.

The recently emerging area of MLOps is concerned with the engineering side of the model deployment.

Deploy to Production

Deployment of machine learning models is a complex question, application dependent.

The recently emerging area of MLOps is concerned with the engineering side of the model deployment.

From the technical point of view, the typical issues solved by ML Ops teams are

- ▶ how to extract/process data in real-time
- ▶ how much storage is required
- ▶ how to store/collect model (and data) artifacts/predictions
- ▶ how to set up APIs, tools, and software environments
- ▶ What the period of predictions (instantaneous or batch predictions) should be
- ▶ how to set up hardware requirements (or cloud requirements for on-cloud environments) by the computational resources required
- ▶ how to set up a pipeline for continuous training and parameter tuning

Deploy to Production

From the user's point of view:

- ▶ How to get a sensible and valuable user output?
 - ▶ AI researchers will be satisfied with tons of running text in terminals.
 - ▶ “Normal” people need a graphical interface with understandable output.
 - ▶ Experts working in other domains typically demand speed and clarity at the extreme.

Deploy to Production

From the user's point of view:

- ▶ How to get a sensible and valuable user output?
 - ▶ AI researchers will be satisfied with tons of running text in terminals.
 - ▶ “Normal” people need a graphical interface with understandable output.
 - ▶ Experts working in other domains typically demand speed and clarity at the extreme.
- ▶ How do you persuade users that the AI is working for them?
 - ▶ Especially if safety is at stake, you need to have outstanding arguments and explanations ready for end-users.
 - ▶ In many areas, the devices need to be certified (medicine, automotive) for ML-based systems.

This complex subject will be only touched on in this course.

Monitor, collect Data

Deployed machine learning models must be constantly monitored.

Because of the influx of new data, ML models work in highly dynamic environments.

For example, an image-processing medical diagnostic model suddenly misdiagnosed a patient because a nurse marked the sample with a marker pen.

Every customer has a different infrastructure and may produce data slightly differently.

Data for retraining and improvement should be stored.

Also, many areas allow the *active learning* where users provide feedback for (continuous) retraining of the models.

Data

Data Science Example

You receive data from a medical researcher concerning a project that you are eager to work on.

Data Science Example

You receive data from a medical researcher concerning a project that you are eager to work on.

The data consists of a 1000 lines table with five columns:

012	232	33.5	0	10.7
020	121	16.9	2	210.1
027	165	24.0	0	427.6
...				

The aim is to predict the last field given the others.

Data Science Example

You receive data from a medical researcher concerning a project that you are eager to work on.

The data consists of a 1000 lines table with five columns:

012	232	33.5	0	10.7
020	121	16.9	2	210.1
027	165	24.0	0	427.6
...				

The aim is to predict the last field given the others.

The medical researcher does not elaborate further on the data, but they seem to be pretty easy to work with, right?

Data Science Example

You receive data from a medical researcher concerning a project that you are eager to work on.

The data consists of a 1000 lines table with five columns:

012	232	33.5	0	10.7
020	121	16.9	2	210.1
027	165	24.0	0	427.6
...				

The aim is to predict the last field given the others.

The medical researcher does not elaborate further on the data, but they seem to be pretty easy to work with, right?

After a few days, you have trained a model that predicts numbers resembling the ones in the table.

You contact the medical researcher and discuss the results.

Model Discussion

Researcher: So, you got the data for all the patients?

Model Discussion

Researcher: So, you got the data for all the patients?

Data Miner: Yes. I haven't had much time for analysis, but I do have a few interesting results.

Model Discussion

Researcher: So, you got the data for all the patients?

Data Miner: Yes. I haven't had much time for analysis, but I do have a few interesting results.

Researcher: Amazing. There were so many data issues with this set of patients that I couldn't do much.

Model Discussion

Researcher: So, you got the data for all the patients?

Data Miner: Yes. I haven't had much time for analysis, but I do have a few interesting results.

Researcher: Amazing. There were so many data issues with this set of patients that I couldn't do much.

Data Miner: Oh? I didn't hear about any possible problems.

Model Discussion

Researcher: So, you got the data for all the patients?

Data Miner: Yes. I haven't had much time for analysis, but I do have a few interesting results.

Researcher: Amazing. There were so many data issues with this set of patients that I couldn't do much.

Data Miner: Oh? I didn't hear about any possible problems.

Researcher: Well, first, there is field 5, the variable we want to predict. It's common knowledge among people who analyze this type of data that results are better if you work with the log of the values, but I didn't discover this until later. Was it mentioned to you?

Model Discussion

Researcher: So, you got the data for all the patients?

Data Miner: Yes. I haven't had much time for analysis, but I do have a few interesting results.

Researcher: Amazing. There were so many data issues with this set of patients that I couldn't do much.

Data Miner: Oh? I didn't hear about any possible problems.

Researcher: Well, first, there is field 5, the variable we want to predict. It's common knowledge among people who analyze this type of data that results are better if you work with the log of the values, but I didn't discover this until later. Was it mentioned to you?

Data Miner: No.

Model Dicsuccion

Researcher: But surely you heard about what happened to field 4? It's supposed to be measured on a scale from 1 to 10, with 0 indicating a missing value, but because of a data entry error, all 10's were changed into 0's. Unfortunately, since some of the patients have missing values for this field, it's impossible to say whether a 0 in this field is a real 0 or a 10. Quite a few of the records have that problem.

Model Dicsuccion

Researcher: But surely you heard about what happened to field 4? It's supposed to be measured on a scale from 1 to 10, with 0 indicating a missing value, but because of a data entry error, all 10's were changed into 0's. Unfortunately, since some of the patients have missing values for this field, it's impossible to say whether a 0 in this field is a real 0 or a 10. Quite a few of the records have that problem.

Data Miner: Interesting. Were there any other problems?

Model Dicsuccion

Researcher: But surely you heard about what happened to field 4? It's supposed to be measured on a scale from 1 to 10, with 0 indicating a missing value, but because of a data entry error, all 10's were changed into 0's. Unfortunately, since some of the patients have missing values for this field, it's impossible to say whether a 0 in this field is a real 0 or a 10. Quite a few of the records have that problem.

Data Miner: Interesting. Were there any other problems?

Researcher: Yes, fields 2 and 3 are basically the same, but I assume that you probably noticed that.

Model Dicsuccion

Researcher: But surely you heard about what happened to field 4? It's supposed to be measured on a scale from 1 to 10, with 0 indicating a missing value, but because of a data entry error, all 10's were changed into 0's. Unfortunately, since some of the patients have missing values for this field, it's impossible to say whether a 0 in this field is a real 0 or a 10. Quite a few of the records have that problem.

Data Miner: Interesting. Were there any other problems?

Researcher: Yes, fields 2 and 3 are basically the same, but I assume that you probably noticed that.

Data Miner: Yes, but these fields were only weak predictors of field 5.

Model Discussion

Researcher: Anyway, given all those problems, I'm surprised you were able to accomplish anything.

Data Miner: True, but my results are really quite good. Field 1 is a very strong predictor of field 5. I'm surprised that this wasn't noticed before.

Researcher: What? Field 1 is just an identification number.

Data Miner: Nonetheless, my results speak for themselves.

Researcher: Oh, no! I just remembered. We assigned ID numbers after we sorted the records based on field 5. There is a strong connection, but it isn't very sensible. Sorry.

Model Discussion

Researcher: Anyway, given all those problems, I'm surprised you were able to accomplish anything.

Data Miner: True, but my results are really quite good. Field 1 is a very strong predictor of field 5. I'm surprised that this wasn't noticed before.

Researcher: What? Field 1 is just an identification number.

Data Miner: Nonetheless, my results speak for themselves.

Researcher: Oh, no! I just remembered. We assigned ID numbers after we sorted the records based on field 5. There is a strong connection, but it isn't very sensible. Sorry.

OK, what's the point?

You have to

Understand the task you want to solve and the data!

Data Objects

Data objects represent entities we work with (e.g., classify them).

For example, in cancer prediction, the data objects are patients. In fruit classification, the data objects are individual fruits.

Data Objects

Data objects represent entities we work with (e.g., classify them).

For example, in cancer prediction, the data objects are patients. In fruit classification, the data objects are individual fruits.

Data objects are described by *attributes* (or *features* or *variables*).

For example, the age, weight, genetic profile, and other patient characteristics. Or the width and height of a fruit.

Attributes vs Features vs Variables

The name differs from field to field.

Attributes vs Features vs Variables

The name differs from field to field.

So, the following names are usually used as synonyms:

- ▶ *Attributes* - used mostly by database and data mining experts.
- ▶ *Features* - used mostly by machine learning experts.
- ▶ *Variables* - used mostly by statisticians.

Attributes vs Features vs Variables

The name differs from field to field.

So, the following names are usually used as synonyms:

- ▶ *Attributes* - used mostly by database and data mining experts.
- ▶ *Features* - used mostly by machine learning experts.
- ▶ *Variables* - used mostly by statisticians.

One may make some distinctions

- ▶ *Attributes* represent information about the object without any additional assumptions.
- ▶ *Features* assume that their values are somewhat characteristic of the object.
- ▶ *Variables* assume that there is some process behind them (typically a random process in the case of statistics).

Data Types - Categorical Attributes

Categorical attributes (nominal attributes) are symbols or names of things.

- ▶ Each value represents some kind of category, code, or state.
- ▶ Values are not ordered and should not be used quantitatively (in computer science, the values are known as enumerations).

Data Types - Categorical Attributes

Categorical attributes (nominal attributes) are symbols or names of things.

- ▶ Each value represents some kind of category, code, or state.
- ▶ Values are not ordered and should not be used quantitatively (in computer science, the values are known as enumerations).
- ▶ **Examples:**

$\text{hair_color} \in \{\text{black, brown, blond, red, auburn, gray, white}\}$

$\text{marital_status} \in \{\text{single, married, divorced, widowed}\}$

$\text{customer_ID} \in \{0, 1, 2, \dots\}$

Even though the last one is usually expressed using numbers, it should not be used quantitatively.

Data Types - Categorical Attributes

Categorical attributes (nominal attributes) are symbols or names of things.

- ▶ Each value represents some kind of category, code, or state.
- ▶ Values are not ordered and should not be used quantitatively (in computer science, the values are known as enumerations).
- ▶ **Examples:**

$\text{hair_color} \in \{\text{black, brown, blond, red, auburn, gray, white}\}$

$\text{marital_status} \in \{\text{single, married, divorced, widowed}\}$

$\text{customer_ID} \in \{0, 1, 2, \dots\}$

Even though the last one is usually expressed using numbers, it should not be used quantitatively.

Binary attributes are categorical attributes with only two values.

DataTypes - Ordinal Attributes

Ordinal attribute is an attribute with values that have a meaningful order or ranking among them.

DataTypes - Ordinal Attributes

Ordinal attribute is an attribute with values that have a meaningful order or ranking among them.

Examples:

$\text{drink_size} \in \{\text{small, medium, large}\}$

$\text{grades} \in \{\text{A, B, C, D, E, F}\}$

It can also be obtained by discretizing numeric quantities into series of intervals.

Ordinal attributes do not allow arithmetic operations.

DataTypes - Ordinal Attributes

Ordinal attribute is an attribute with values that have a meaningful order or ranking among them.

Examples:

$\text{drink_size} \in \{\text{small, medium, large}\}$

$\text{grades} \in \{\text{A, B, C, D, E, F}\}$

It can also be obtained by discretizing numeric quantities into series of intervals.

Ordinal attributes do not allow arithmetic operations.

Categorical and ordinal attributes are called *qualitative* attributes.

Next, we look at numeric, i.e., *quantitative* attributes.

Data Types - Numeric Attributes

Numeric attributes are quantities represented by numbers.

Data Types - Numeric Attributes

Numeric attributes are quantities represented by numbers.

Distinguish two types: *Interval-scale* and *ratio-scale*.

	INTERVAL SCALE	RATIO SCALE
Measurement interval	Equal intervals between consecutive points.	Equal intervals with the presence of a true zero.
Absolute zero	Lacks a true zero point.	Possesses a true zero point.
Statistical analysis	Limited to addition and subtraction	Allows for meaningful multiplication and division.
Meaningful ratios	Ratios are not meaningful due to the lack of zero.	Ratios are meaningful due to the presence of zero.
Examples	Celsius temperature (20 degrees not twice as hot as 10 degrees)	Height, duration, etc. (20 meters is twice as large as 10 meters)

Discrete vs Continuous Attributes

Often, two kinds of numeric attributes are distinguished:

Discrete vs Continuous Attributes

Often, two kinds of numeric attributes are distinguished:

- ▶ *Discrete*

A finite or countably infinite range of values, i.e., integers may represent the values.

Some (but not all) authors count the qualitative (categorical, ordinal) attributes among the discrete attributes.

Discrete vs Continuous Attributes

Often, two kinds of numeric attributes are distinguished:

- ▶ *Discrete*

A finite or countably infinite range of values, i.e., integers may represent the values.

Some (but not all) authors count the qualitative (categorical, ordinal) attributes among the discrete attributes.

- ▶ *Continuous*

A continuous range of values, an interval.

There are several more or less formal definitions of continuous attributes in the literature. For example:

- ▶ All non-discrete variables.
- ▶ Have an infinite number of values between any two values.
- ▶ Their values are measured (??).

Deeper characteristics of data (statistical properties, etc.) will be examined at tutorials.

Decision Trees

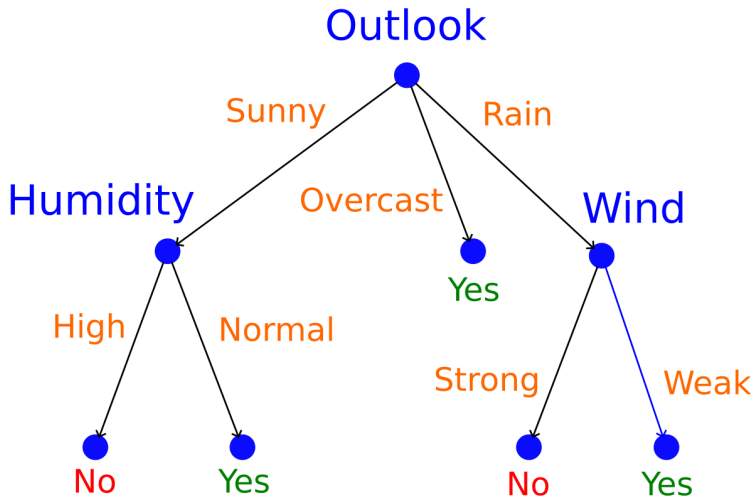
Decision Trees

- ▶ One of the widely used methods for machine learning.
- ▶ Intuitively simple, directly explainable.
- ▶ Basis for random forests (a powerful model).

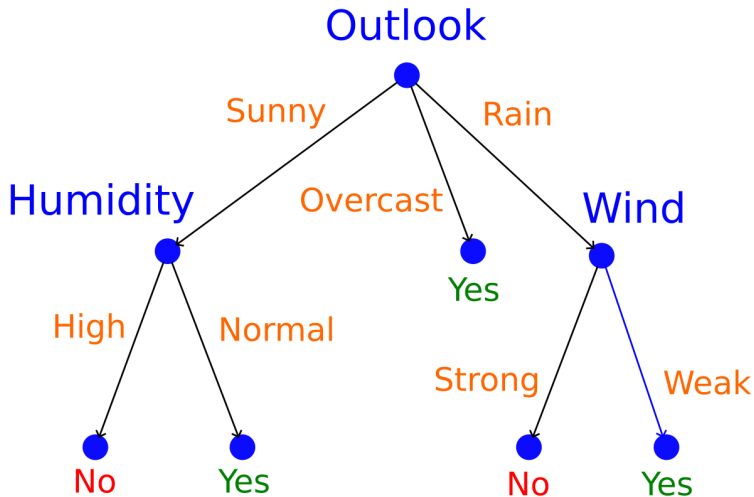
Decision Trees

- ▶ One of the widely used methods for machine learning.
- ▶ Intuitively simple, directly explainable.
- ▶ Basis for random forests (a powerful model).
- ▶ We will consider the ID3 algorithm.
Quinlan, 1979
- ▶ Various adjustments that appear in C4.5, CART, etc.

Consider the weather forecast for tennis playing. How would you decide whether to play today?



Consider the weather forecast for tennis playing. How would you decide whether to play today?



How do we obtain such a tree based on experience/data?

Learning Decision Trees

Consider data represented as follows:

- ▶ A finite set of *attributes* $\mathcal{A} = \{A_1, \dots, A_n\}$.
- ▶ Each attribute $A \in \mathcal{A}$ has its *set of values* $V(A)$.

We start with trees on discrete datasets, that is, assume $V(A)$ finite for all $A \in \mathcal{A}$.

Learning Decision Trees

Consider data represented as follows:

- ▶ A finite set of *attributes* $\mathcal{A} = \{A_1, \dots, A_n\}$.
- ▶ Each attribute $A \in \mathcal{A}$ has its *set of values* $V(A)$.

We start with trees on discrete datasets, that is, assume $V(A)$ finite for all $A \in \mathcal{A}$.

Objects to be classified are described by vectors of values of all attributes:

$$\vec{x} = (x_1, \dots, x_n) \in V(A_1) \times \dots \times V(A_n)$$

Given \vec{x} and an attribute A_k we denote by $A_k(\vec{x})$ the value x_k of the attribute A_k in \vec{x} .

Learning Decision Trees

Consider data represented as follows:

- ▶ A finite set of *attributes* $\mathcal{A} = \{A_1, \dots, A_n\}$.
- ▶ Each attribute $A \in \mathcal{A}$ has its *set of values* $V(A)$.

We start with trees on discrete datasets, that is, assume $V(A)$ finite for all $A \in \mathcal{A}$.

Objects to be classified are described by vectors of values of all attributes:

$$\vec{x} = (x_1, \dots, x_n) \in V(A_1) \times \dots \times V(A_n)$$

Given \vec{x} and an attribute A_k we denote by $A_k(\vec{x})$ the value x_k of the attribute A_k in \vec{x} .

Consider a set C of *classes*.

We consider a multiclass classification in general, i.e., C is an arbitrary finite set.

Example

The tennis problem:

- ▶ The attributes are:

$$A_1 = \textit{Outlook}, A_2 = \textit{Temperature}, A_3 = \textit{Humidity}, A_4 = \textit{Wind}$$

Example

The tennis problem:

- ▶ The attributes are:

$A_1 = \text{Outlook}, A_2 = \text{Temperature}, A_3 = \text{Humidity}, A_4 = \text{Wind}$

- ▶ The sets of values of the attributes:

- ▶ $V(A_1) = \{\text{Sunny}, \text{Overcast}, \text{Rain}\}$
- ▶ $V(A_2) = \{\text{Hot}, \text{Mild}, \text{Cool}\}$
- ▶ $V(A_3) = \{\text{High}, \text{Normal}\}$
- ▶ $V(A_4) = \{\text{Strong}, \text{Weak}\}$

Example

The tennis problem:

- ▶ The attributes are:

$$A_1 = \text{Outlook}, A_2 = \text{Temperature}, A_3 = \text{Humidity}, A_4 = \text{Wind}$$

- ▶ The sets of values of the attributes:

- ▶ $V(A_1) = \{\text{Sunny}, \text{Overcast}, \text{Rain}\}$

- ▶ $V(A_2) = \{\text{Hot}, \text{Mild}, \text{Cool}\}$

- ▶ $V(A_3) = \{\text{High}, \text{Normal}\}$

- ▶ $V(A_4) = \{\text{Strong}, \text{Weak}\}$

- ▶ Consider

$$\vec{x} = (\text{Overcast}, \text{Hot}, \text{Normal}, \text{Weak})$$

$$\in V(A_1) \times V(A_2) \times V(A_3) \times V(A_4)$$

Example

The tennis problem:

- ▶ The attributes are:

$$A_1 = \text{Outlook}, A_2 = \text{Temperature}, A_3 = \text{Humidity}, A_4 = \text{Wind}$$

- ▶ The sets of values of the attributes:

- ▶ $V(A_1) = \{\text{Sunny}, \text{Overcast}, \text{Rain}\}$

- ▶ $V(A_2) = \{\text{Hot}, \text{Mild}, \text{Cool}\}$

- ▶ $V(A_3) = \{\text{High}, \text{Normal}\}$

- ▶ $V(A_4) = \{\text{Strong}, \text{Weak}\}$

- ▶ Consider

$$\vec{x} = (\text{Overcast}, \text{Hot}, \text{Normal}, \text{Weak})$$

$$\in V(A_1) \times V(A_2) \times V(A_3) \times V(A_4)$$

Then

$$A_3(\vec{x}) = \text{Humidity}(\vec{x}) = \text{Normal}$$

$$A_4(\vec{x}) = \text{Wind}(\vec{x}) = \text{Weak}$$

Example

The tennis problem:

- ▶ The attributes are:

$$A_1 = \text{Outlook}, A_2 = \text{Temperature}, A_3 = \text{Humidity}, A_4 = \text{Wind}$$

- ▶ The sets of values of the attributes:

- ▶ $V(A_1) = \{\text{Sunny}, \text{Overcast}, \text{Rain}\}$

- ▶ $V(A_2) = \{\text{Hot}, \text{Mild}, \text{Cool}\}$

- ▶ $V(A_3) = \{\text{High}, \text{Normal}\}$

- ▶ $V(A_4) = \{\text{Strong}, \text{Weak}\}$

- ▶ Consider

$$\vec{x} = (\text{Overcast}, \text{Hot}, \text{Normal}, \text{Weak})$$

$$\in V(A_1) \times V(A_2) \times V(A_3) \times V(A_4)$$

Then

$$A_3(\vec{x}) = \text{Humidity}(\vec{x}) = \text{Normal}$$

$$A_4(\vec{x}) = \text{Wind}(\vec{x}) = \text{Weak}$$

- ▶ $C = \{\text{Yes}, \text{No}\}$

Decision Trees

Consider (directed, rooted) *trees* $\mathcal{T} = (T, E)$ where T is a set of nodes and $E \subseteq T \times T$ is a set of directed edges.

Decision Trees

Consider (directed, rooted) *trees* $\mathcal{T} = (T, E)$ where T is a set of nodes and $E \subseteq T \times T$ is a set of directed edges.

Denote by $T_{leaf} \subseteq T$ the set of all *leaves* of the tree and by T_{int} the set $T \setminus T_{leaf}$ of *internal nodes*.

Decision Trees

Consider (directed, rooted) *trees* $\mathcal{T} = (T, E)$ where T is a set of nodes and $E \subseteq T \times T$ is a set of directed edges.

Denote by $T_{leaf} \subseteq T$ the set of all *leaves* of the tree and by T_{int} the set $T \setminus T_{leaf}$ of *internal nodes*.

A *decision tree* is

- ▶ a tree $\mathcal{T} = (T, E)$ where
- ▶ each leaf $\tau \in T_{leaf}$ is assigned a class $C(\tau) \in C$,
- ▶ each internal node $\tau \in T_{int}$ is assigned an attribute $A(\tau) \in \mathcal{A}$,
- ▶ and there is a bijection between edges from τ and values of the attribute $A(\tau)$. Given an edge $(\tau, \tau') \in E$ we write $V(\tau, \tau')$ to denote the value of the attribute $A(\tau)$ assigned to the edge.

Decision Trees

Consider (directed, rooted) *trees* $\mathcal{T} = (T, E)$ where T is a set of nodes and $E \subseteq T \times T$ is a set of directed edges.

Denote by $T_{leaf} \subseteq T$ the set of all *leaves* of the tree and by T_{int} the set $T \setminus T_{leaf}$ of *internal nodes*.

A *decision tree* is

- ▶ a tree $\mathcal{T} = (T, E)$ where
- ▶ each leaf $\tau \in T_{leaf}$ is assigned a class $C(\tau) \in C$,
- ▶ each internal node $\tau \in T_{int}$ is assigned an attribute $A(\tau) \in \mathcal{A}$,
- ▶ and there is a bijection between edges from τ and values of the attribute $A(\tau)$. Given an edge $(\tau, \tau') \in E$ we write $V(\tau, \tau')$ to denote the value of the attribute $A(\tau)$ assigned to the edge.

Inference: Given an input \vec{x} , we traverse the tree from the root to a leaf, always choosing edges labeled with values of attributes from \vec{x} . The output is the class labeling the leaf.

Example

$$T = \{O, H, W, z_1, z_2, z_3, z_4, z_5\}$$

$$T_{leaf} = \{z_1, z_2, z_3, z_4, z_5\}, T_{int} = \{O, H, W\}$$

$$E = \{(O, H), (O, W), (H, z_1), (H, z_2), \\ (O, z_3), (W, z_4), (W, z_5)\}$$

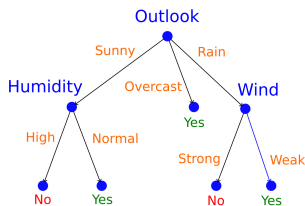
$$C(z_1) = C(z_4) = \text{No}, C(z_2) = C(z_3) = C(z_5) = \text{Yes}$$

$$A(O) = \text{Outlook}, A(H) = \text{Humidity}, A(W) = \text{Wind}$$

$$V(O, H) = \text{Sunny}, V(O, z_3) = \text{Overcast}, V(O, W) = \text{Rain}$$

$$V(H, z_1) = \text{High}, V(H, z_2) = \text{Normal}$$

$$V(W, z_4) = \text{Strong}, V(W, z_5) = \text{Weak}$$



Inference: For (*Rain, Hot, High, Strong*) we reach z_4 , yielding *No*.

Training Dataset

Consider a *training dataset*

$$\mathcal{D} = \{(\vec{x}_k, c_k) \mid k = 1, \dots, p\}$$

Here $\vec{x}_k \in V(A_1) \times \dots \times V(A_k)$ and $c_k \in C$ for every k .

Technically \mathcal{D} can be a multiset containing several occurrences of the same vector.

Index	Outlook	Temperature	Humidity	Wind	PlayTennis
1	Sunny	Hot	High	Weak	No
2	Sunny	Hot	High	Strong	No
3	Overcast	Hot	High	Weak	Yes
4	Rain	Mild	High	Weak	Yes
5	Rain	Cool	Normal	Weak	Yes
6	Rain	Cool	Normal	Strong	No
7	Overcast	Cool	Normal	Strong	Yes
8	Sunny	Mild	High	Weak	No
9	Sunny	Cool	Normal	Weak	Yes
10	Rain	Mild	Normal	Weak	Yes
11	Sunny	Mild	Normal	Strong	Yes
12	Overcast	Mild	High	Strong	Yes
13	Overcast	Hot	Normal	Weak	Yes
14	Rain	Mild	High	Strong	No

$$\begin{aligned}
\mathcal{D} = \{ & ((Sunny, Hot, High, Weak), No), \\
& ((Sunny, Hot, High, Strong), No) \\
& \dots \\
& ((Rain, Mild, High, Strong), No) \}
\end{aligned}$$

Learning Decision Trees

The learning algorithm ID3 works as follows:

- ▶ Start with the whole training dataset \mathcal{D} .

Learning Decision Trees

The learning algorithm ID3 works as follows:

- ▶ Start with the whole training dataset \mathcal{D} .
- ▶ If there is just a single class in \mathcal{D} , create a single node decision tree that returns the class.

Learning Decision Trees

The learning algorithm ID3 works as follows:

- ▶ Start with the whole training dataset \mathcal{D} .
- ▶ If there is just a single class in \mathcal{D} , create a single node decision tree that returns the class.
- ▶ Otherwise, identify an attribute $A \in \mathcal{A}$ which *best classifies* the examples in \mathcal{D} . For every $v \in V(A)$ we obtain

$$\mathcal{D}_v = \{\vec{x} \mid \vec{x} \in \mathcal{D}, A(\vec{x}) = v\}$$

We aim to have each \mathcal{D}_v as pure as possible, that is, ideally, to contain examples of just a single class.

Learning Decision Trees

The learning algorithm ID3 works as follows:

- ▶ Start with the whole training dataset \mathcal{D} .
- ▶ If there is just a single class in \mathcal{D} , create a single node decision tree that returns the class.
- ▶ Otherwise, identify an attribute $A \in \mathcal{A}$ which *best classifies* the examples in \mathcal{D} . For every $v \in V(A)$ we obtain

$$\mathcal{D}_v = \{\vec{x} \mid \vec{x} \in \mathcal{D}, A(\vec{x}) = v\}$$

We aim to have each \mathcal{D}_v as pure as possible, that is, ideally, to contain examples of just a single class.

- ▶ Finally,
 - ▶ create a root node τ of a decision tree,

Learning Decision Trees

The learning algorithm ID3 works as follows:

- ▶ Start with the whole training dataset \mathcal{D} .
- ▶ If there is just a single class in \mathcal{D} , create a single node decision tree that returns the class.
- ▶ Otherwise, identify an attribute $A \in \mathcal{A}$ which *best classifies* the examples in \mathcal{D} . For every $v \in V(A)$ we obtain

$$\mathcal{D}_v = \{\vec{x} \mid \vec{x} \in \mathcal{D}, A(\vec{x}) = v\}$$

We aim to have each \mathcal{D}_v as pure as possible, that is, ideally, to contain examples of just a single class.

- ▶ Finally,
 - ▶ create a root node τ of a decision tree,
 - ▶ assign the attribute A to τ ,

Learning Decision Trees

The learning algorithm ID3 works as follows:

- ▶ Start with the whole training dataset \mathcal{D} .
- ▶ If there is just a single class in \mathcal{D} , create a single node decision tree that returns the class.
- ▶ Otherwise, identify an attribute $A \in \mathcal{A}$ which *best classifies* the examples in \mathcal{D} . For every $v \in V(A)$ we obtain

$$\mathcal{D}_v = \{\vec{x} \mid \vec{x} \in \mathcal{D}, A(\vec{x}) = v\}$$

We aim to have each \mathcal{D}_v as pure as possible, that is, ideally, to contain examples of just a single class.

- ▶ Finally,
 - ▶ create a root node τ of a decision tree,
 - ▶ assign the attribute A to τ ,
 - ▶ for every $v \in V(A)$, recursively construct a decision tree with a root τ_v using \mathcal{D}_v ,

Learning Decision Trees

The learning algorithm ID3 works as follows:

- ▶ Start with the whole training dataset \mathcal{D} .
- ▶ If there is just a single class in \mathcal{D} , create a single node decision tree that returns the class.
- ▶ Otherwise, identify an attribute $A \in \mathcal{A}$ which *best classifies* the examples in \mathcal{D} . For every $v \in V(A)$ we obtain

$$\mathcal{D}_v = \{\vec{x} \mid \vec{x} \in \mathcal{D}, A(\vec{x}) = v\}$$

We aim to have each \mathcal{D}_v as pure as possible, that is, ideally, to contain examples of just a single class.

- ▶ Finally,
 - ▶ create a root node τ of a decision tree,
 - ▶ assign the attribute A to τ ,
 - ▶ for every $v \in V(A)$, recursively construct a decision tree with a root τ_v using \mathcal{D}_v ,
 - ▶ for every $v \in V(A)$ introduce an edge (τ, τ_v) assigned v .

```
1: function ID3(dataset  $\mathcal{D}$ , attribute set  $\mathcal{A}$ )
2:   Create a root node  $\tau$  for the tree
3:   if  $\mathcal{D} = \emptyset$  then
4:     Return the single node  $\tau$  assigned with a default class.
5:   else if all examples in  $\mathcal{D}$  are of the same class  $c$  then
6:     Return the single-node tree, where  $\tau$  is assigned  $c$ 
7:   else if set of attributes  $\mathcal{A}$  is empty then
8:     Return the single-node tree where  $\tau$  is assigned
       the most common class in  $\mathcal{D}$ 
9:   else
10:    Choose attribute  $A \in \mathcal{A}$  best classifying examples in  $\mathcal{D}$ .
11:    Set the decision attribute for  $\tau$  to  $A$ 
12:    for each value  $v \in D(A)$  do
13:      Compute a decision tree  $\text{ID3}(\mathcal{D}_v, \mathcal{A} \setminus \{A\})$  with root  $\tau_v$ ,
14:      add a new edge  $(\tau, \tau_v)$  assigned  $v$ .
15:    end for
16:  end if
17: return  $\tau$ 
18: end function
```

Best Classifying Attribute

We aim to choose an attribute that best informs us about the class.

As a result, we would possibly use as few attributes as possible and obtain a small tree containing only class-relevant decisions.

How to choose an attribute that best classifies examples in \mathcal{D} ?

There are several measures used in practice.

The most common are

- ▶ *information gain*
- ▶ *Gini impurity decrease*

Information Gain

The information gain is based on the notion of entropy.

Information Gain

The information gain is based on the notion of entropy.

We need some notation:

- ▶ Given a training dataset \mathcal{D} and a class $c \in \mathcal{C}$ we denote by p_c the proportion of examples with class c in \mathcal{D} .

Information Gain

The information gain is based on the notion of entropy.

We need some notation:

- ▶ Given a training dataset \mathcal{D} and a class $c \in \mathcal{C}$ we denote by p_c the proportion of examples with class c in \mathcal{D} .
- ▶ We define the *entropy* of \mathcal{D} by

$$Entropy(\mathcal{D}) = \sum_{c \in \mathcal{C}} -p_c \log_2 p_c$$

Information Gain

The information gain is based on the notion of entropy.

We need some notation:

- ▶ Given a training dataset \mathcal{D} and a class $c \in \mathcal{C}$ we denote by p_c the proportion of examples with class c in \mathcal{D} .
- ▶ We define the *entropy* of \mathcal{D} by

$$Entropy(\mathcal{D}) = \sum_{c \in \mathcal{C}} -p_c \log_2 p_c$$

- ▶ The *information gain* of an attribute A is then defined by

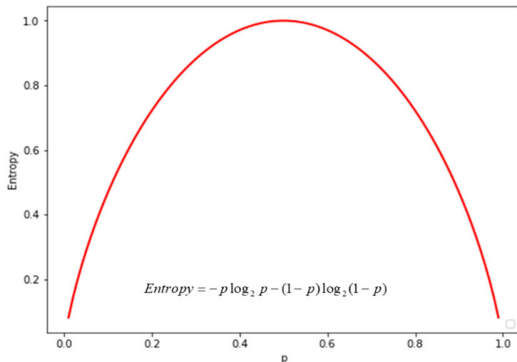
$$Gain(\mathcal{D}, A) = Entropy(\mathcal{D}) - \sum_{v \in V(A)} \frac{|\mathcal{D}_v|}{|\mathcal{D}|} Entropy(\mathcal{D}_v)$$

In every step of the ID3 algorithm, we choose an attribute *maximizing* the information gain for the current dataset \mathcal{D} .

Information Gain

The intuition behind information gain:

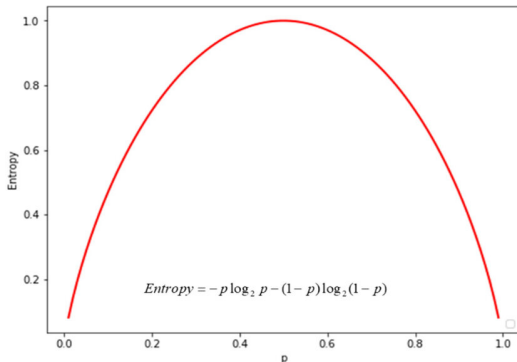
- Consider $C = \{0, 1\}$ and p the proportion of examples of class 1. p measures the “uncertainty” of the class:



Information Gain

The intuition behind information gain:

- Consider $C = \{0, 1\}$ and p the proportion of examples of class 1. p measures the “uncertainty” of the class:

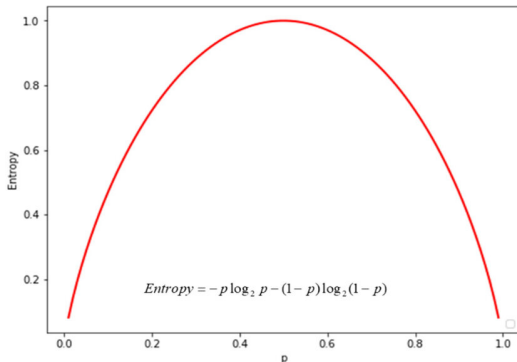


- $\sum_{v \in V(A)} \frac{|\mathcal{D}_v|}{|\mathcal{D}|} Entropy(\mathcal{D}_v)$ is weighted uncertainty of classes in each \mathcal{D}_v (weighted by the relative size of \mathcal{D}_v).

Information Gain

The intuition behind information gain:

- Consider $C = \{0, 1\}$ and p the proportion of examples of class 1. p measures the “uncertainty” of the class:



- $\sum_{v \in V(A)} \frac{|\mathcal{D}_v|}{|\mathcal{D}|} Entropy(\mathcal{D}_v)$ is weighted uncertainty of classes in each \mathcal{D}_v (weighted by the relative size of \mathcal{D}_v).
- $Gain(\mathcal{D}, A)$ measures reduction in uncertainty of classes by splitting \mathcal{D} according to A .

Gini Impurity

- We define *Gini impurity* of \mathcal{D} by

$$Gini(\mathcal{D}) = 1 - \sum_{c \in C} p_c^2$$

Gini Impurity

- ▶ We define *Gini impurity* of \mathcal{D} by

$$Gini(\mathcal{D}) = 1 - \sum_{c \in C} p_c^2$$

- ▶ The *impurity decrease* of an attribute A is then defined similarly to the gain in the entropy case

$$ImpDec(\mathcal{D}, A) = Gini(\mathcal{D}) - \sum_{v \in V(A)} \frac{|\mathcal{D}_v|}{|\mathcal{D}|} Gini(\mathcal{D}_v)$$

In every step of the ID3 algorithm, we choose an attribute *maximizing* the impurity decrease for the current dataset \mathcal{D} .

Gini Impurity

What is the intuition behind $Gini(\mathcal{D})$?

Gini Impurity

What is the intuition behind $Gini(\mathcal{D})$?

Assume we randomly independently choose objects from \mathcal{D} .

$1 - \sum_{c \in C} p_c^2$ is the probability of choosing two objects of different classes in two consecutive independent trials.

Indeed, p_c is the probability of choosing an object of class c , p_c^2 the probability of choosing objects of the class c twice, and $\sum_{c \in C} p_c^2$ the probability of choosing two objects of the same class.

In what follows (and at the exam), we will work only with the Gini impurity as it is easier to compute by hand.

Example

Consider our tennis example (see the table).

- ▶ Consider the whole dataset \mathcal{D} .
 - ▶ $p_{Yes} = 9/14$
 - ▶ $p_{No} = 5/14$
 - ▶ $Gini(\mathcal{D}) = 1 - (9/14)^2 - (5/14)^2 = 0.459$

Example

Consider our tennis example (see the table).

- ▶ Consider the whole dataset \mathcal{D} .
 - ▶ $p_{Yes} = 9/14$
 - ▶ $p_{No} = 5/14$
 - ▶ $Gini(\mathcal{D}) = 1 - (9/14)^2 - (5/14)^2 = 0.459$
- ▶ For $A = Outlook$ we get
 - ▶ $Gini(\mathcal{D}_{Sunny}) = 1 - (2/5)^2 - (3/5)^2 = 0.48$
 - ▶ $Gini(\mathcal{D}_{Overcast}) = 1 - 1^2 - 0^2 = 0$
 - ▶ $Gini(\mathcal{D}_{Rain}) = 1 - (3/5)^2 - (2/5)^2 = 0.48$

Example

Consider our tennis example (see the table).

- ▶ Consider the whole dataset \mathcal{D} .
 - ▶ $p_{Yes} = 9/14$
 - ▶ $p_{No} = 5/14$
 - ▶ $Gini(\mathcal{D}) = 1 - (9/14)^2 - (5/14)^2 = 0.459$
- ▶ For $A = Outlook$ we get
 - ▶ $Gini(\mathcal{D}_{Sunny}) = 1 - (2/5)^2 - (3/5)^2 = 0.48$
 - ▶ $Gini(\mathcal{D}_{Overcast}) = 1 - 1^2 - 0^2 = 0$
 - ▶ $Gini(\mathcal{D}_{Rain}) = 1 - (3/5)^2 - (2/5)^2 = 0.48$

Thus

$$\begin{aligned} ImpDec(\mathcal{D}, Outlook) &= \\ &0.459 - (5/14) \cdot 0.48 - (4/14) \cdot 0 - (5/14) \cdot 0.48 \\ &= 0.117 \end{aligned}$$

- ▶ $ImpDec(\mathcal{D}, Temperature) = 0.018$
- ▶ $ImpDec(\mathcal{D}, Humidity) = 0.091$
- ▶ $ImpDec(\mathcal{D}, Wind) = 0.030$

So the largest information gain is given by the *Outlook*.

Example

Going further on, consider $\mathcal{D} = \mathcal{D}_{Sunny}$. We get

- ▶ $ImpDec(\mathcal{D}, Temperature) = 0.279$
- ▶ $ImpDec(\mathcal{D}, Humidity) = 0.48$
- ▶ $ImpDec(\mathcal{D}, Wind) = 0.013$

The best choice attribute after *Sunny* in *Outlook* is *Humidity*.

Example

Going further on, consider $\mathcal{D} = \mathcal{D}_{Sunny}$. We get

- ▶ $ImpDec(\mathcal{D}, Temperature) = 0.279$
- ▶ $ImpDec(\mathcal{D}, Humidity) = 0.48$
- ▶ $ImpDec(\mathcal{D}, Wind) = 0.013$

The best choice attribute after *Sunny* in *Outlook* is *Humidity*.

Now consider $\mathcal{D} = \mathcal{D}_{Rain}$.

- ▶ $ImpDec(\mathcal{D}, Temperature) = 0.013$
- ▶ $ImpDec(\mathcal{D}, Humidity) = 0.013$
- ▶ $ImpDec(\mathcal{D}, Wind) = 0.48$

The best choice attribute after *Rain* in *Outlook* is *Wind*.

Attribute Importance Computation

How important are attributes for the trained tree \mathcal{T} ? Depends on

- ▶ how close they are to the root of \mathcal{T} ,
- ▶ how large information gain/decrease in impurity they give.

There are several formulae for computing the importance.

Attribute Importance Computation

How important are attributes for the trained tree \mathcal{T} ? Depends on

- ▶ how close they are to the root of \mathcal{T} ,
- ▶ how large information gain/decrease in impurity they give.

There are several formulae for computing the importance.

One example is *mean decrease impurity* defined as follows:

Consider a decision tree T trained on a dataset \mathcal{D} using the ID3.

Attribute Importance Computation

How important are attributes for the trained tree \mathcal{T} ? Depends on

- ▶ how close they are to the root of \mathcal{T} ,
- ▶ how large information gain/decrease in impurity they give.

There are several formulae for computing the importance.

One example is *mean decrease impurity* defined as follows:

Consider a decision tree T trained on a dataset \mathcal{D} using the ID3.

For every node τ of \mathcal{T} , denote by $\mathcal{D}[\tau]$ the subset of \mathcal{D} which was used in the ID3 procedure when the node τ was created (line 2).

Attribute Importance Computation

How important are attributes for the trained tree \mathcal{T} ? Depends on

- ▶ how close they are to the root of \mathcal{T} ,
- ▶ how large information gain/decrease in impurity they give.

There are several formulae for computing the importance.

One example is *mean decrease impurity* defined as follows:

Consider a decision tree T trained on a dataset \mathcal{D} using the ID3.

For every node τ of \mathcal{T} , denote by $\mathcal{D}[\tau]$ the subset of \mathcal{D} which was used in the ID3 procedure when the node τ was created (line 2).

Consider an attribute A and denote by $T[A] \subseteq T_{int}$ the set of all nodes of \mathcal{T} assigned the attribute A by ID3 (line 11).

Attribute Importance Computation

How important are attributes for the trained tree \mathcal{T} ? Depends on

- ▶ how close they are to the root of \mathcal{T} ,
- ▶ how large information gain/decrease in impurity they give.

There are several formulae for computing the importance.

One example is *mean decrease impurity* defined as follows:

Consider a decision tree T trained on a dataset \mathcal{D} using the ID3.

For every node τ of \mathcal{T} , denote by $\mathcal{D}[\tau]$ the subset of \mathcal{D} which was used in the ID3 procedure when the node τ was created (line 2).

Consider an attribute A and denote by $T[A] \subseteq T_{int}$ the set of all nodes of \mathcal{T} assigned the attribute A by ID3 (line 11).

Then define the importance as the average decrease in Gini impurity (i.e., average *ImpDec*) in the nodes of $T[A]$:

$$GiniImportance(A) = \sum_{\tau \in T[A]} \frac{|\mathcal{D}[\tau]|}{|\mathcal{D}|} ImpDec(\mathcal{D}[\tau], A)$$

Continuous-Valued Attributes

What if values of an attribute A come from a continuous variable?

A is a numerical attribute that can take any value in an interval, such as temperature, size, time, etc.

Continuous-Valued Attributes

What if values of an attribute A come from a continuous variable?

A is a numerical attribute that can take any value in an interval, such as temperature, size, time, etc.

Consider an internal node $\tau \in T_{int}$ assigned such a continuous attribute A . Then

- ▶ τ is assigned a threshold value called a *cut point* $H \in \mathbb{R}$,
- ▶ there are two edges $e_{\text{true}}, e_{\text{false}}$ from τ ,
- ▶ e_{true} labeled with True and e_{false} labeled with False.

Continuous-Valued Attributes

What if values of an attribute A come from a continuous variable?

A is a numerical attribute that can take any value in an interval, such as temperature, size, time, etc.

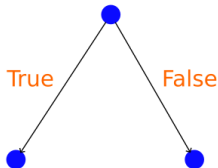
Consider an internal node $\tau \in T_{int}$ assigned such a continuous attribute A . Then

- ▶ τ is assigned a threshold value called a *cut point* $H \in \mathbb{R}$,
- ▶ there are two edges $e_{\text{true}}, e_{\text{false}}$ from τ ,
- ▶ e_{true} labeled with True and e_{false} labeled with False.

During inference, when considering an example \vec{x} in the node τ ,

- ▶ evaluate $A(\vec{x}) \leq H$,
- ▶ if $A(\vec{x}) \leq H$, then follow e_{true} ,
- ▶ else follow e_{false} .

Temperature ≤ 15



In training, the cut point is chosen from the attribute values in the training set using information gain/impurity decrease similar to discrete attributes.

Iris Example

iris setosa



petal

sepal

iris versicolor



petal

sepal

iris virginica



petal

sepal

Attributes

Sepal.Length, Sepal.Width, Petal.Length, Petal.Width

Classes (Variety)

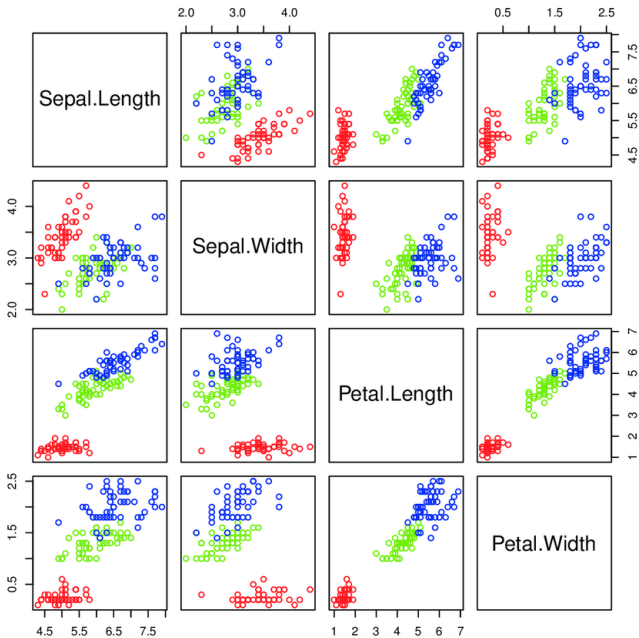
Setosa, Versicolor, Virginica

Iris Example

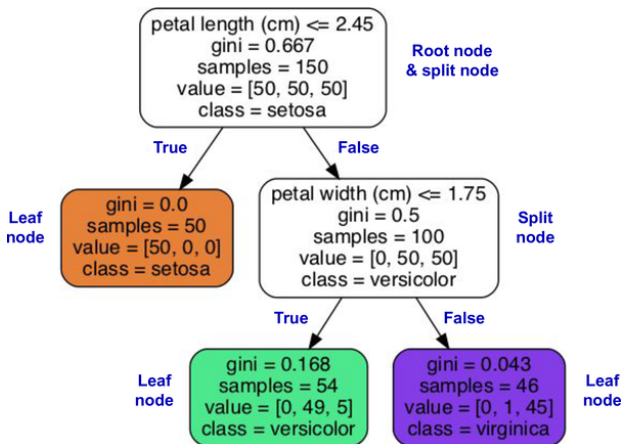
The dataset (150 examples):

Sepal.Length	Sepal.Width	Petal.Length	Petal.Width	Variety
5.5	3.5	1.3	0.2	Setosa
6.8	2.8	4.8	1.4	Versicolor
6.7	3.1	4.7	1.5	Versicolor
6.9	3.1	5.1	2.3	Virginica
7.3	2.9	6.3	1.8	Virginica
5.4	3.7	1.5	0.2	Setosa
4.6	3.4	1.4	0.3	Setosa
6.2	2.8	4.8	1.8	Virginica
5.4	3.0	4.5	1.5	Versicolor
4.7	3.2	1.6	0.2	Setosa
6.7	3.3	5.7	2.1	Virginica
5.0	3.4	1.5	0.2	Setosa
5.0	3.0	1.6	0.2	Setosa
4.4	2.9	1.4	0.2	Setosa
6.0	3.4	4.5	1.6	Versicolor
5.1	3.5	1.4	0.2	Setosa
6.6	3.0	4.4	1.4	Versicolor
5.9	3.2	4.8	1.8	Versicolor
5.6	2.8	4.9	2.0	Virginica
...				

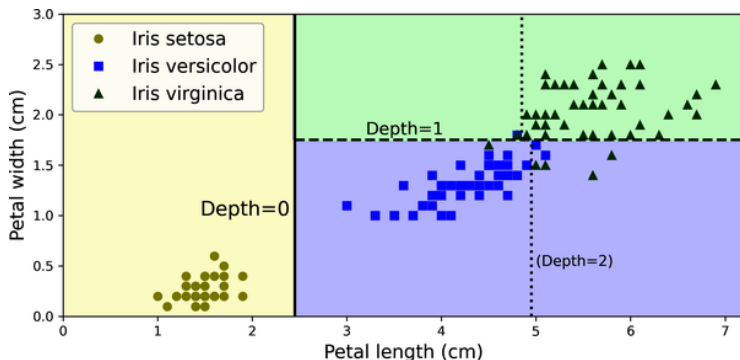
Iris Example



Iris Example - Decision Tree

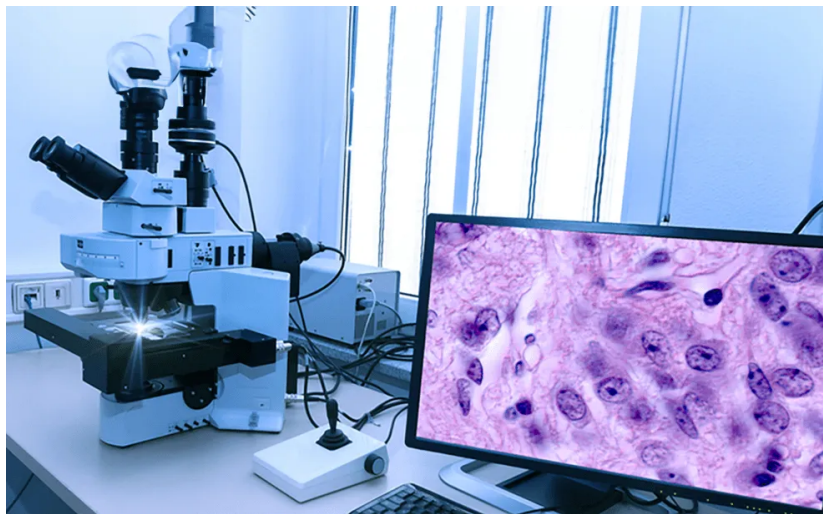


Iris Example - Decision Tree Boudaries

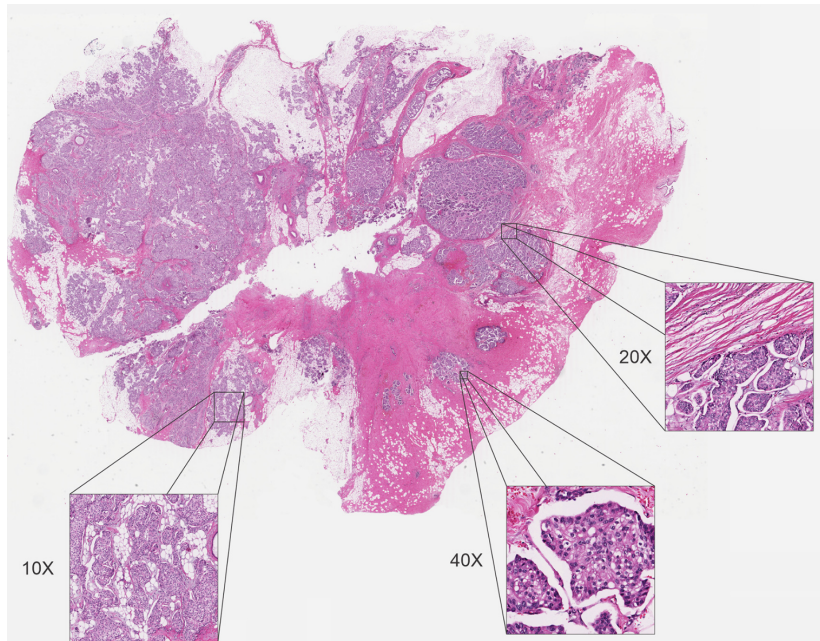


If the leaves are split further, the Depth = 2 boundary would be added.

Example: Wisconsin Breast Cancer Dataset



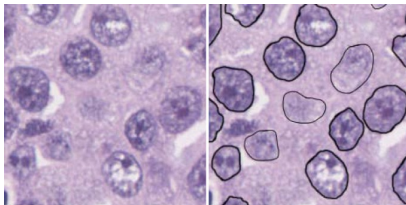
Wisconsin Breast Cancer Dataset



Example: Wisconsin Breast Cancer Dataset

- ▶ The Wisconsin Breast Cancer Dataset (WBCD) is used for binary classification of tumors.
- ▶ Origin: Created by Dr. William H. Wolberg at the University of Wisconsin-Madison.
- ▶ **Samples:** 569 breast cancer cases.
- ▶ **Target Variable:**
 - ▶ 0 = Benign (non-cancerous)
 - ▶ 1 = Malignant (cancerous)
- ▶ **Features:** 30 numerical attributes describing cell nuclei.

Example: Wisconsin Breast Cancer Dataset



- ▶ Each sample is characterized by 10 different measurements:
 - ▶ Radius, Texture, Perimeter, Area
 - ▶ Smoothness, Compactness, Concavity
 - ▶ Concave Points, Symmetry, Fractal Dimension
- ▶ Each feature has three statistical measures:
 - ▶ Mean (Average over cells in the sample)
 - ▶ Standard Error (SE) (Variation within the sample)
 - ▶ Worst (Largest value in the sample)

	id	diagnosis	radius_mean	texture_mean	perimeter_mean	area_mean	smoothness_mean	compactness_mean	concavity_mean	concave points_mean	...
0	842302	M	17.99	10.38	122.80	1001.0	0.11840	0.27760	0.3001	0.14710	...
1	842517	M	20.57	17.77	132.90	1326.0	0.08474	0.07864	0.0869	0.07017	...
2	84300903	M	19.69	21.25	130.00	1203.0	0.10960	0.15990	0.1974	0.12790	...
3	84348301	M	11.42	20.38	77.58	386.1	0.14250	0.28390	0.2414	0.10520	...
4	84358402	M	20.29	14.34	135.10	1297.0	0.10030	0.13280	0.1980	0.10430	...

Example: Wisconsin Breast Cancer Dataset

Just asked ChatGPT: Train a decision tree predicting the diagnosis and show me graphically the result.

Example: Wisconsin Breast Cancer Dataset

Just asked ChatGPT: Train a decision tree predicting the diagnosis and show me graphically the result.

I'll train a decision tree classifier to predict the **diagnosis** from your **cleaned_breast_cancer_dataset.csv** and visualize the tree. Let me first inspect the dataset to confirm its structure.

The dataset has **569 samples** and **32 columns**, including an **ID column** and a **diagnosis column (target variable)**. The diagnosis is encoded as **1 (Malignant)** and **0 (Benign)**.

I will:

1. Remove the **ID** column.
2. Train a **Decision Tree Classifier** on the remaining features.
3. Visualize the trained tree.

(see the result shown separately; too large for a slide)

Training Accuracy: 0.9956

Test Accuracy: 0.9474

```

import pandas as pd
import matplotlib.pyplot as plt
from sklearn.model_selection import train_test_split
from sklearn.tree import DecisionTreeClassifier, plot_tree
from sklearn.metrics import accuracy_score

# Load dataset
file_path = "cleaned_breast_cancer_dataset.csv"
df = pd.read_csv(file_path)

# Remove ID column
df = df.drop(columns=["id"])

# Split data into features and target
X = df.drop(columns=["diagnosis"])
y = df["diagnosis"]

# Split into training and test sets
X_train, X_test, y_train, y_test = train_test_split(
    X, y, test_size=0.2, random_state=42)

# Train decision tree classifier
clf = DecisionTreeClassifier(max_depth=4, random_state=42)
clf.fit(X_train, y_train)

# Evaluate model
train_accuracy = accuracy_score(y_train, clf.predict(X_train))
test_accuracy = accuracy_score(y_test, clf.predict(X_test))

# Print accuracy
print(f"Training Accuracy: {train_accuracy:.4f}")
print(f"Test Accuracy: {test_accuracy:.4f}")

# Plot decision tree
plt.figure(figsize=(20,10))
plot_tree(clf, feature_names=X.columns, class_names=["Benign", "Malignant"],
          filled=True, rounded=True)
plt.show()

```

Decision Trees

Practical Issues

Practical Issues

- ▶ Data preprocessing
- ▶ Model tuning (overfitting and underfitting)
- ▶ Sensitivity to changes in data/hyperparameters
- ▶ Learning representation problems (the XOR)

Data Preprocessing

Little preprocessing is needed for decision trees.

Data Preprocessing

Little preprocessing is needed for decision trees.

Of course, ideally, clean up the data, but

- ▶ Missing values are not such a big issue
(considering a concrete example, exclude the attributes with missing values)

Data Preprocessing

Little preprocessing is needed for decision trees.

Of course, ideally, clean up the data, but

- ▶ Missing values are not such a big issue
(considering a concrete example, exclude the attributes with missing values)
- ▶ Outliers are not such a big issue either; the division of nodes is done based on relative counts, not concrete values.

Data Preprocessing

Little preprocessing is needed for decision trees.

Of course, ideally, clean up the data, but

- ▶ Missing values are not such a big issue
(considering a concrete example, exclude the attributes with missing values)
- ▶ Outliers are not such a big issue either; the division of nodes is done based on relative counts, not concrete values.
- ▶ Decision trees can cope with continuous and categorical values directly (i.e., no encoding necessary)

Data Preprocessing

Little preprocessing is needed for decision trees.

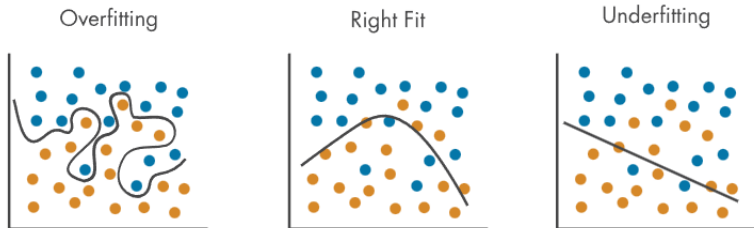
Of course, ideally, clean up the data, but

- ▶ Missing values are not such a big issue (considering a concrete example, exclude the attributes with missing values)
- ▶ Outliers are not such a big issue either; the division of nodes is done based on relative counts, not concrete values.
- ▶ Decision trees can cope with continuous and categorical values directly (i.e., no encoding necessary)

Imbalanced classes might cause problems because of small information gain/impurity decrease in splitting.

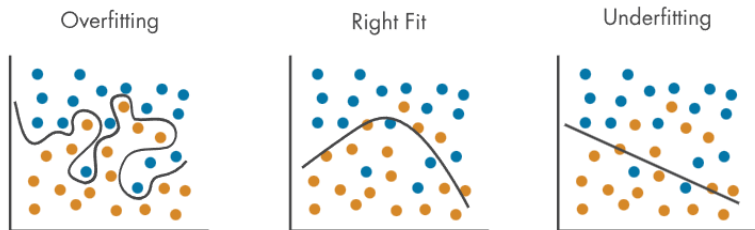
Model Tuning - Over/Under Fitting

The behavior of the model on the training set:



Model Tuning - Over/Under Fitting

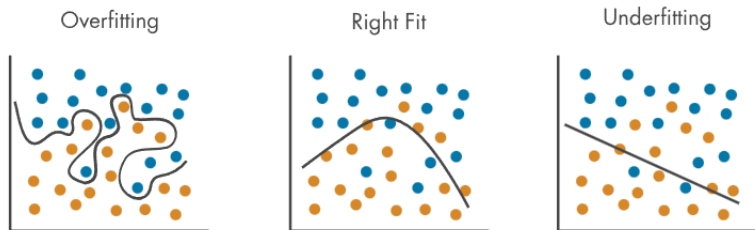
The behavior of the model on the training set:



- ▶ The left one is strongly overfitting. It would possibly not work well on new data.

Model Tuning - Over/Under Fitting

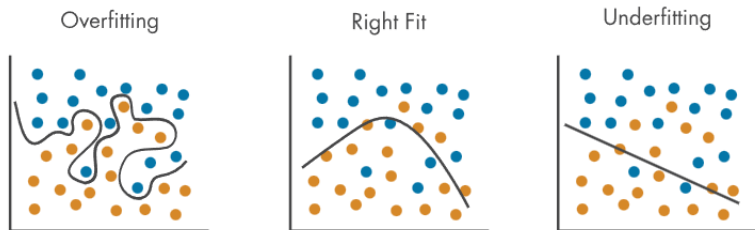
The behavior of the model on the training set:



- ▶ The left one is strongly overfitting. It would possibly not work well on new data.
- ▶ The right one is strongly underfitting. It would probably give poor classification results.

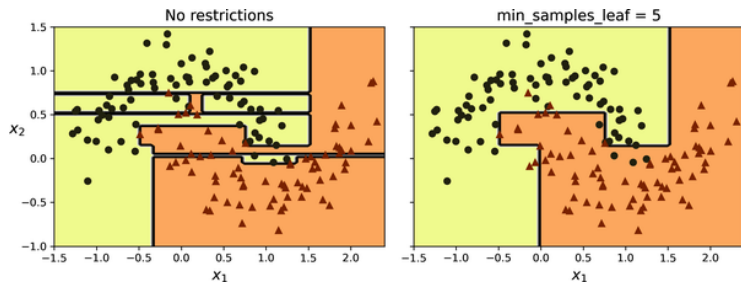
Model Tuning - Over/Under Fitting

The behavior of the model on the training set:



- ▶ The left one is strongly overfitting. It would possibly not work well on new data.
- ▶ The right one is strongly underfitting. It would probably give poor classification results.
- ▶ The middle one seems good (but still needs to be tested on fresh data).

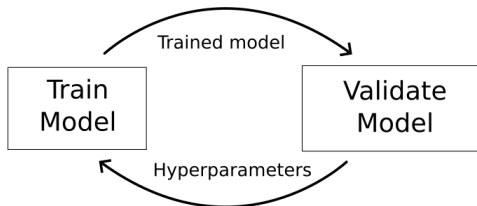
Model Tuning - Overfitting in Decision Trees



See the overfitting on the left and the “nice” model on the right. Both overfitting and underfitting are best avoided. But how do we find out?

Model Tuning (In General)

Recall from the first lecture:

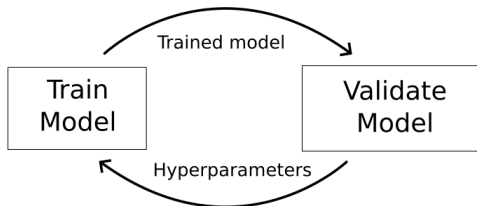


The validation should be done on a **validation set** separated from the training set.

We will discuss more sophisticated techniques later.

Model Tuning (In General)

Recall from the first lecture:



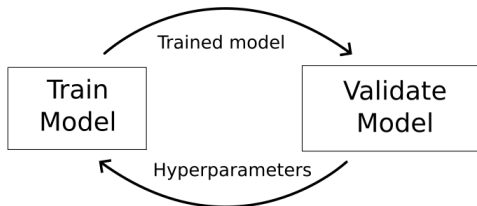
The validation should be done on a **validation set** separated from the training set.

We will discuss more sophisticated techniques later.

What hyperparameters to set? (see the next slide)

Model Tuning (In General)

Recall from the first lecture:



The validation should be done on a **validation set** separated from the training set.

We will discuss more sophisticated techniques later.

What hyperparameters to set? (see the next slide)

What to observe? In the case of decision trees, one should observe the difference between performance measures (e.g., classification accuracy) on the training and validation sets.

The too-large difference implies an improperly fitting model.

How to Fit Decision Trees?

There are several approaches available for decision trees.

Generally, the overfitting can be either prevented or resolved.

How to Fit Decision Trees?

There are several approaches available for decision trees.

Generally, the overfitting can be either prevented or resolved.

- ▶ **Pre-pruning:** Build the tree so it does not overfit by restricting its size.

How to Fit Decision Trees?

There are several approaches available for decision trees.

Generally, the overfitting can be either prevented or resolved.

- ▶ **Pre-pruning**: Build the tree so it does not overfit by restricting its size.
- ▶ **Post-pruning**: Overfit with a large tree and remove subtrees to make it smaller.

How to Fit Decision Trees?

There are several approaches available for decision trees.

Generally, the overfitting can be either prevented or resolved.

- ▶ **Pre-pruning**: Build the tree so it does not overfit by restricting its size.
- ▶ **Post-pruning**: Overfit with a large tree and remove subtrees to make it smaller.
- ▶ **Ensemble methods**: Fit several different trees and let them classify together (e.g., using majority voting).

How to Fit Decision Trees?

There are several approaches available for decision trees.

Generally, the overfitting can be either prevented or resolved.

- ▶ **Pre-pruning**: Build the tree so it does not overfit by restricting its size.
- ▶ **Post-pruning**: Overfit with a large tree and remove subtrees to make it smaller.
- ▶ **Ensemble methods**: Fit several different trees and let them classify together (e.g., using majority voting).

The post-pruning approach has been more successful in practice than the pre-pruning because it is usually hard to say when to stop growing the tree.

We shall meet this controversy also in deep learning, where recent history shows a similar phenomenon.

The ensemble methods will be covered later when we discuss random forests.

Pre-Pruning - Hyperparameters

Hyperparameters controlling the size of the tree:

- ▶ Maximum depth - do not grow the tree beyond the max depth
The deeper the tree, the more complex models you can create \Rightarrow overfitting. Low depth may restrict expressivity.

Pre-Pruning - Hyperparameters

Hyperparameters controlling the size of the tree:

- ▶ Maximum depth - do not grow the tree beyond the max depth
The deeper the tree, the more complex models you can create \Rightarrow overfitting. Low depth may restrict expressivity.
- ▶ Minimum number of examples to split a node - if $\mathcal{D}[\tau]$ is small, τ becomes a leaf (labeled with the majority class)
Higher values prevent a model from learning relations specific only for a few examples. Too high values may result in underfitting.

Pre-Pruning - Hyperparameters

Hyperparameters controlling the size of the tree:

- ▶ Maximum depth - do not grow the tree beyond the max depth
The deeper the tree, the more complex models you can create \Rightarrow overfitting. Low depth may restrict expressivity.
- ▶ Minimum number of examples to split a node - if $\mathcal{D}[\tau]$ is small, τ becomes a leaf (labeled with the majority class)
Higher values prevent a model from learning relations specific only for a few examples. Too high values may result in underfitting.
- ▶ Minimum number of examples required to be in a leaf
Similar to the previous one. A higher number means we cannot have very specific branches concerned with particular combinations of values.

Pre-Pruning - Hyperparameters

Hyperparameters controlling the size of the tree:

- ▶ Maximum depth - do not grow the tree beyond the max depth
The deeper the tree, the more complex models you can create \Rightarrow overfitting. Low depth may restrict expressivity.
- ▶ Minimum number of examples to split a node - if $\mathcal{D}[\tau]$ is small, τ becomes a leaf (labeled with the majority class)
Higher values prevent a model from learning relations specific only for a few examples. Too high values may result in underfitting.
- ▶ Minimum number of examples required to be in a leaf
Similar to the previous one. A higher number means we cannot have very specific branches concerned with particular combinations of values.
- ▶ Minimum information gain/impurity decrease
A small impurity decrease means that the split does not contribute too much to the classification (their proportions after a split are similar to proportions before a split). However, keep in mind that it is *weighted average impurity* after the split.

Post-Pruning - Reduced Error Pruning

Train a large tree and then remove nodes that make classification worse on the validation set.

Post-Pruning - Reduced Error Pruning

Train a large tree and then remove nodes that make classification worse on the validation set.

Given a decision tree \mathcal{T} and its internal node $\tau \in T_{int}$, we denote by $\mathcal{T}_{-\tau}$ the tree obtained from \mathcal{T} by removing the subtree rooted in τ , i.e., τ is a leaf of $\mathcal{T}_{-\tau}$.

Post-Pruning - Reduced Error Pruning

Train a large tree and then remove nodes that make classification worse on the validation set.

Given a decision tree \mathcal{T} and its internal node $\tau \in T_{int}$, we denote by $\mathcal{T}_{-\tau}$ the tree obtained from \mathcal{T} by removing the subtree rooted in τ , i.e., τ is a leaf of $\mathcal{T}_{-\tau}$.

- 1: Train \mathcal{T} to maximum fit on the *training dataset*.
- 2: **while** true **do**
- 3: $Err[\mathcal{T}] \leftarrow$ the error of \mathcal{T} on the validation set.
- 4: **for** $\tau \in T_{int}$ **do**
- 5: $Err[\mathcal{T}_{-\tau}] \leftarrow$ the error of $\mathcal{T}_{-\tau}$ on the validation set.
- 6: **end for**
- 7: **if** $Err[\mathcal{T}] \leq \min\{Err[\mathcal{T}_{-\tau}] \mid \tau \in T_{int}\}$ **then return** \mathcal{T}
- 8: **else**
- 9: $\mathcal{T} \leftarrow \operatorname{argmin}\{Err[\mathcal{T}_{-\tau}] \mid \tau \in T_{int}\}$
- 10: **end if**
- 11: **end while**

The error $Err[\mathcal{T}]$ can be any measure of the “badness” of the decision tree \mathcal{T} . For example, $1 - Accuracy$.

Other Pruning Methods

There are more pruning methods.

- ▶ Rule Post-Pruning:
 - ▶ Transform the tree into a set of rules.
Rules correspond to paths in the tree; they have a form of implication: Specific values of attributes imply a class.
 - ▶ Remove the attribute conditions from the premises of the implications.

This gives a more refined pruning: Instead of removing the whole subtree, each path of the subtree can be pruned individually.

Other Pruning Methods

There are more pruning methods.

- ▶ Rule Post-Pruning:

- ▶ Transform the tree into a set of rules.

Rules correspond to paths in the tree; they have a form of implication: Specific values of attributes imply a class.

- ▶ Remove the attribute conditions from the premises of the implications.

This gives a more refined pruning: Instead of removing the whole subtree, each path of the subtree can be pruned individually.

- ▶ Using cost complexity measure: Evaluate trees not only based on the classification error but also based on their size.

Other Pruning Methods

There are more pruning methods.

- ▶ Rule Post-Pruning:
 - ▶ Transform the tree into a set of rules.
Rules correspond to paths in the tree; they have a form of implication: Specific values of attributes imply a class.
 - ▶ Remove the attribute conditions from the premises of the implications.

This gives a more refined pruning: Instead of removing the whole subtree, each path of the subtree can be pruned individually.

- ▶ Using cost complexity measure: Evaluate trees not only based on the classification error but also based on their size.

Typically introduce regularization into the error functions:
Given a decision tree \mathcal{T}

$$Err_{\alpha}(\mathcal{T}) = Err(\mathcal{T}) + \alpha|\mathcal{T}|$$

The original paper by Breiman et al. (1984) defined $Err(\mathcal{T})$ to be the misclassification rate on the training dataset, and $|\mathcal{T}|$ is the number of nodes of the tree \mathcal{T} .

Sensitivity to Small Changes and Randomness

- ▶ Decision trees are sensitive to small changes in data and hyperparameters.
Several attributes may provide (almost) identical information gain but divide the training dataset very differently.
- ▶ Some implementations choose attributes partially in random (sci-kit-learn). You may get completely different trees.

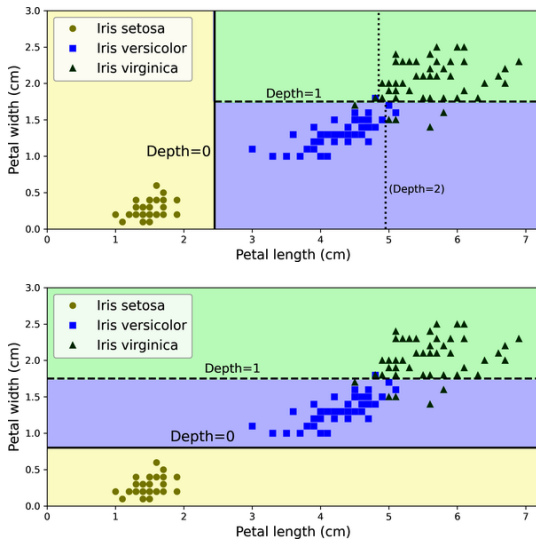
Sensitivity to Small Changes and Randomness

- ▶ Decision trees are sensitive to small changes in data and hyperparameters.
Several attributes may provide (almost) identical information gain but divide the training dataset very differently.
- ▶ Some implementations choose attributes partially in random (sci-kit-learn). You may get completely different trees.

A solution is to train an ensemble of many decision trees and then use majority voting for classification.

This is the fundamental idea behind random forests (see later lectures).

Iris - Illustration

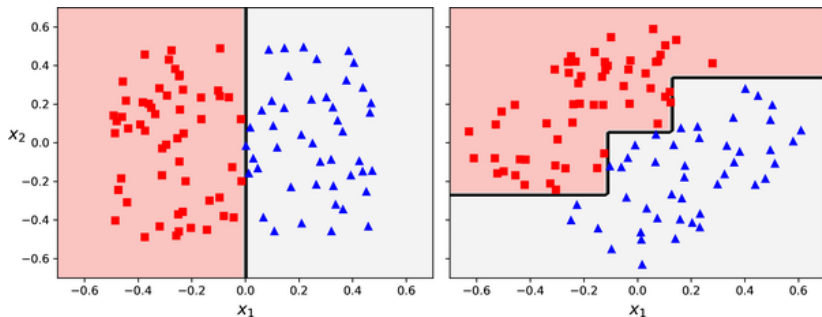


Decision trees trained on the Iris dataset.

Iris Setosa is perfectly separated by many choices for the first split.

Axis Sensitivity

The decision makes divisions along particular axes:

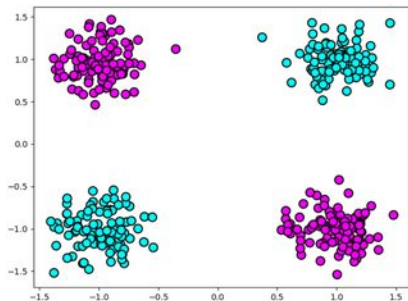


That is, rotated data may result in a completely different model.

That is why decision trees are often preceded by the *principal component analysis (PCA)* transformation, which aligns data along the axes of maximum data variance.

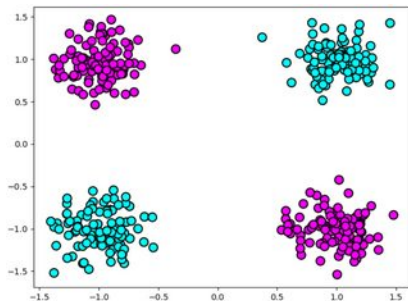
XOR Training Problem

Consider the following training dataset:

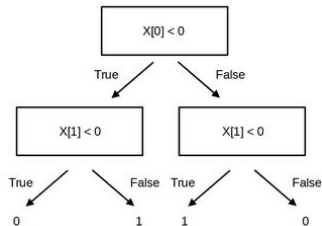


XOR Training Problem

Consider the following training dataset:

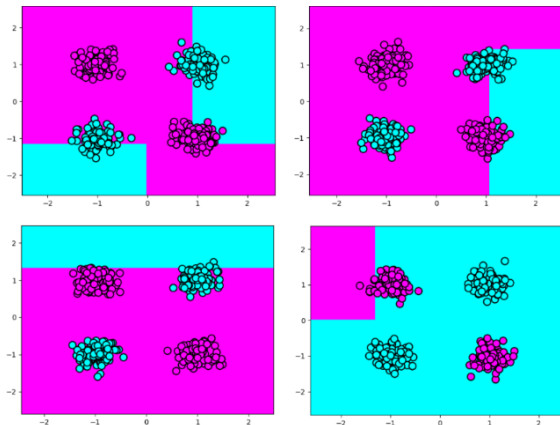


An ideal decision tree would look like this:



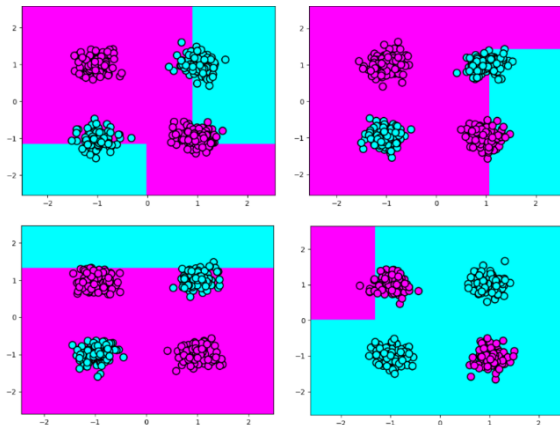
Attempts at Training on XOR

Max depth = 2:



Attempts at Training on XOR

Max depth = 2:

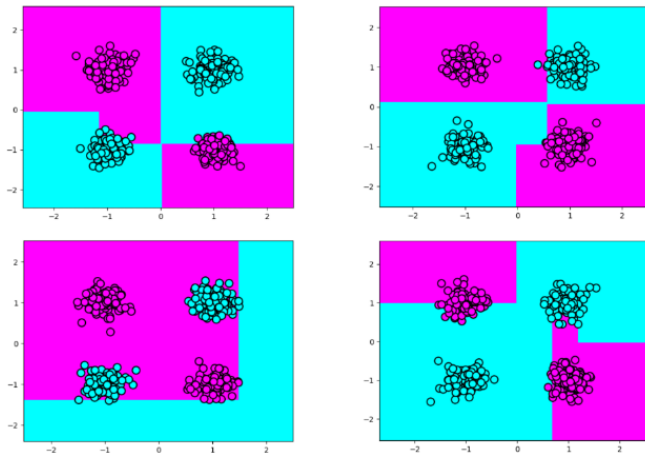


The problem: Both information gain and decrease in impurity consider only the relationship of a *single* attribute and the class.

However, there is no relationship between a single attribute and the class; both attributes need to be considered together!

More Attempts at Training on XOR

Max depth = 3:



It's better but still fails occasionally.

Advantages of Decision Trees

- ▶ Simple to understand and interpret; trees can be visualized.
- ▶ Uses a white box model, where conditions are easily explained by boolean logic.

Advantages of Decision Trees

- ▶ Simple to understand and interpret; trees can be visualized.
- ▶ Uses a white box model, where conditions are easily explained by boolean logic.
- ▶ Can approximate an arbitrary (reasonable) boundary and capture complex non-linear relationships between attributes.
- ▶ Capable of handling multi-class problems.

Advantages of Decision Trees

- ▶ Simple to understand and interpret; trees can be visualized.
- ▶ Uses a white box model, where conditions are easily explained by boolean logic.
- ▶ Can approximate an arbitrary (reasonable) boundary and capture complex non-linear relationships between attributes.
- ▶ Capable of handling multi-class problems.
- ▶ Little data preparation, unlike other techniques requiring normalization, dummy variables, or missing value removal.
- ▶ Handles numerical and categorical data.
- ▶ Not sensitive to outliers since the splitting is based on the proportion of examples within the split ranges and not on absolute values.

Advantages of Decision Trees

- ▶ Simple to understand and interpret; trees can be visualized.
- ▶ Uses a white box model, where conditions are easily explained by boolean logic.
- ▶ Can approximate an arbitrary (reasonable) boundary and capture complex non-linear relationships between attributes.
- ▶ Capable of handling multi-class problems.
- ▶ Little data preparation, unlike other techniques requiring normalization, dummy variables, or missing value removal.
- ▶ Handles numerical and categorical data.
- ▶ Not sensitive to outliers since the splitting is based on the proportion of examples within the split ranges and not on absolute values.
- ▶ The cost of using a well-balanced tree is logarithmic in the number of data points used to train it.

Disadvantages of Decision Trees

- ▶ Overfitting: Trees can be over-complex and not generalize well, needing pruning or limits on tree depth.
- ▶ Instability: Small data variations can result in very different trees. This is mitigated in ensemble methods.
- ▶ Non-smooth predictions: Decision trees make piecewise constant approximations, which are not suitable for extrapolation.

Disadvantages of Decision Trees

- ▶ Overfitting: Trees can be over-complex and not generalize well, needing pruning or limits on tree depth.
- ▶ Instability: Small data variations can result in very different trees. This is mitigated in ensemble methods.
- ▶ Non-smooth predictions: Decision trees make piecewise constant approximations, which are not suitable for extrapolation.
- ▶ Difficulty learning certain concepts, such as XOR.
- ▶ Bias in trees: Decision trees can create biased trees if some classes dominate. Balancing the dataset is recommended.

Disadvantages of Decision Trees

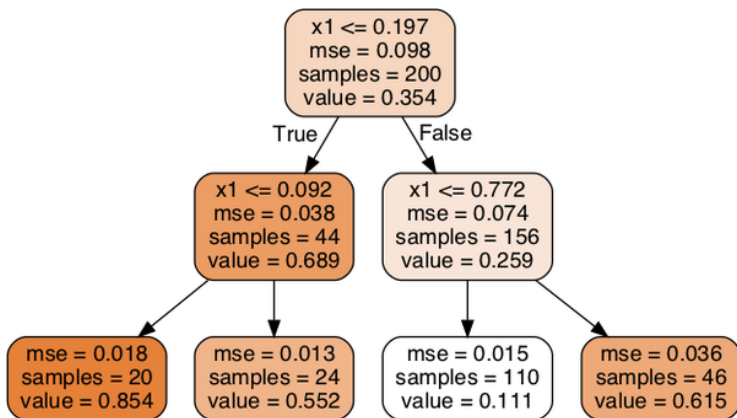
- ▶ Overfitting: Trees can be over-complex and not generalize well, needing pruning or limits on tree depth.
- ▶ Instability: Small data variations can result in very different trees. This is mitigated in ensemble methods.
- ▶ Non-smooth predictions: Decision trees make piecewise constant approximations, which are not suitable for extrapolation.
- ▶ Difficulty learning certain concepts, such as XOR.
- ▶ Bias in trees: Decision trees can create biased trees if some classes dominate. Balancing the dataset is recommended.
- ▶ Learning optimal trees is NP-complete: Heuristic algorithms like greedy algorithms are used, which do not guarantee globally optimal trees. Ensemble methods can help.

History of Decision Trees

- ▶ Hunt and colleagues use exhaustive search decision-tree methods (CLS) to model human concept learning in the 1960's.
- ▶ In the late 70's, Quinlan developed ID3 with the information gain heuristic to learn expert systems from examples.
- ▶ Simultaneously, Breiman, Friedman, and colleagues develop CART (Classification and Regression Trees), similar to ID3.
- ▶ In the 1980s, various improvements were introduced to handle noise, continuous features, missing features, and improved splitting criteria. Various expert-system development tools results.
- ▶ Quinlan's updated decision-tree package (C4.5) released in 1993.

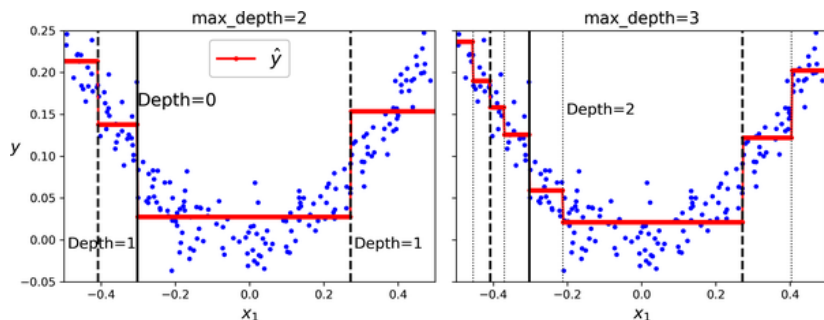
Comment on Regression Trees

Decision trees can also be used to approximate functions. Assign a function value to the leaves instead of classes.



Here, “mse” is the mean-squared-error.

Comment on Regression Trees



Intuitively, for every subinterval of x_1 , the value (the red line) is at the average y over the subinterval.

How are the subintervals being set?

Regression Trees

A *regression tree* is a decision tree whose leaves are labeled by values from \mathbb{R} .

We follow the same procedure as in decision trees during inference on an input \vec{x} .

Regression Trees

A *regression tree* is a decision tree whose leaves are labeled by values from \mathbb{R} .

We follow the same procedure as in decision trees during inference on an input \vec{x} .

How exactly are such trees trained?

We aim to minimize the mean squared error.

Regression Trees

A *regression tree* is a decision tree whose leaves are labeled by values from \mathbb{R} .

We follow the same procedure as in decision trees during inference on an input \vec{x} .

How exactly are such trees trained?

We aim to minimize the mean squared error.

Assume, for the moment, that we want to train a single-node tree. What will be the best labeling value for the single node?

Regression Trees

A *regression tree* is a decision tree whose leaves are labeled by values from \mathbb{R} .

We follow the same procedure as in decision trees during inference on an input \vec{x} .

How exactly are such trees trained?

We aim to minimize the mean squared error.

Assume, for the moment, that we want to train a single-node tree. What will be the best labeling value for the single node?

Now, consider the training for arbitrary trees.

Assume ordinal attributes.

The algorithm also works for discrete attributes; ordinal attributes, however, allow us to make binary splits.

Regression Trees

A *regression tree* is a decision tree whose leaves are labeled by values from \mathbb{R} .

We follow the same procedure as in decision trees during inference on an input \vec{x} .

How exactly are such trees trained?

We aim to minimize the mean squared error.

Assume, for the moment, that we want to train a single-node tree. What will be the best labeling value for the single node?

Now, consider the training for arbitrary trees.

Assume ordinal attributes.

The algorithm also works for discrete attributes; ordinal attributes, however, allow us to make binary splits.

The training procedure is the same as for the decision trees, except that the splits and cut points are selected differently.

Regression Trees

Given a dataset $\mathcal{D} = \{(\vec{x}_1, d_1), \dots, (\vec{x}_p, d_p)\}$, we denote by $\bar{\mathcal{D}}$ the average *desired* value in \mathcal{D} , that is $\bar{\mathcal{D}} = \frac{1}{p} \sum_{k=1}^p d_k$.

Regression Trees

Given a dataset $\mathcal{D} = \{(\vec{x}_1, d_1), \dots, (\vec{x}_p, d_p)\}$, we denote by $\bar{\mathcal{D}}$ the average *desired* value in \mathcal{D} , that is $\bar{\mathcal{D}} = \frac{1}{p} \sum_{k=1}^p d_k$.

Consider a dataset \mathcal{D} and let us find a cut point for A in \mathcal{D} .

Regression Trees

Given a dataset $\mathcal{D} = \{(\vec{x}_1, d_1), \dots, (\vec{x}_p, d_p)\}$, we denote by \bar{D} the average *desired* value in \mathcal{D} , that is $\bar{D} = \frac{1}{p} \sum_{k=1}^p d_k$.

Consider a dataset \mathcal{D} and let us find a cut point for A in \mathcal{D} .

We are looking for a value H of the attribute A such that the split:

$$\mathcal{D}_{\leq H} = \{(\vec{x}, d) \in \mathcal{D} \mid A(\vec{x}) \leq H\} \quad \mathcal{D}_{>H} = \{(\vec{x}, d) \in \mathcal{D} \mid A(\vec{x}) > H\}$$

Minimizes the following *split error*:

$$\frac{1}{|\mathcal{D}_{\leq H}|} \sum_{(\vec{x}, d) \in \mathcal{D}_{\leq H}} (d - \bar{D}_{\leq H})^2 + \frac{1}{|\mathcal{D}_{>H}|} \sum_{(\vec{x}, d) \in \mathcal{D}_{>H}} (d - \bar{D}_{>H})^2$$

Regression Trees

Given a dataset $\mathcal{D} = \{(\vec{x}_1, d_1), \dots, (\vec{x}_p, d_p)\}$, we denote by \bar{D} the average *desired* value in \mathcal{D} , that is $\bar{D} = \frac{1}{p} \sum_{k=1}^p d_k$.

Consider a dataset \mathcal{D} and let us find a cut point for A in \mathcal{D} .

We are looking for a value H of the attribute A such that the split:

$$\mathcal{D}_{\leq H} = \{(\vec{x}, d) \in \mathcal{D} \mid A(\vec{x}) \leq H\} \quad \mathcal{D}_{> H} = \{(\vec{x}, d) \in \mathcal{D} \mid A(\vec{x}) > H\}$$

Minimizes the following *split error*:

$$\frac{1}{|\mathcal{D}_{\leq H}|} \sum_{(\vec{x}, d) \in \mathcal{D}_{\leq H}} (d - \bar{D}_{\leq H})^2 + \frac{1}{|\mathcal{D}_{> H}|} \sum_{(\vec{x}, d) \in \mathcal{D}_{> H}} (d - \bar{D}_{> H})^2$$

Denote by $\text{MinE}(\mathcal{D})$ the minimum of the split error over all attributes A and all cut points H .

Regression Trees

Given a dataset $\mathcal{D} = \{(\vec{x}_1, d_1), \dots, (\vec{x}_p, d_p)\}$, we denote by \bar{D} the average *desired* value in \mathcal{D} , that is $\bar{D} = \frac{1}{p} \sum_{k=1}^p d_k$.

Consider a dataset \mathcal{D} and let us find a cut point for A in \mathcal{D} .

We are looking for a value H of the attribute A such that the split:

$$\mathcal{D}_{\leq H} = \{(\vec{x}, d) \in \mathcal{D} \mid A(\vec{x}) \leq H\} \quad \mathcal{D}_{> H} = \{(\vec{x}, d) \in \mathcal{D} \mid A(\vec{x}) > H\}$$

Minimizes the following *split error*:

$$\frac{1}{|\mathcal{D}_{\leq H}|} \sum_{(\vec{x}, d) \in \mathcal{D}_{\leq H}} (d - \bar{D}_{\leq H})^2 + \frac{1}{|\mathcal{D}_{> H}|} \sum_{(\vec{x}, d) \in \mathcal{D}_{> H}} (d - \bar{D}_{> H})^2$$

Denote by $\text{MinE}(\mathcal{D})$ the minimum of the split error over all attributes A and all cut points H . Compute

$$\Delta = \left(\frac{1}{|\mathcal{D}|} \sum_{(\vec{x}, d) \in \mathcal{D}} (d - \bar{D})^2 \right) - \text{MinE}$$

Regression Trees

Given a dataset $\mathcal{D} = \{(\vec{x}_1, d_1), \dots, (\vec{x}_p, d_p)\}$, we denote by \bar{D} the average *desired* value in \mathcal{D} , that is $\bar{D} = \frac{1}{p} \sum_{k=1}^p d_k$.

Consider a dataset \mathcal{D} and let us find a cut point for A in \mathcal{D} .

We are looking for a value H of the attribute A such that the split:

$$\mathcal{D}_{\leq H} = \{(\vec{x}, d) \in \mathcal{D} \mid A(\vec{x}) \leq H\} \quad \mathcal{D}_{> H} = \{(\vec{x}, d) \in \mathcal{D} \mid A(\vec{x}) > H\}$$

Minimizes the following *split error*:

$$\frac{1}{|\mathcal{D}_{\leq H}|} \sum_{(\vec{x}, d) \in \mathcal{D}_{\leq H}} (d - \bar{D}_{\leq H})^2 + \frac{1}{|\mathcal{D}_{> H}|} \sum_{(\vec{x}, d) \in \mathcal{D}_{> H}} (d - \bar{D}_{> H})^2$$

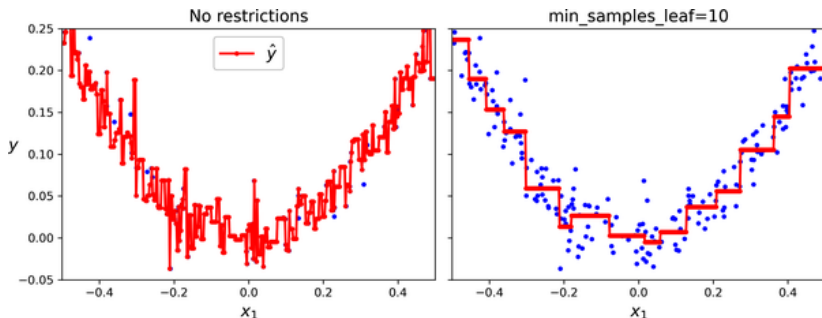
Denote by $\text{MinE}(\mathcal{D})$ the minimum of the split error over all attributes A and all cut points H . Compute

$$\Delta = \left(\frac{1}{|\mathcal{D}|} \sum_{(\vec{x}, d) \in \mathcal{D}} (d - \bar{D})^2 \right) - \text{MinE}$$

If Δ is large enough, split on A and H that minimize the split error. Otherwise, stop splitting and label the leaf with \bar{D} .

Regression Tress

Without any lower bound on the number of examples in the leaves, the algorithm will eventually overfit by splitting into (possibly) singleton leaves.



Probabilistic Classification

Probabilistic Classification – Idea

Imagine that

- ▶ I look out of a window and see a bird,
- ▶ it is black, approx. 25 cm long and has a rather yellow beak.

Probabilistic Classification – Idea

Imagine that

- ▶ I look out of a window and see a bird,
- ▶ it is black, approx. 25 cm long and has a rather yellow beak.

My daughter asks: What kind of bird is this?

Probabilistic Classification – Idea

Imagine that

- ▶ I look out of a window and see a bird,
- ▶ it is black, approx. 25 cm long and has a rather yellow beak.

My daughter asks: What kind of bird is this?

My usual answer: This is *probably* a kind of blackbird (kos černý in Czech).

Probabilistic Classification – Idea

Imagine that

- ▶ I look out of a window and see a bird,
- ▶ it is black, approx. 25 cm long and has a rather yellow beak.

My daughter asks: What kind of bird is this?

My usual answer: This is *probably* a kind of blackbird (kos černý in Czech).

Here *probably* means that out of my extensive catalog of four kinds of birds that I can recognize, "blackbird" gets the highest degree of belief based on *features* of this particular bird.

Frequentists might say that the largest proportion of birds with similar features I have ever seen were blackbirds.

Probabilistic Classification – Idea

Imagine that

- ▶ I look out of a window and see a bird,
- ▶ it is black, approx. 25 cm long and has a rather yellow beak.

My daughter asks: What kind of bird is this?

My usual answer: This is *probably* a kind of blackbird (kos černý in Czech).

Here *probably* means that out of my extensive catalog of four kinds of birds that I can recognize, "blackbird" gets the highest degree of belief based on *features* of this particular bird.

Frequentists might say that the largest proportion of birds with similar features I have ever seen were blackbirds.

The degree of belief (Bayesian), or the relative frequency (frequentists), is the *probability*.

Basic Discrete Probability Theory

- ▶ A finite or countably infinite set Ω of *possible outcomes*, Ω is called *sample space*.

Experiment: Roll one dice once. Sample space: $\Omega = \{1, \dots, 6\}$

Basic Discrete Probability Theory

- ▶ A finite or countably infinite set Ω of *possible outcomes*, Ω is called *sample space*.

Experiment: Roll one dice once. Sample space: $\Omega = \{1, \dots, 6\}$

- ▶ Each element ω of Ω is assigned a "probability" value $f(\omega)$, here f must satisfy
 - ▶ $f(\omega) \in [0, 1]$ for all $\omega \in \Omega$,
 - ▶ $\sum_{\omega \in \Omega} f(\omega) = 1$.

If the dice is fair, then $f(\omega) = \frac{1}{6}$ for all $\omega \in \{1, \dots, 6\}$.

Basic Discrete Probability Theory

- ▶ A finite or countably infinite set Ω of *possible outcomes*, Ω is called *sample space*.

Experiment: Roll one dice once. Sample space: $\Omega = \{1, \dots, 6\}$

- ▶ Each element ω of Ω is assigned a "probability" value $f(\omega)$, here f must satisfy
 - ▶ $f(\omega) \in [0, 1]$ for all $\omega \in \Omega$,
 - ▶ $\sum_{\omega \in \Omega} f(\omega) = 1$.

If the dice is fair, then $f(\omega) = \frac{1}{6}$ for all $\omega \in \{1, \dots, 6\}$.

- ▶ An *event* is any subset E of Ω .
- ▶ The *probability* of a given event $E \subseteq \Omega$ is defined as

$$P(E) = \sum_{\omega \in E} f(\omega)$$

Let E be the event that an odd number is rolled, i.e., $E = \{1, 3, 5\}$. Then $P(E) = \frac{1}{2}$.

Basic Discrete Probability Theory

- ▶ A finite or countably infinite set Ω of *possible outcomes*, Ω is called *sample space*.

Experiment: Roll one dice once. Sample space: $\Omega = \{1, \dots, 6\}$

- ▶ Each element ω of Ω is assigned a "probability" value $f(\omega)$, here f must satisfy
 - ▶ $f(\omega) \in [0, 1]$ for all $\omega \in \Omega$,
 - ▶ $\sum_{\omega \in \Omega} f(\omega) = 1$.

If the dice is fair, then $f(\omega) = \frac{1}{6}$ for all $\omega \in \{1, \dots, 6\}$.

- ▶ An *event* is any subset E of Ω .
- ▶ The *probability* of a given event $E \subseteq \Omega$ is defined as

$$P(E) = \sum_{\omega \in E} f(\omega)$$

Let E be the event that an odd number is rolled, i.e., $E = \{1, 3, 5\}$. Then $P(E) = \frac{1}{2}$.

- ▶ **Basic laws:** $P(\Omega) = 1$, $P(\emptyset) = 0$, given disjoint sets A, B we have $P(A \cup B) = P(A) + P(B)$, $P(\Omega \setminus A) = 1 - P(A)$.

Conditional Probability and Independence

- ▶ $P(A \mid B)$ is the probability of A given B (assume $P(B) > 0$) defined by

$$P(A \mid B) = P(A \cap B) / P(B)$$

(We assume that B is all and only information known.)

A fair dice: what is the probability that 3 is rolled assuming that an odd number is rolled? ... and assuming that an even number is rolled?

Conditional Probability and Independence

- ▶ $P(A | B)$ is the probability of A given B (assume $P(B) > 0$) defined by

$$P(A | B) = P(A \cap B) / P(B)$$

(We assume that B is all and only information known.)

A fair dice: what is the probability that 3 is rolled assuming that an odd number is rolled? ... and assuming that an even number is rolled?

- ▶ A and B are **independent** if $P(A \cap B) = P(A) \cdot P(B)$.

It is easy to show that if $P(B) > 0$, then

A, B are independent iff $P(A | B) = P(A)$.

Random Variables and Random Vectors

- ▶ A *random variable* X is a function $X : \Omega \rightarrow \mathbb{R}$.
A dice: $X : \{1, \dots, 6\} \rightarrow \{0, 1\}$ such that $X(n) = n \bmod 2$.
- ▶ A *random vector* is a function $X : \Omega \rightarrow \mathbb{R}^d$.

Random Variables and Random Vectors

- ▶ A *random variable* X is a function $X : \Omega \rightarrow \mathbb{R}$.

A dice: $X : \{1, \dots, 6\} \rightarrow \{0, 1\}$ such that $X(n) = n \bmod 2$.

- ▶ A *random vector* is a function $X : \Omega \rightarrow \mathbb{R}^d$.

We use $X = (X_1, \dots, X_d)$ where X_i is a random variable returning the i -th component of X .

- ▶ Consider random variables X_1, X_2 and Y . The variables X_1, X_2 are *conditionally independent given Y* if for all x_1, x_2 and y we have that

$$P(X_1 = x_1, X_2 = x_2 \mid Y = y) = \\ P(X_1 = x_1 \mid Y = y) \cdot P(X_2 = x_2 \mid Y = y)$$

Random Vectors – Example

Let Ω be a space of colored geometric shapes that are divided into two categories (**1** and **0**).

Assume a random vector $X = (X_{color}, X_{shape}, X_{cat})$ where

- ▶ $X_{color} : \Omega \rightarrow \{red, blue\}$,
- ▶ $X_{shape} : \Omega \rightarrow \{circle, square\}$,
- ▶ $X_{cat} : \Omega \rightarrow \{\mathbf{1}, \mathbf{0}\}$.

The following tables give probability distribution of values:

category **1**:

	circle	square
red	0.2	0.02
blue	0.02	0.01

category **0**:

	circle	square
red	0.05	0.3
blue	0.2	0.2

Random Vectors – Example

Example:

$$P(\text{red}, \text{circle}, \mathbf{1}) = P(X_{\text{color}} = \text{red}, X_{\text{shape}} = \text{circle}, X_{\text{cat}} = \mathbf{1}) = 0.2$$

Random Vectors – Example

Example:

$$P(\text{red}, \text{circle}, \mathbf{1}) = P(X_{\text{color}} = \text{red}, X_{\text{shape}} = \text{circle}, X_{\text{cat}} = \mathbf{1}) = 0.2$$

"Summing over" all possible values of some variable(s) gives the distribution of the rest:

$$\begin{aligned} P(\text{red}, \text{circle}) &= P(X_{\text{color}} = \text{red}, X_{\text{shape}} = \text{circle}) \\ &= P(\text{red}, \text{circle}, \mathbf{1}) + P(\text{red}, \text{circle}, \mathbf{0}) \\ &= 0.2 + 0.05 = 0.25 \end{aligned}$$

Random Vectors – Example

Example:

$$P(\text{red}, \text{circle}, \mathbf{1}) = P(X_{\text{color}} = \text{red}, X_{\text{shape}} = \text{circle}, X_{\text{cat}} = \mathbf{1}) = 0.2$$

"Summing over" all possible values of some variable(s) gives the distribution of the rest:

$$\begin{aligned} P(\text{red}, \text{circle}) &= P(X_{\text{color}} = \text{red}, X_{\text{shape}} = \text{circle}) \\ &= P(\text{red}, \text{circle}, \mathbf{1}) + P(\text{red}, \text{circle}, \mathbf{0}) \\ &= 0.2 + 0.05 = 0.25 \end{aligned}$$

$$P(\text{red}) = 0.2 + 0.02 + 0.05 + 0.3 = 0.57$$

Random Vectors – Example

Example:

$$P(\text{red}, \text{circle}, \mathbf{1}) = P(X_{\text{color}} = \text{red}, X_{\text{shape}} = \text{circle}, X_{\text{cat}} = \mathbf{1}) = 0.2$$

"Summing over" all possible values of some variable(s) gives the distribution of the rest:

$$\begin{aligned} P(\text{red}, \text{circle}) &= P(X_{\text{color}} = \text{red}, X_{\text{shape}} = \text{circle}) \\ &= P(\text{red}, \text{circle}, \mathbf{1}) + P(\text{red}, \text{circle}, \mathbf{0}) \\ &= 0.2 + 0.05 = 0.25 \end{aligned}$$

$$P(\text{red}) = 0.2 + 0.02 + 0.05 + 0.3 = 0.57$$

Thus also, all conditional probabilities can be computed:

$$P(\mathbf{1} \mid \text{red}, \text{circle}) = \frac{P(\text{red}, \text{circle}, \mathbf{1})}{P(\text{red}, \text{circle})} = \frac{0.2}{0.25} = 0.8$$

Bayesian Classification

Let Ω be a sample space (a universum) of all objects that can be classified. We assume a probability P on Ω .

Bayesian Classification

Let Ω be a sample space (a universum) of all objects that can be classified. We assume a probability P on Ω .

We consider the problem of **binary classification**:

- ▶ Let Y be the random variable for the category which takes values in $\{0, 1\}$.

Bayesian Classification

Let Ω be a sample space (a universum) of all objects that can be classified. We assume a probability P on Ω .

We consider the problem of **binary classification**:

- ▶ Let Y be the random variable for the category which takes values in $\{\mathbf{0}, \mathbf{1}\}$.
- ▶ Let X be the random vector describing n features of a given instance, i.e., $X = (X_1, \dots, X_n)$
 - ▶ Denote by $\vec{x} \in \mathbb{R}^n$ values of X ,
 - ▶ and by $x_i \in \mathbb{R}$ values of X_i .

Bayesian Classification

Let Ω be a sample space (a universum) of all objects that can be classified. We assume a probability P on Ω .

We consider the problem of **binary classification**:

- ▶ Let Y be the random variable for the category which takes values in $\{\mathbf{0}, \mathbf{1}\}$.
- ▶ Let X be the random vector describing n features of a given instance, i.e., $X = (X_1, \dots, X_n)$
 - ▶ Denote by $\vec{x} \in \mathbb{R}^n$ values of X ,
 - ▶ and by $x_i \in \mathbb{R}$ values of X_i .

Bayes classifier: Given a vector of feature values \vec{x} ,

$$C^{Bayes}(\vec{x}) := \begin{cases} \mathbf{1} & \text{if } P(Y = \mathbf{1} \mid X = \vec{x}) \geq P(Y = \mathbf{0} \mid X = \vec{x}) \\ \mathbf{0} & \text{otherwise.} \end{cases}$$

Intuitively, C^{Bayes} assigns to \vec{x} the most probable category it might be in.

Bayesian Classification – Example

Imagine a conveyor belt with apples and apricots.

A machine is supposed to correctly distinguish apples from apricots based on their weight and diameter.

Bayesian Classification – Example

Imagine a conveyor belt with apples and apricots.

A machine is supposed to correctly distinguish apples from apricots based on their weight and diameter.

That is,

- ▶ $Y \in \{1, 0\}$
(here our interpretation is **1** = apple, **0** = apricot)
- ▶ $X = (X_{weight}, X_{diam})$

Bayesian Classification – Example

Imagine a conveyor belt with apples and apricots.

A machine is supposed to correctly distinguish apples from apricots based on their weight and diameter.

That is,

- ▶ $Y \in \{1, 0\}$
(here our interpretation is **1** = apple, **0** = apricot)
- ▶ $X = (X_{weight}, X_{diam})$

We are given a fruit of a diameter of 5cm that weighs 40g.

Bayesian Classification – Example

Imagine a conveyor belt with apples and apricots.

A machine is supposed to correctly distinguish apples from apricots based on their weight and diameter.

That is,

- ▶ $Y \in \{1, 0\}$
(here our interpretation is **1** = apple, **0** = apricot)
- ▶ $X = (X_{weight}, X_{diam})$

We are given a fruit of a diameter of 5cm that weighs 40g.

The Bayes classifier compares $P(Y = 1 \mid X = (40g, 5cm))$ with $P(Y = 0 \mid X = (40g, 5cm))$ and selects the more probable category given the features.

Crucial question: Is such a classifier good?

There are other classifiers, e.g., one which compares the weight divided by 10 with the diameter and decides based on the answer, or maybe a classifier that sums the weight and the diameter and compares the result with a constant, etc.

Bayes Classifier

Let C be an arbitrary *classifier*, that is a function that to every feature vector $\vec{x} \in \mathbb{R}^n$ assigns a class from $\{\mathbf{0}, \mathbf{1}\}$.

Bayes Classifier

Let C be an arbitrary *classifier*, that is a function that to every feature vector $\vec{x} \in \mathbb{R}^n$ assigns a class from $\{\mathbf{0}, \mathbf{1}\}$.

Define the error of the classifier C by

$$E_C = P(Y \neq C)$$

(Here we slightly abuse notation and apply C to samples, technically we apply the composition $C \circ X$ of C and X which first determines the features using X and then classifies according to C).

Bayes Classifier

Let C be an arbitrary *classifier*, that is a function that to every feature vector $\vec{x} \in \mathbb{R}^n$ assigns a class from $\{\mathbf{0}, \mathbf{1}\}$.

Define the error of the classifier C by

$$E_C = P(Y \neq C)$$

(Here we slightly abuse notation and apply C to samples, technically we apply the composition $C \circ X$ of C and X which first determines the features using X and then classifies according to C).

Theorem

The Bayes classifier C^{Bayes} minimizes E_C , that is

$$E_{C^{Bayes}} = \min_{C \text{ is a classifier}} E_C$$

Practical Use of Bayes Classifier

The crucial problem: The probability P is not known!

In particular, where to get $P(Y = \mathbf{1} \mid X = \vec{x})$?

Note that $P(Y = \mathbf{0} \mid X = \vec{x}) = 1 - P(Y = \mathbf{1} \mid X = \vec{x})$

Practical Use of Bayes Classifier

The crucial problem: The probability P is not known!

In particular, where to get $P(Y = \mathbf{1} \mid X = \vec{x})$?

Note that $P(Y = \mathbf{0} \mid X = \vec{x}) = 1 - P(Y = \mathbf{1} \mid X = \vec{x})$

Given no other assumptions, this requires a table showing the probability of the category $\mathbf{1}$ for each possible feature vector \vec{x} .

Where do you get these probabilities?

Practical Use of Bayes Classifier

The crucial problem: The probability P is not known!

In particular, where to get $P(Y = \mathbf{1} \mid X = \vec{x})$?

Note that $P(Y = \mathbf{0} \mid X = \vec{x}) = 1 - P(Y = \mathbf{1} \mid X = \vec{x})$

Given no other assumptions, this requires a table showing the probability of the category $\mathbf{1}$ for each possible feature vector \vec{x} .

Where do you get these probabilities?

In some cases, the probabilities might come from the knowledge of the solved problem (e.g., applications in physics might be supported by a theory giving the probabilities).

In most cases, however, P is estimated from sampled data by

$$\bar{P}(Y = \mathbf{1} \mid X = \vec{x}) = \frac{\text{number of samples with } Y = \mathbf{1} \text{ and } X = \vec{x}}{\text{number of samples with } X = \vec{x}}$$

(We use \bar{P} to denote an estimate of P from data.)

Estimating P

Consider a problem with $X = (X_1, X_2, X_3)$ where each X_i returns either 0 or 1. What might the data look like?

Estimating P

Consider a problem with $X = (X_1, X_2, X_3)$ where each X_i returns either 0 or 1. What might the data look like?

Part of the data table:

Y	X_1	X_2	X_3
1	1	0	1
1	0	1	1
0	1	0	1
0	0	0	1
1	0	0	0
0	1	1	1
...			

All data with $X_1 = 1, X_2 = 0, X_3 = 1$:

Y	X_1	X_2	X_3
1	1	0	1
1	1	0	1
0	1	0	1
0	1	0	1
1	1	0	1
1	1	0	1

Estimate: $\bar{P}(\mathbf{1} \mid 1, 0, 1) = 2/3$

Estimating P

Consider a problem with $X = (X_1, X_2, X_3)$ where each X_i returns either 0 or 1. What might the data look like?

Part of the data table:

Y	X_1	X_2	X_3
1	1	0	1
1	0	1	1
0	1	0	1
0	0	0	1
1	0	0	0
0	1	1	1
...			

All data with $X_1 = 1, X_2 = 0, X_3 = 1$:

Y	X_1	X_2	X_3
1	1	0	1
1	1	0	1
0	1	0	1
0	1	0	1
1	1	0	1
1	1	0	1

Estimate: $\bar{P}(\mathbf{1} \mid 1, 0, 1) = 2/3$

The probability table and the necessary data are typically too large!

Concretely, if all X_1, \dots, X_n are binary, there are 2^n probabilities $P(Y = \mathbf{1} \mid X = \vec{x})$, one for each possible $\vec{x} \in \{0, 1\}^n$.

Let's Look at It the Other Way Round

Theorem (Bayes,1764)

$$P(A | B) = \frac{P(B | A) \cdot P(A)}{P(B)}$$

Let's Look at It the Other Way Round

Theorem (Bayes, 1764)

$$P(A | B) = \frac{P(B | A) \cdot P(A)}{P(B)}$$

Proof.

$$P(A | B) = \frac{P(A \cap B)}{P(B)} = \frac{\frac{P(A \cap B)}{P(A)} \cdot P(A)}{P(B)} = \frac{P(B | A) \cdot P(A)}{P(B)}$$



Bayesian Classification

Determine the category for \vec{x} by computing

$$P(Y = y \mid X = \vec{x}) = \frac{P(Y = y) \cdot P(X = \vec{x} \mid Y = y)}{P(X = \vec{x})}$$

for both $y \in \{\mathbf{0}, \mathbf{1}\}$ and deciding whether or not the following holds:

$$P(Y = \mathbf{1} \mid X = \vec{x}) \geq P(Y = \mathbf{0} \mid X = \vec{x})$$

Bayesian Classification

Determine the category for \vec{x} by computing

$$P(Y = y \mid X = \vec{x}) = \frac{P(Y = y) \cdot P(X = \vec{x} \mid Y = y)}{P(X = \vec{x})}$$

for both $y \in \{\mathbf{0}, \mathbf{1}\}$ and deciding whether or not the following holds:

$$P(Y = \mathbf{1} \mid X = \vec{x}) \geq P(Y = \mathbf{0} \mid X = \vec{x})$$

So, to make the classifier, we need to compute the following:

- ▶ **The prior** $P(Y = \mathbf{1})$ (then $P(Y = \mathbf{0}) = 1 - P(Y = \mathbf{1})$)
- ▶ **The conditionals** $P(X = \vec{x} \mid Y = y)$ for $y \in \{\mathbf{0}, \mathbf{1}\}$ and for every \vec{x}

Estimating the Prior and Conditionals

- ▶ $P(Y = \mathbf{1})$ can be easily estimated from data by

$$\bar{P}(Y = \mathbf{1}) = \frac{\text{number of samples with } Y = \mathbf{1}}{\text{number of all samples}}$$

Estimating the Prior and Conditionals

- ▶ $P(Y = \mathbf{1})$ can be easily estimated from data by

$$\bar{P}(Y = \mathbf{1}) = \frac{\text{number of samples with } Y = \mathbf{1}}{\text{number of all samples}}$$

- ▶ If the dimension of features is small, $P(X = \vec{x} \mid Y = y)$ can be estimated from data similarly as $P(Y = \mathbf{1} \mid X = \vec{x})$ by

$$\bar{P}(X = \vec{x} \mid Y = y) = \frac{\text{number of samples with } Y = y \text{ and } X = \vec{x}}{\text{number of samples with } Y = y}$$

Unfortunately, for higher dimensional data too many samples are needed to estimate all $P(X = \vec{x} \mid Y = y)$ (there are too many \vec{x} 's).

So where is the advantage of using the Bayes thm.??

Estimating the Prior and Conditionals

- ▶ $P(Y = \mathbf{1})$ can be easily estimated from data by

$$\bar{P}(Y = \mathbf{1}) = \frac{\text{number of samples with } Y = \mathbf{1}}{\text{number of all samples}}$$

- ▶ If the dimension of features is small, $P(X = \vec{x} \mid Y = y)$ can be estimated from data similarly as $P(Y = \mathbf{1} \mid X = \vec{x})$ by

$$\bar{P}(X = \vec{x} \mid Y = y) = \frac{\text{number of samples with } Y = y \text{ and } X = \vec{x}}{\text{number of samples with } Y = y}$$

Unfortunately, for higher dimensional data too many samples are needed to estimate all $P(X = \vec{x} \mid Y = y)$ (there are too many \vec{x} 's).

So where is the advantage of using the Bayes thm.??

We introduce *independence assumptions* about the features!

Naive Bayes

- ▶ We assume that features are (conditionally) independent *given the category*. That is for all $\vec{x} = (x_1, \dots, x_n)$ and $y \in \{0, 1\}$ we **assume**:

$$\begin{aligned} P(X = x \mid Y = y) &= P(X_1 = x_1, \dots, X_n = x_n \mid Y) \\ &= \prod_{i=1}^n P(X_i = x_i \mid Y = y) \end{aligned}$$

Naive Bayes

- ▶ We assume that features are (conditionally) independent *given the category*. That is for all $\vec{x} = (x_1, \dots, x_n)$ and $y \in \{\mathbf{0}, \mathbf{1}\}$ we **assume**:

$$\begin{aligned} P(X = x \mid Y = y) &= P(X_1 = x_1, \dots, X_n = x_n \mid Y) \\ &= \prod_{i=1}^n P(X_i = x_i \mid Y = y) \end{aligned}$$

- ▶ Therefore, we only need to specify $P(X_i = x_i \mid Y = y)$ for each possible pair of a feature-value x_i and $y \in \{\mathbf{0}, \mathbf{1}\}$.

Naive Bayes

- ▶ We assume that features are (conditionally) independent *given the category*. That is for all $\vec{x} = (x_1, \dots, x_n)$ and $y \in \{0, 1\}$ we **assume**:

$$\begin{aligned} P(X = x \mid Y = y) &= P(X_1 = x_1, \dots, X_n = x_n \mid Y) \\ &= \prod_{i=1}^n P(X_i = x_i \mid Y = y) \end{aligned}$$

- ▶ Therefore, we only need to specify $P(X_i = x_i \mid Y = y)$ for each possible pair of a feature-value x_i and $y \in \{0, 1\}$.

Note that if all X_i are binary (values in $\{0, 1\}$), this requires specifying only $2n$ parameters:

$$P(X_i = 1 \mid Y = \mathbf{1}) \text{ and } P(X_i = 1 \mid Y = \mathbf{0}) \text{ for each } X_i$$

as $P(X_i = 0 \mid Y = y) = 1 - P(X_i = 1 \mid Y = y)$ for $y \in \{0, 1\}$.

Compared to specifying 2^n parameters without any independence assumption.

Estimating the marginal probabilities

Estimate the probabilities $P(X_i = x_i \mid Y = y)$ by

$$\bar{P}(X_i = x_i \mid Y = y) = \frac{\text{number of samples with } X_i = x_i \text{ and } Y = y}{\text{number of samples with } Y = y}$$

Estimating the marginal probabilities

Estimate the probabilities $P(X_i = x_i \mid Y = y)$ by

$$\bar{P}(X_i = x_i \mid Y = y) = \frac{\text{number of samples with } X_i = x_i \text{ and } Y = y}{\text{number of samples with } Y = y}$$

Example: Consider a problem with $X = (X_1, X_2, X_3)$ where each X_i returns either 0 or 1. The data is

Y	X_1	X_2	X_3
1	1	0	1
1	0	1	1
0	1	0	1
0	0	0	1
1	0	0	0
0	1	1	1

$$\bar{P}(X_1 = 1 \mid Y = \mathbf{1}) = 1/3 \quad \bar{P}(X_1 = 1 \mid Y = \mathbf{0}) = 2/3$$

$$\bar{P}(X_2 = 1 \mid Y = \mathbf{1}) = 1/3 \quad \bar{P}(X_2 = 1 \mid Y = \mathbf{0}) = 1/3$$

$$\bar{P}(X_3 = 1 \mid Y = \mathbf{1}) = 2/3 \quad \bar{P}(X_3 = 1 \mid Y = \mathbf{0}) = 1$$

Naive Bayes – Example

Consider classification of geometric shapes:

$X_1 \in \{small, medium, large\}$

$X_2 \in \{red, blue, green\}$

$X_3 \in \{square, triangle, circle\}$

Naive Bayes – Example

Consider classification of geometric shapes:

$X_1 \in \{small, medium, large\}$

$X_2 \in \{red, blue, green\}$

$X_3 \in \{square, triangle, circle\}$

Assume that we have already estimated the following probabilities:

	$Y = \mathbf{1}$	$Y = \mathbf{0}$
$\bar{P}(Y)$	0.5	0.5
$\bar{P}(small \mid Y)$	0.4	0.4
$\bar{P}(medium \mid Y)$	0.1	0.2
$\bar{P}(large \mid Y)$	0.5	0.4
$\bar{P}(red \mid Y)$	0.9	0.3
$\bar{P}(blue \mid Y)$	0.05	0.3
$\bar{P}(green \mid Y)$	0.05	0.4
$\bar{P}(square \mid Y)$	0.05	0.4
$\bar{P}(triangle \mid Y)$	0.05	0.3
$\bar{P}(circle \mid Y)$	0.9	0.3

Does $(medium, red, circle)$ belong to the category **1** ?

	$Y = \mathbf{1}$	$Y = \mathbf{0}$
$\bar{P}(Y)$	0.5	0.5
$\bar{P}(\textit{medium} \mid Y)$	0.1	0.2
$\bar{P}(\textit{red} \mid Y)$	0.9	0.3
$\bar{P}(\textit{circle} \mid Y)$	0.9	0.3

Denote $\vec{x} = (\textit{medium}, \textit{red}, \textit{circle})$.

	$Y = \mathbf{1}$	$Y = \mathbf{0}$
$\bar{P}(Y)$	0.5	0.5
$\bar{P}(\textit{medium} \mid Y)$	0.1	0.2
$\bar{P}(\textit{red} \mid Y)$	0.9	0.3
$\bar{P}(\textit{circle} \mid Y)$	0.9	0.3

Denote $\vec{x} = (\textit{medium}, \textit{red}, \textit{circle})$.

$$\begin{aligned}
 P(Y = \mathbf{1} \mid X = \vec{x}) &= \\
 &= P(\mathbf{1}) \cdot P(\textit{medium} \mid \mathbf{1}) \cdot P(\textit{red} \mid \mathbf{1}) \cdot P(\textit{circle} \mid \mathbf{1}) / P(X = \vec{x}) \\
 &\doteq 0.5 \cdot 0.1 \cdot 0.9 \cdot 0.9 / P(X = \vec{x}) = 0.0405 / P(X = \vec{x})
 \end{aligned}$$

	$Y = \mathbf{1}$	$Y = \mathbf{0}$
$\bar{P}(Y)$	0.5	0.5
$\bar{P}(\text{medium} \mid Y)$	0.1	0.2
$\bar{P}(\text{red} \mid Y)$	0.9	0.3
$\bar{P}(\text{circle} \mid Y)$	0.9	0.3

Denote $\vec{x} = (\text{medium}, \text{red}, \text{circle})$.

$$\begin{aligned}
 P(Y = \mathbf{1} \mid X = \vec{x}) &= \\
 &= P(\mathbf{1}) \cdot P(\text{medium} \mid \mathbf{1}) \cdot P(\text{red} \mid \mathbf{1}) \cdot P(\text{circle} \mid \mathbf{1}) / P(X = \vec{x}) \\
 &\doteq 0.5 \cdot 0.1 \cdot 0.9 \cdot 0.9 / P(X = \vec{x}) = 0.0405 / P(X = \vec{x})
 \end{aligned}$$

$$\begin{aligned}
 P(Y = \mathbf{0} \mid X = \vec{x}) &= \\
 &= P(\mathbf{0}) \cdot P(\text{medium} \mid \mathbf{0}) \cdot P(\text{red} \mid \mathbf{0}) \cdot P(\text{circle} \mid \mathbf{0}) / P(X = \vec{x}) \\
 &\doteq 0.5 \cdot 0.2 \cdot 0.3 \cdot 0.3 / P(X = \vec{x}) = 0.009 / P(X = \vec{x})
 \end{aligned}$$

(Note that we used the estimates \bar{P} of P to finish the computation above.)

	$Y = \mathbf{1}$	$Y = \mathbf{0}$
$\bar{P}(Y)$	0.5	0.5
$\bar{P}(\text{medium} \mid Y)$	0.1	0.2
$\bar{P}(\text{red} \mid Y)$	0.9	0.3
$\bar{P}(\text{circle} \mid Y)$	0.9	0.3

Denote $\vec{x} = (\text{medium}, \text{red}, \text{circle})$.

$$\begin{aligned}
 P(Y = \mathbf{1} \mid X = \vec{x}) &= \\
 &= P(\mathbf{1}) \cdot P(\text{medium} \mid \mathbf{1}) \cdot P(\text{red} \mid \mathbf{1}) \cdot P(\text{circle} \mid \mathbf{1}) / P(X = \vec{x}) \\
 &\doteq 0.5 \cdot 0.1 \cdot 0.9 \cdot 0.9 / P(X = \vec{x}) = 0.0405 / P(X = \vec{x})
 \end{aligned}$$

$$\begin{aligned}
 P(Y = \mathbf{0} \mid X = \vec{x}) &= \\
 &= P(\mathbf{0}) \cdot P(\text{medium} \mid \mathbf{0}) \cdot P(\text{red} \mid \mathbf{0}) \cdot P(\text{circle} \mid \mathbf{0}) / P(X = \vec{x}) \\
 &\doteq 0.5 \cdot 0.2 \cdot 0.3 \cdot 0.3 / P(X = \vec{x}) = 0.009 / P(X = \vec{x})
 \end{aligned}$$

(Note that we used the estimates \bar{P} of P to finish the computation above.)

Apparently,

$$P(Y = \mathbf{1} \mid X = \vec{x}) \doteq 0.0405 / P(X = \vec{x}) > 0.009 / P(X = \vec{x}) \doteq P(\mathbf{0} \mid X = \vec{x})$$

So we classify \vec{x} to the category $\mathbf{1}$.

Estimating Probabilities in Practice

We already know that $P(X_i = x_i \mid Y = y)$ can be estimated by

$$\bar{P}(X_i = x_i \mid Y = y) = \ell_{y,x_i} / \ell_y$$

where

- ▶ ℓ_{y,x_i} = number of samples with $Y = y$ and $X_i = x_i$
- ▶ ℓ_y = number of samples with $Y = y$

Estimating Probabilities in Practice

We already know that $P(X_i = x_i \mid Y = y)$ can be estimated by

$$\bar{P}(X_i = x_i \mid Y = y) = \ell_{y,x_i} / \ell_y$$

where

- ▶ ℓ_{y,x_i} = number of samples with $Y = y$ and $X_i = x_i$
- ▶ ℓ_y = number of samples with $Y = y$

Problem: If, by chance, a rare value x_i of a feature X_i never occurs in the training data, we get

$$\bar{P}(X_i = x_i \mid Y = y) = 0 \quad \text{for both } y \in \{\mathbf{0}, \mathbf{1}\}$$

But then $\bar{P}(X = x) = 0$ for x containing the value x_i for X_i , and thus $\bar{P}(Y = y \mid X = x)$ is not well defined.

Moreover, $\bar{P}(Y = y) \cdot \bar{P}(X = x \mid Y = y) = 0$ (for $y \in \{\mathbf{0}, \mathbf{1}\}$) so even this cannot be used for classification.

Probability Estimation Example

Training data:

Size	Color	Shape	Class
small	red	circle	1
large	red	circle	1
small	red	triangle	0
large	blue	circle	0

Estimated probabilities:

	$Y = 1$	$Y = 0$
$\bar{P}(Y)$	0.5	0.5
$\bar{P}(\text{small} \mid Y)$	0.5	0.5
$\bar{P}(\text{medium} \mid Y)$	0	0
$\bar{P}(\text{large} \mid Y)$	0.5	0.5
$\bar{P}(\text{red} \mid Y)$	1	0.5
$\bar{P}(\text{blue} \mid Y)$	0	0.5
$\bar{P}(\text{green} \mid Y)$	0	0
$\bar{P}(\text{square} \mid Y)$	0	0
$\bar{P}(\text{triangle} \mid Y)$	0	0.5
$\bar{P}(\text{circle} \mid Y)$	1	0.5

Note that $\bar{P}(\text{medium} \mid \mathbf{1}) = P(\text{medium} \mid \mathbf{0}) = 0$ and thus also $\bar{P}(\text{medium}, \text{red}, \text{circle}) = 0$.

So what is $\bar{P}(\mathbf{1} \mid \text{medium}, \text{red}, \text{circle})$?

Smoothing

- ▶ To account for estimation from small samples, probability estimates are adjusted or *smoothed*.

Smoothing

- ▶ To account for estimation from small samples, probability estimates are adjusted or *smoothed*.
- ▶ *Laplace smoothing* adds one to every count of feature values

$$\tilde{P}(X_i = x_i \mid Y = y) = \frac{\ell_{y,x_i} + 1}{\ell_y + v_i}$$

where

- ▶ ℓ_y = number of training samples with $Y = y$,
- ▶ ℓ_{y,x_i} = number of training samples with $Y = y$ and $X_i = x_i$,
- ▶ v_i is the number of all distinct values of the variable X_i .

To understand note that

$$\ell_y = \sum_{x_i \text{ is a value of } X_i} \ell_{y,x_i}$$

and thus

$$\bar{P}(X_i = x_i \mid Y = y) = \ell_{y,x_i} / \sum_{x_i \text{ is a value of } X_i} \ell_{y,x_i}$$

$$\tilde{P}(X_i = x_i \mid Y = y) = (\ell_{y,x_i} + 1) / \sum_{x_i \text{ is a value of } X_i} (\ell_{y,x_i} + 1)$$

Laplace Smoothing Example

- ▶ Assume training set contains 10 samples of category **1**:
 - ▶ 4 small
 - ▶ 0 medium
 - ▶ 6 large

Laplace Smoothing Example

- ▶ Assume training set contains 10 samples of category **1**:
 - ▶ 4 small
 - ▶ 0 medium
 - ▶ 6 large
- ▶ Estimate parameters as follows
 - ▶ $\tilde{P}(\text{small} \mid \mathbf{1}) = (4 + 1)/(10 + 3) = 0.384$
 - ▶ $\tilde{P}(\text{medium} \mid \mathbf{1}) = (0 + 1)/(10 + 3) = 0.0769$
 - ▶ $\tilde{P}(\text{large} \mid \mathbf{1}) = (6 + 1)/(10 + 3) = 0.538$

Continuous Features

Ω may be (potentially) continuous, X_i may assign a continuum of values in \mathbb{R} .

Continuous Features

Ω may be (potentially) continuous, X_i may assign a continuum of values in \mathbb{R} .

- The probabilities are computed using *probability density*

$p : \mathbb{R} \rightarrow \mathbb{R}^+$.

A random variable $X : \Omega \rightarrow \mathbb{R}^+$ has a density $p : \mathbb{R} \rightarrow \mathbb{R}^+$ if for every interval $[a, b]$ we have

$$P(a \leq X \leq b) = \int_a^b p(x) dx$$

Usually, $P(X_i \mid Y = y)$ is used to denote the *density* of X_i conditioned on $Y = y$.

Continuous Features

Ω may be (potentially) continuous, X_i may assign a continuum of values in \mathbb{R} .

- ▶ The probabilities are computed using *probability density*

$p : \mathbb{R} \rightarrow \mathbb{R}^+$.

A random variable $X : \Omega \rightarrow \mathbb{R}^+$ has a density $p : \mathbb{R} \rightarrow \mathbb{R}^+$ if for every interval $[a, b]$ we have

$$P(a \leq X \leq b) = \int_a^b p(x) dx$$

Usually, $P(X_i | Y = y)$ is used to denote the *density* of X_i conditioned on $Y = y$.

- ▶ The densities $P(X_i | Y = y)$ are usually estimated using Gaussian densities as follows:
 - ▶ Estimate the mean μ_{iy} and the standard deviation σ_{iy} based on training data.
 - ▶ Then put

$$\bar{P}(X_i | Y = y) = \frac{1}{\sigma_{iy} \sqrt{2\pi}} \exp \left(\frac{-(X_i - \mu_{iy})^2}{2\sigma_{iy}^2} \right)$$

Comments on Naive Bayes

- ▶ Tends to work well despite rather a strong assumption of conditional independence of features.

Comments on Naive Bayes

- ▶ Tends to work well despite rather a strong assumption of conditional independence of features.
- ▶ Experiments show that it is quite competitive with other classification methods.

Even if the probabilities are not accurately estimated, it often picks the correct maximum probability category.

Comments on Naive Bayes

- ▶ Tends to work well despite rather a strong assumption of conditional independence of features.
- ▶ Experiments show that it is quite competitive with other classification methods.
Even if the probabilities are not accurately estimated, it often picks the correct maximum probability category.
- ▶ Directly constructs a model from parameter estimates that are calculated from the training data.

Comments on Naive Bayes

- ▶ Tends to work well despite rather a strong assumption of conditional independence of features.
- ▶ Experiments show that it is quite competitive with other classification methods.
Even if the probabilities are not accurately estimated, it often picks the correct maximum probability category.
- ▶ Directly constructs a model from parameter estimates that are calculated from the training data.
- ▶ Typically handles outliers and noise well in the discrete version. In the continuous case depends on the way the parameters are estimated (e.g., the mean is very sensitive to outliers).

Comments on Naive Bayes

- ▶ Tends to work well despite rather a strong assumption of conditional independence of features.
- ▶ Experiments show that it is quite competitive with other classification methods.
Even if the probabilities are not accurately estimated, it often picks the correct maximum probability category.
- ▶ Directly constructs a model from parameter estimates that are calculated from the training data.
- ▶ Typically handles outliers and noise well in the discrete version. In the continuous case depends on the way the parameters are estimated (e.g., the mean is very sensitive to outliers).
- ▶ Missing values are easy to deal with; use only non-missing values in the computation of $\bar{P}(X_i = x_i \mid Y = y)$.

Bayesian Networks (Basic Information)

In the Naive Bayes, we have assumed that *all* features X_1, \dots, X_n are independent.

Bayesian Networks (Basic Information)

In the Naive Bayes, we have assumed that *all* features X_1, \dots, X_n are independent.

This is usually not realistic.

E.g. Variables "rain" and "grass wet" are (usually) strongly dependent.

Bayesian Networks (Basic Information)

In the Naive Bayes, we have assumed that *all* features X_1, \dots, X_n are independent.

This is usually not realistic.

E.g. Variables "rain" and "grass wet" are (usually) strongly dependent.

What if we return some dependencies?

(But now in a well-defined sense.)

Bayesian Networks (Basic Information)

In the Naive Bayes, we have assumed that *all* features X_1, \dots, X_n are independent.

This is usually not realistic.

E.g. Variables "rain" and "grass wet" are (usually) strongly dependent.

What if we return some dependencies?

(But now in a well-defined sense.)

Bayesian networks are a graphical model that uses a directed acyclic graph to specify dependencies among variables.

Bayesian Networks – Example

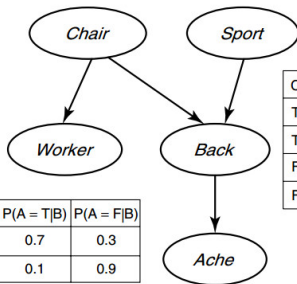
$P(C = T)$	$P(C = F)$
0.8	0.2

$P(S = T)$	$P(S = F)$
0.02	0.98

C	$P(W = T C)$	$P(W = F C)$
T	0.9	0.1
F	0.01	0.99

B	$P(A = T B)$	$P(A = F B)$
T	0.7	0.3
F	0.1	0.9

C	S	$P(B = T C,S)$	$P(B = F C,S)$
T	T	0.9	0.1
T	F	0.2	0.8
F	T	0.9	0.1
F	F	0.01	0.99



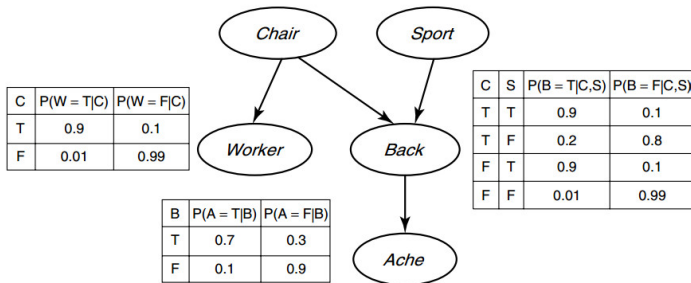
Now, e.g.,

$$P(C, S, W, B, A) = P(C) \cdot P(S) \cdot P(W | C) \cdot P(B | C, S) \cdot P(A | B)$$

Bayesian Networks – Example

$P(C = T)$	$P(C = F)$
0.8	0.2

$P(S = T)$	$P(S = F)$
0.02	0.98



Now, e.g.,

$$P(C, S, W, B, A) = P(C) \cdot P(S) \cdot P(W | C) \cdot P(B | C, S) \cdot P(A | B)$$

Now, we may, e.g., infer the probability $P(C = T | A = T)$ that we sit in the wrong chair, assuming that our back aches.

Bayesian Networks – Example

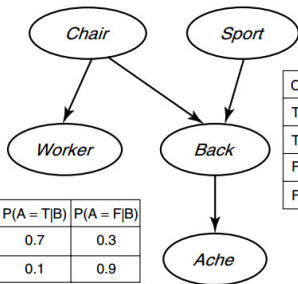
$P(C = T)$	$P(C = F)$
0.8	0.2

$P(S = T)$	$P(S = F)$
0.02	0.98

C	$P(W = T C)$	$P(W = F C)$
T	0.9	0.1
F	0.01	0.99

B	$P(A = T B)$	$P(A = F B)$
T	0.7	0.3
F	0.1	0.9

C	S	$P(B = T C,S)$	$P(B = F C,S)$
T	T	0.9	0.1
T	F	0.2	0.8
F	T	0.9	0.1
F	F	0.01	0.99



Now, e.g.,

$$P(C, S, W, B, A) = P(C) \cdot P(S) \cdot P(W | C) \cdot P(B | C, S) \cdot P(A | B)$$

Now, we may, e.g., infer the probability $P(C = T | A = T)$ that we sit in the wrong chair, assuming that our back aches.

We have to store only 10 numbers as opposed to $2^5 - 1$ possible probabilities for all vectors of values of C, S, W, B, A .

Bayesian Networks – Learning & Naive Bayes

Many algorithms have been developed for learning:

- ▶ the structure of the graph of the network,
- ▶ the *conditional probability tables*.

The methods are based on maximum-likelihood estimation, gradient descent, etc.

Automatic procedures are usually combined with expert knowledge.

Bayesian Networks – Learning & Naive Bayes

Many algorithms have been developed for learning:

- ▶ the structure of the graph of the network,
- ▶ the *conditional probability tables*.

The methods are based on maximum-likelihood estimation, gradient descent, etc.

Automatic procedures are usually combined with expert knowledge.

Can you express the naive Bayes for Y, X_1, \dots, X_n using a Bayesian network?

Classifier Evaluation

Classifier

Assume binary classification into two classes $\{0, 1\}$.

Classifier

Assume binary classification into two classes $\{0, 1\}$.

Consider a classification dataset:

$$\{(\vec{x}_k, c_k) \mid k = 1, \dots, p\}$$

Here \vec{x}_k is a vector of attributes/features and $c_k \in \{0, 1\}$ for all k .

Classifier

Assume binary classification into two classes $\{0, 1\}$.

Consider a classification dataset:

$$\{(\vec{x}_k, c_k) \mid k = 1, \dots, p\}$$

Here \vec{x}_k is a vector of attributes/features and $c_k \in \{0, 1\}$ for all k .

Consider a sequence of predictions generated by a classifier:

$$h_1, \dots, h_p \in \{0, 1\}$$

Here each h_k has been predicted for the k -th example (\vec{x}_k, c_k) .

Classifier

Assume binary classification into two classes $\{0, 1\}$.

Consider a classification dataset:

$$\{(\vec{x}_k, c_k) \mid k = 1, \dots, p\}$$

Here \vec{x}_k is a vector of attributes/features and $c_k \in \{0, 1\}$ for all k .

Consider a sequence of predictions generated by a classifier:

$$h_1, \dots, h_p \in \{0, 1\}$$

Here each h_k has been predicted for the k -th example (\vec{x}_k, c_k) .

How good are the predictions h_1, \dots, h_p w.r.t. c_1, \dots, c_p ?

There are many possible metrics ...

Classifier

Assume binary classification into two classes $\{0, 1\}$.

Consider a classification dataset:

$$\{(\vec{x}_k, c_k) \mid k = 1, \dots, p\}$$

Here \vec{x}_k is a vector of attributes/features and $c_k \in \{0, 1\}$ for all k .

Consider a sequence of predictions generated by a classifier:

$$h_1, \dots, h_p \in \{0, 1\}$$

Here each h_k has been predicted for the k -th example (\vec{x}_k, c_k) .

How good are the predictions h_1, \dots, h_p w.r.t. c_1, \dots, c_p ?

There are many possible metrics ...

I will call the class 1 *positive* and the class 0 *negative*.

Note that the class 0 is not negative in the numerical sense but in the absence of something (e.g., predicted illness).

Confusion Matrix for Binary Classifier

		Predicted	
		1	0
Actual	1	TP	FN
	0	FP	TN

Confusion Matrix for Binary Classifier

		Predicted	
		1	0
Actual	1	TP	FN
	0	FP	TN

- ▶ TP = number of correctly classified examples with actual class 1

$$TP = |\{k \mid h_k = 1 \wedge c_k = 1\}|$$

Confusion Matrix for Binary Classifier

		Predicted	
		1	0
Actual	1	TP	FN
	0	FP	TN

- ▶ TP = number of correctly classified examples with actual class 1

$$TP = |\{k \mid h_k = 1 \wedge c_k = 1\}|$$

- ▶ TN = number of correctly classified examples with actual class 0

$$TN = |\{k \mid h_k = 0 \wedge c_k = 0\}|$$

Confusion Matrix for Binary Classifier

		Predicted	
		1	0
Actual	1	TP	FN
	0	FP	TN

- ▶ TP = number of correctly classified examples with actual class 1

$$TP = |\{k \mid h_k = 1 \wedge c_k = 1\}|$$

- ▶ TN = number of correctly classified examples with actual class 0

$$TN = |\{k \mid h_k = 0 \wedge c_k = 0\}|$$

- ▶ FP = number of incorrectly classified examples with actual class 0

$$FP = |\{k \mid h_k = 1 \wedge c_k = 0\}|$$

Confusion Matrix for Binary Classifier

		Predicted	
		1	0
Actual	1	TP	FN
	0	FP	TN

- ▶ TP = number of correctly classified examples with actual class 1

$$TP = |\{k \mid h_k = 1 \wedge c_k = 1\}|$$

- ▶ TN = number of correctly classified examples with actual class 0

$$TN = |\{k \mid h_k = 0 \wedge c_k = 0\}|$$

- ▶ FP = number of incorrectly classified examples with actual class 0

$$FP = |\{k \mid h_k = 1 \wedge c_k = 0\}|$$

- ▶ FN = number of incorrectly classified examples with actual class 1

$$FN = |\{k \mid h_k = 0 \wedge c_k = 1\}|$$

Example

Given a sample of 12 individuals, eight have cancer, and four are cancer-free.

Example

Given a sample of 12 individuals, eight have cancer, and four are cancer-free.

Assume that we have trained a classifier with the following results:

Index	1	2	3	4	5	6	7	8	9	10	11	12
Actual	1	1	1	1	1	1	1	1	0	0	0	0
Predicted	0	0	1	1	1	1	1	1	1	0	0	0
Result	FN	FN	TP	TP	TP	TP	TP	TP	FP	TN	TN	TN

Example

Given a sample of 12 individuals, eight have cancer, and four are cancer-free.

Assume that we have trained a classifier with the following results:

Index	1	2	3	4	5	6	7	8	9	10	11	12
Actual	1	1	1	1	1	1	1	1	0	0	0	0
Predicted	0	0	1	1	1	1	1	1	1	0	0	0
Result	FN	FN	TP	TP	TP	TP	TP	TP	FP	TN	TN	TN

Actual condition	Predicted condition	
	Cancer	Non-cancer
Cancer	TP = 6	FN = 2
Non-cancer	FP = 1	TN = 3
Total	8 + 4 = 12	

Terminology

- ▶ TP aka hit
- ▶ TN aka correct rejection
- ▶ FP aka type I error, false alarm, overestimation
- ▶ FN aka type II error, miss, underestimation

Usually, TP, TN, FP, and FN are used to denote the individual examples of a particular kind and the number of these examples.

Terminology

- ▶ TP aka hit
- ▶ TN aka correct rejection
- ▶ FP aka type I error, false alarm, overestimation
- ▶ FN aka type II error, miss, underestimation

Usually, TP, TN, FP, and FN are used to denote the individual examples of a particular kind and the number of these examples.

In what follows, we also use

- ▶ $P = TP + FN$ of all cases with the *actual* class 1
- ▶ $N = TN + FP$ of all cases with the *actual* class 0
- ▶ $PP = TP + FP$ of all cases with the *predicted* class 1
- ▶ $PN = TN + FN$ of all cases with the *predicted* class 0

Note that $P + N = PP + PN$ is the number of all cases.

Terminology

- ▶ TP aka hit
- ▶ TN aka correct rejection
- ▶ FP aka type I error, false alarm, overestimation
- ▶ FN aka type II error, miss, underestimation

Usually, TP, TN, FP, and FN are used to denote the individual examples of a particular kind and the number of these examples.

In what follows, we also use

- ▶ $P = TP + FN$ of all cases with the *actual* class 1
- ▶ $N = TN + FP$ of all cases with the *actual* class 0
- ▶ $PP = TP + FP$ of all cases with the *predicted* class 1
- ▶ $PN = TN + FN$ of all cases with the *predicted* class 0

Note that $P + N = PP + PN$ is the number of all cases.

There is a large number of derived metrics. We consider some of the most used in practice.

Accuracy

$$\text{Accuracy} = \frac{TP + TN}{P + N}$$

Intuitively, Accuracy is the proportion of correctly classified cases w.r.t. all cases.

Accuracy

$$\text{Accuracy} = \frac{TP + TN}{P + N}$$

Intuitively, Accuracy is the proportion of correctly classified cases w.r.t. all cases.

Example: Consider our cancer predictor with the confusion matrix

Actual condition	Predicted condition	
	Cancer	Non-cancer
Cancer	TP = 6	FN = 2
Non-cancer	FP = 1	TN = 3
Total	8 + 4 = 12	

Accuracy

$$\text{Accuracy} = \frac{TP + TN}{P + N}$$

Intuitively, Accuracy is the proportion of correctly classified cases w.r.t. all cases.

Example: Consider our cancer predictor with the confusion matrix

Actual condition	Predicted condition	
	Cancer	Non-cancer
Cancer	TP = 6	FN = 2
Non-cancer	FP = 1	TN = 3
Total	8 + 4 = 12	

The Accuracy is

$$\text{ACC} = \frac{TP + TN}{P + N} = \frac{6 + 3}{12} = \frac{3}{4}$$

Accuracy - Imbalanced Classes

Accuracy can be misleading when the classes are imbalanced:

- ▶ Consider 100 cases, 90 in the class 0 and 10 in the class 1,
- ▶ consider a classifier that returns 1 for a single sample of class 1 and 0 for all other samples.

Accuracy - Imbalanced Classes

Accuracy can be misleading when the classes are imbalanced:

- ▶ Consider 100 cases, 90 in the class 0 and 10 in the class 1,
- ▶ consider a classifier that returns 1 for a single sample of class 1 and 0 for all other samples.

Actual	Predicted	
	Pos	Neg
Pos	1	9
Neg	0	90
Total	$90 + 10 = 100$	

Accuracy - Imbalanced Classes

Accuracy can be misleading when the classes are imbalanced:

- ▶ Consider 100 cases, 90 in the class 0 and 10 in the class 1,
- ▶ consider a classifier that returns 1 for a single sample of class 1 and 0 for all other samples.

Actual	Predicted	
	Pos	Neg
Pos	1	9
Neg	0	90
Total	$90 + 10 = 100$	

The Accuracy is $91/100 > 0.9$. Pretty good, right?

Accuracy - Imbalanced Classes

Accuracy can be misleading when the classes are imbalanced:

- ▶ Consider 100 cases, 90 in the class 0 and 10 in the class 1,
- ▶ consider a classifier that returns 1 for a single sample of class 1 and 0 for all other samples.

Actual	Predicted	
	Pos	Neg
Pos	1	9
Neg	0	90
Total	90 + 10 = 100	

The Accuracy is $91/100 > 0.9$. Pretty good, right?

However, the classifier is pretty bad in the positive cases.

In the case of cancer prediction, such a classifier would be a disaster.

Precision & Recall

To mitigate the defect of the Accuracy, we may compute the following metrics:

$$\text{Precision} = \frac{TP}{PP} \quad (= \text{how often is predicted positive actually positive})$$

Precision is also known as positive predictive value (PPV)

Precision & Recall

To mitigate the defect of the Accuracy, we may compute the following metrics:

$$\text{Precision} = \frac{TP}{PP} \quad (= \text{how often is predicted positive actually positive})$$

Precision is also known as positive predictive value (PPV)

$$\text{Recall} = \frac{TP}{P} \quad (= \text{how often is actually positive predicted positive})$$

Recall is also known as true positive rate, sensitivity, hit rate, and power.

Precision & Recall - Example

Example: In our cancer example:

Actual condition	Predicted condition	
	Cancer	Non-cancer
Cancer	TP = 6	FN = 2
Non-cancer	FP = 1	TN = 3
Total	8 + 4 = 12	

Precision & Recall - Example

Example: In our cancer example:

Actual condition	Predicted condition	
	Cancer	Non-cancer
Cancer	TP = 6	FN = 2
Non-cancer	FP = 1	TN = 3
Total	8 + 4 = 12	

- Precision measures how often is the patient predicted to be ill truly ill (in our case, 6/7)

Precision & Recall - Example

Example: In our cancer example:

Actual condition	Predicted condition	
	Cancer	Non-cancer
Cancer	TP = 6	FN = 2
Non-cancer	FP = 1	TN = 3
Total	8 + 4 = 12	

- ▶ Precision measures how often is the patient predicted to be ill truly ill (in our case, 6/7)
- ▶ Recall measures how often is an ill patient found to be ill (in our case, 6/8)

Precision & Recall - Imbalanced Classes

- ▶ Consider 100 cases, 90 in the class 0 and 10 in the class 1,

Precision & Recall - Imbalanced Classes

- ▶ Consider 100 cases, 90 in the class 0 and 10 in the class 1,
- ▶ consider a classifier that returns 1 for a single sample of class 1 and 0 for all other samples.

Actual	Predicted	
	Pos	Neg
Pos	1	9
Neg	0	90
Total	$90 + 10 = 100$	

Precision & Recall - Imbalanced Classes

- ▶ Consider 100 cases, 90 in the class 0 and 10 in the class 1,
- ▶ consider a classifier that returns 1 for a single sample of class 1 and 0 for all other samples.

Actual	Predicted	
	Pos	Neg
Pos	1	9
Neg	0	90
Total	90 + 10 = 100	

$$\text{Precision} = 1$$

$$\text{Recall} = \frac{1}{10}$$

You can see that the predictor is very precise (on the class 1) but useless due to the weak Recall.

Precision & Recall - Relative Importance

Let us get back to our cancer example:

Actual condition	Predicted condition	
	Cancer	Non-cancer
Cancer	TP = 6	FN = 2
Non-cancer	FP = 1	TN = 3
Total	8 + 4 = 12	

Consider *Precision* and *Recall*.

By now, you should remember what they measure.

Precision & Recall - Relative Importance

Let us get back to our cancer example:

Actual condition	Predicted condition	
	Cancer	Non-cancer
Cancer	TP = 6	FN = 2
Non-cancer	FP = 1	TN = 3
Total	8 + 4 = 12	

Consider *Precision* and *Recall*.

By now, you should remember what they measure.

Which of the two is more important in medicine?

Precision & Recall - Relative Importance

Let us get back to our cancer example:

Actual condition	Predicted condition	
	Cancer	Non-cancer
Cancer	TP = 6	FN = 2
Non-cancer	FP = 1	TN = 3
Total	8 + 4 = 12	

Consider *Precision* and *Recall*.

By now, you should remember what they measure.

Which of the two is more important in medicine?

Which of the two is more important for plagiarism detectors?

Precision & Recall - Relative Importance

Let us get back to our cancer example:

Actual condition	Predicted condition	
	Cancer	Non-cancer
Cancer	TP = 6	FN = 2
Non-cancer	FP = 1	TN = 3
Total	8 + 4 = 12	

Consider *Precision* and *Recall*.

By now, you should remember what they measure.

Which of the two is more important in medicine?

Which of the two is more important for plagiarism detectors?

Can we get a single number summarizing both Precision and Recall?

For example, to compare two classifiers.

F_1 Score

F_1 score is the harmonic mean of Recall and Precision:

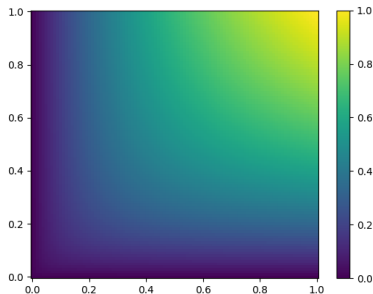
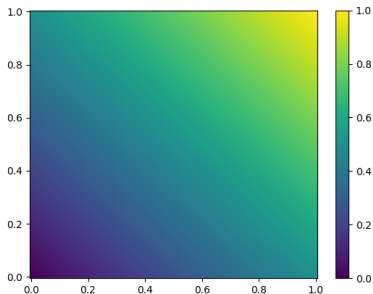
$$F_1 = \frac{2}{\text{Recall}^{-1} + \text{Precision}^{-1}} = \frac{2TP}{2TP + FP + FN}$$

F_1 Score

F_1 score is the harmonic mean of Recall and Precision:

$$F_1 = \frac{2}{\text{Recall}^{-1} + \text{Precision}^{-1}} = \frac{2TP}{2TP + FP + FN}$$

Compare the arithmetic (left) and harmonic (right) mean:



The harmonic mean prefers the two values closer to each other.

For example, the harmonic mean of $2/3$ and $1/3$ is (approx) 0.44444.

F_1 Score - Examples

Consider the cancer example:

Actual condition	Predicted condition	
	Cancer	Non-cancer
Cancer	TP = 6	FN = 2
Non-cancer	FP = 1	TN = 3
Total	8 + 4 = 12	

Here $F_1 = \frac{2TP}{2TP+FP+FN} = (2 \cdot 6)/((2 \cdot 6) + 1 + 2) = 0.8$.

F_1 Score - Examples

Consider the cancer example:

Actual condition	Predicted condition	
	Cancer	Non-cancer
Cancer	TP = 6	FN = 2
Non-cancer	FP = 1	TN = 3
Total	8 + 4 = 12	

Here $F_1 = \frac{2TP}{2TP+FP+FN} = (2 \cdot 6)/((2 \cdot 6) + 1 + 2) = 0.8$.

Our imbalanced example:

Actual	Predicted	
	Pos	Neg
Pos	1	9
Neg	0	90
Total	90 + 10 = 100	

Here $F_1 = \frac{2TP}{2TP+FP+FN} = (2 \cdot 1)/((2 \cdot 1) + 0 + 9) = 0.18$.

Note that the average of Precision and Recall is 0.55, which would give us a much less severe warning that the classifier is bad.

Imbalanced Classes Once More

Note that the standard definitions of Precision and Recall for binary classifiers reveal only part of the truth.

Imbalanced Classes Once More

Note that the standard definitions of Precision and Recall for binary classifiers reveal only part of the truth.

In particular, *true negatives are not used* in the definition of F_1 .

Imbalanced Classes Once More

Note that the standard definitions of Precision and Recall for binary classifiers reveal only part of the truth.

In particular, *true negatives are not used* in the definition of F_1 .

Consider

Actual	Predicted	
	Pos	Neg
Pos	90	0
Neg	9	1
Total	90 + 10 = 100	

Imbalanced Classes Once More

Note that the standard definitions of Precision and Recall for binary classifiers reveal only part of the truth.

In particular, *true negatives are not used* in the definition of F_1 .

Consider

Actual	Predicted	
	Pos	Neg
Pos	90	0
Neg	9	1
Total	90 + 10 = 100	

$$\text{Precision} = 90/99 \quad \text{Recall} = 90/90$$

$$F_1 = \frac{2TP}{2TP + FP + FN} = (2 \cdot 90)/(2 \cdot 90 + 9 + 0) = 0.95$$

Imbalanced Classes Once More

Note that the standard definitions of Precision and Recall for binary classifiers reveal only part of the truth.

In particular, *true negatives are not used* in the definition of F_1 .

Consider

Actual	Predicted	
	Pos	Neg
Pos	90	0
Neg	9	1
Total	90 + 10 = 100	

$$\text{Precision} = 90/99 \quad \text{Recall} = 90/90$$

$$F_1 = \frac{2TP}{2TP + FP + FN} = (2 \cdot 90)/(2 \cdot 90 + 9 + 0) = 0.95$$

All great, except that the classifier sucks on the negative cases.

If you are concerned with the negative cases, swap the classes and compute another set of metrics.

F_1 Score

- ▶ F_1 is often used as a summary score for binary classifiers instead of Accuracy.

Works better with imbalanced classes.

F_1 Score

- ▶ F_1 is often used as a summary score for binary classifiers instead of Accuracy.
Works better with imbalanced classes.
- ▶ Criticised for giving Precision and Recall the same importance.
- ▶ Is not symmetric, ignores true negatives, i.e., is misleading for some cases of imbalanced classes.

F_1 Score

- ▶ F_1 is often used as a summary score for binary classifiers instead of Accuracy.
Works better with imbalanced classes.
- ▶ Criticised for giving Precision and Recall the same importance.
- ▶ Is not symmetric, ignores true negatives, i.e., is misleading for some cases of imbalanced classes.
- ▶ *Fowlkes-Mallows index* is a geometric mean of Precision and Recall (used in clustering).
The geometric mean is between the arithmetic and harmonic mean. For example, the geometric mean of $2/3$ and $1/3$ is (approx) 0.4714.

More Derived Metrics

<p>Positive predictive value (PPV), precision</p> $= \frac{TP}{PP} = 1 - FDR$	<p>False omission rate (FOR)</p> $= \frac{FN}{PN} = 1 - NPV$
<p>False discovery rate (FDR)</p> $= \frac{FP}{PP} = 1 - PPV$	<p>Negative predictive value (NPV)</p> $= \frac{TN}{PN} = 1 - FOR$

You can see that the negative predictive value becomes the Precision when we swap the classes (and vice versa).

More Derived Metrics

<p>True positive rate (TPR), recall, sensitivity (SEN), probability of detection, hit rate, power</p> $= \frac{TP}{P} = 1 - FNR$	<p>False negative rate (FNR), miss rate</p> $= \frac{FN}{P} = 1 - TPR$
<p>False positive rate (FPR), probability of false alarm, fall-out</p> $= \frac{FP}{N} = 1 - TNR$	<p>True negative rate (TNR), specificity (SPC), selectivity</p> $= \frac{TN}{N} = 1 - FPR$

Note that *specificity* becomes Recall when we swap the classes (and vice versa).

For example, medical doctors communicate in terms of *sensitivity* and *specificity*.

Actual condition	Predicted condition	
	Cancer	Non-cancer
Cancer	TP = 6	FN = 2
Non-cancer	FP = 1	TN = 3
Total	8 + 4 = 12	

Actual condition	Predicted condition	
	Cancer	Non-cancer
Cancer	TP = 6	FN = 2
Non-cancer	FP = 1	TN = 3
Total	8 + 4 = 12	

$$\text{TPR} = \text{Sensitivity} = \text{Recall} = \text{TP}/P = 6/8$$

How often is positive predicted positive?

Actual condition	Predicted condition	
	Cancer	Non-cancer
Cancer	TP = 6	FN = 2
Non-cancer	FP = 1	TN = 3
Total	8 + 4 = 12	

$$\text{TPR} = \text{Sensitivity} = \text{Recall} = \text{TP}/P = 6/8$$

How often is positive predicted positive?

$$\text{TNR} = \text{Specificity} = \text{TN}/N = 3/4$$

How often is negative predicted negative?

Actual condition	Predicted condition	
	Cancer	Non-cancer
Cancer	TP = 6	FN = 2
Non-cancer	FP = 1	TN = 3
Total	8 + 4 = 12	

$$\text{TPR} = \text{Sensitivity} = \text{Recall} = \text{TP}/P = 6/8$$

How often is positive predicted positive?

$$\text{TNR} = \text{Specificity} = \text{TN}/N = 3/4$$

How often is negative predicted negative?

$$\text{FPR} = \text{Prob. of false alarm} = \text{FP}/N = 1/4$$

How often is negative predicted positive?

Actual condition	Predicted condition	
	Cancer	Non-cancer
Cancer	TP = 6	FN = 2
Non-cancer	FP = 1	TN = 3
Total	8 + 4 = 12	

$$\text{TPR} = \text{Sensitivity} = \text{Recall} = \text{TP}/P = 6/8$$

How often is positive predicted positive?

$$\text{TNR} = \text{Specificity} = \text{TN}/N = 3/4$$

How often is negative predicted negative?

$$\text{FPR} = \text{Prob. of false alarm} = \text{FP}/N = 1/4$$

How often is negative predicted positive?

$$\text{FNR} = \text{Miss rate} = \text{FN}/P = 2/8$$

How often is positive predicted negative?

Evaluating Multi-class Classifiers

Classification Into Multiple Classes

Assume classification into classes from a finite set C .

Classification Into Multiple Classes

Assume classification into classes from a finite set C .

Consider a classification dataset:

$$\{(\vec{x}_k, c_k) \mid k = 1, \dots, p\}$$

Here \vec{x}_k is a vector of attributes/features and $c_k \in C$ for all k .

Classification Into Multiple Classes

Assume classification into classes from a finite set C .

Consider a classification dataset:

$$\{(\vec{x}_k, c_k) \mid k = 1, \dots, p\}$$

Here \vec{x}_k is a vector of attributes/features and $c_k \in C$ for all k .

Consider a sequence of predictions generated by a classifier:

$$h_1, \dots, h_p \in C$$

Here each h_k has been predicted for the k -th example (\vec{x}_k, c_k) .

Classification Into Multiple Classes

Assume classification into classes from a finite set C .

Consider a classification dataset:

$$\{(\vec{x}_k, c_k) \mid k = 1, \dots, p\}$$

Here \vec{x}_k is a vector of attributes/features and $c_k \in C$ for all k .

Consider a sequence of predictions generated by a classifier:

$$h_1, \dots, h_p \in C$$

Here each h_k has been predicted for the k -th example (\vec{x}_k, c_k) .

How good are the predictions h_1, \dots, h_p w.r.t. c_1, \dots, c_p ?

There are many possible metrics ...

Consider an arbitrary (finite) number of classes in C .

Confusion Matrix

Assume that $C = \{1, \dots, m\}$.

Confusion Matrix

Assume that $C = \{1, \dots, m\}$.

Now, given two classes $i, j \in C$ we denote by M_{ij} the number of samples of class i classified into the class j .

Confusion Matrix

Assume that $C = \{1, \dots, m\}$.

Now, given two classes $i, j \in C$ we denote by M_{ij} the number of samples of class i classified into the class j .

Formally,

$$M_{ij} = |\{k \mid c_k = i \wedge h_k = j\}|$$

Actual	Predicted				
	1	...	j	...	m
1	M_{11}	...	M_{1j}	...	M_{1m}
\vdots	\vdots		\vdots		\vdots
i	M_{i1}	...	M_{ij}	...	M_{im}
\vdots	\vdots		\vdots		\vdots
m	M_{m1}	...	M_{mj}	...	M_{mm}

Example

Actual	Predicted
big	big
big	big
small	big
medium	medium
big	small
big	big
small	small
small	small
medium	medium
medium	small
small	small
big	big
medium	small
small	medium
big	big

Example

Actual	Predicted
big	big
big	big
small	big
medium	medium
big	small
big	big
small	small
small	small
medium	medium
medium	small
small	small
big	big
medium	small
small	medium
big	big

Actual	Predicted		
	big	medium	small
big	5	0	1
medium	0	2	2
small	1	1	3

Note that the diagonal counts the correctly classified samples.

The off-diagonal elements correspond to misclassified samples.

Metrics

We can easily generalize Accuracy, Precision, Recall, and F_1 -score from the binary classification to multiple classes.

Metrics

We can easily generalize Accuracy, Precision, Recall, and F_1 -score from the binary classification to multiple classes.

Notation

$$\blacktriangleright M_{i\bullet} = \sum_{j=1}^m M_{ij}$$

Metrics

We can easily generalize Accuracy, Precision, Recall, and F_1 -score from the binary classification to multiple classes.

Notation

- ▶ $M_{i\bullet} = \sum_{j=1}^m M_{ij}$
- ▶ $M_{\bullet j} = \sum_{i=1}^m M_{ij}$

Metrics

We can easily generalize Accuracy, Precision, Recall, and F_1 -score from the binary classification to multiple classes.

Notation

- ▶ $M_{i\bullet} = \sum_{j=1}^m M_{ij}$
- ▶ $M_{\bullet j} = \sum_{i=1}^m M_{ij}$
- ▶ $M_{\bullet\bullet} = \sum_{i=1}^m \sum_{j=1}^m M_{ij}$

Metrics

We can easily generalize Accuracy, Precision, Recall, and F_1 -score from the binary classification to multiple classes.

Notation

- ▶ $M_{i\bullet} = \sum_{j=1}^m M_{ij}$
- ▶ $M_{\bullet j} = \sum_{i=1}^m M_{ij}$
- ▶ $M_{\bullet\bullet} = \sum_{i=1}^m \sum_{j=1}^m M_{ij}$

Now, the metrics:

$$\text{Accuracy} = \frac{\sum_{k=1}^m M_{kk}}{M_{\bullet\bullet}}$$

Metrics

We can easily generalize Accuracy, Precision, Recall, and F_1 -score from the binary classification to multiple classes.

Notation

- ▶ $M_{i\bullet} = \sum_{j=1}^m M_{ij}$
- ▶ $M_{\bullet j} = \sum_{i=1}^m M_{ij}$
- ▶ $M_{\bullet\bullet} = \sum_{i=1}^m \sum_{j=1}^m M_{ij}$

Now, the metrics:

$$\text{Accuracy} = \frac{\sum_{k=1}^m M_{kk}}{M_{\bullet\bullet}}$$

For a given class $i \in C$:

$$\text{Precision}[i] = \frac{M_{ii}}{M_{\bullet i}} \quad \text{Recall}[i] = \frac{M_{ii}}{M_{i\bullet}}$$

Metrics

We can easily generalize Accuracy, Precision, Recall, and F_1 -score from the binary classification to multiple classes.

Notation

- ▶ $M_{i\bullet} = \sum_{j=1}^m M_{ij}$
- ▶ $M_{\bullet j} = \sum_{i=1}^m M_{ij}$
- ▶ $M_{\bullet\bullet} = \sum_{i=1}^m \sum_{j=1}^m M_{ij}$

Now, the metrics:

$$\text{Accuracy} = \frac{\sum_{k=1}^m M_{kk}}{M_{\bullet\bullet}}$$

For a given class $i \in C$:

$$\text{Precision}[i] = \frac{M_{ii}}{M_{\bullet i}} \quad \text{Recall}[i] = \frac{M_{ii}}{M_{i\bullet}}$$

$$F_1[i] = \frac{2 * \text{Precision}[i] * \text{Recall}[i]}{\text{Precision}[i] + \text{Recall}[i]}$$

Note that Precision, Recall, and F_1 can be defined only for a given class!

Example

Actual	Predicted		
	big	medium	small
big	5	0	1
medium	0	2	2
small	1	1	3

Compute the metrics.

Example

$$\text{Accuracy} = (5+2+3)/15 = 0.66$$

$$\text{Precision}[\text{big}] = 5/6$$

$$\text{Precision}[\text{medium}] = 2/3$$

$$\text{Precision}[\text{small}] = 3/6$$

$$\text{Recall}[\text{big}] = 5/6$$

$$\text{Recall}[\text{medium}] = 2/4$$

$$\text{Recall}[\text{small}] = 3/5$$

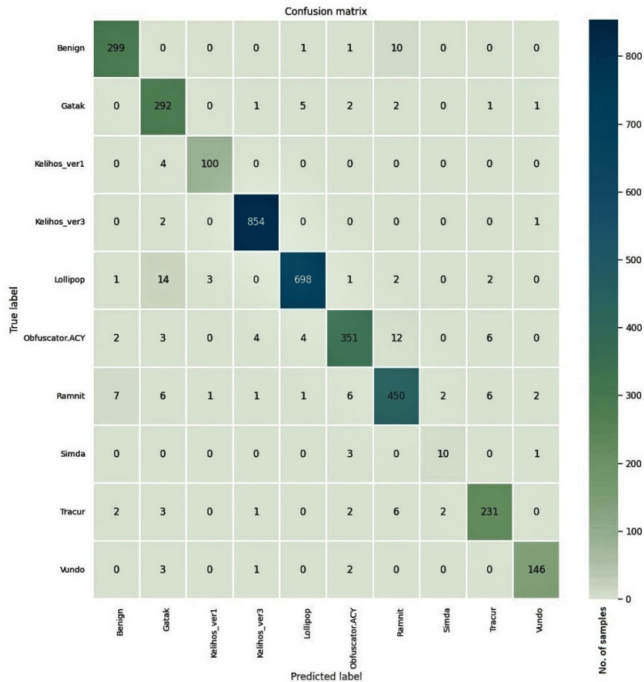
$$F_1[\text{big}] = \frac{2 * (5/6) * (5/6)}{(5/6) + (5/6)} = 5/6 = 0.83$$

$$F_1[\text{medium}] = 0.57$$

$$F_1[\text{small}] = 0.54$$

Actual	Predicted		
	big	medium	small
big	5	0	1
medium	0	2	2
small	1	1	3

How do you get a single number out of these? Average Precision, Recall, and F_1 are usually computed, but one needs to be careful about the variance.



Probabilistic Classifier Evaluation

Binary Probabilistic Classifier

Assume binary classification into two classes $\{0, 1\}$.

Binary Probabilistic Classifier

Assume binary classification into two classes $\{0, 1\}$.

Consider a classification dataset:

$$\{(\vec{x}_k, c_k) \mid k = 1, \dots, p\}$$

Here \vec{x}_k is a vector of attributes/features and $c_k \in C$ for all k .

Binary Probabilistic Classifier

Assume binary classification into two classes $\{0, 1\}$.

Consider a classification dataset:

$$\{(\vec{x}_k, c_k) \mid k = 1, \dots, p\}$$

Here \vec{x}_k is a vector of attributes/features and $c_k \in C$ for all k .

Consider a sequence of predictions generated by a classifier.

Now the classifier returns *probability of class 1* for a given input:

$$h_1, \dots, h_p \in [0, 1]$$

Here each h_k has been predicted for the k -th example (\vec{x}_k, c_k) .

Binary Probabilistic Classifier

Assume binary classification into two classes $\{0, 1\}$.

Consider a classification dataset:

$$\{(\vec{x}_k, c_k) \mid k = 1, \dots, p\}$$

Here \vec{x}_k is a vector of attributes/features and $c_k \in C$ for all k .

Consider a sequence of predictions generated by a classifier.

Now the classifier returns *probability of class 1* for a given input:

$$h_1, \dots, h_p \in [0, 1]$$

Here each h_k has been predicted for the k -th example (\vec{x}_k, c_k) .

How to interpret the predictions h_1, \dots, h_p ?

Binary Probabilistic Classifier

Assume binary classification into two classes $\{0, 1\}$.

Consider a classification dataset:

$$\{(\vec{x}_k, c_k) \mid k = 1, \dots, p\}$$

Here \vec{x}_k is a vector of attributes/features and $c_k \in C$ for all k .

Consider a sequence of predictions generated by a classifier.

Now the classifier returns *probability of class 1* for a given input:

$$h_1, \dots, h_p \in [0, 1]$$

Here each h_k has been predicted for the k -th example (\vec{x}_k, c_k) .

How to interpret the predictions h_1, \dots, h_p ?

How good are the predictions h_1, \dots, h_p w.r.t. c_1, \dots, c_p ?

Probabilistic Classifier

Let us fix predictions h_1, \dots, h_p .

Probabilistic Classifier

Let us fix predictions h_1, \dots, h_p .

Given a threshold $T \in [0, 1]$ we define

$$h_k^T = \begin{cases} 1 & \text{if } h_k \geq T \\ 0 & \text{if } h_k < T \end{cases}$$

For every T we can compute all the metrics (Precision, Recall, etc.)

Probabilistic Classifier

Let us fix predictions h_1, \dots, h_p .

Given a threshold $T \in [0, 1]$ we define

$$h_k^T = \begin{cases} 1 & \text{if } h_k \geq T \\ 0 & \text{if } h_k < T \end{cases}$$

For every T we can compute all the metrics (Precision, Recall, etc.)

Given a metric MET and a threshold T , we denote by $\text{MET}[T]$ the metric MET evaluated on h_1^T, \dots, h_p^T .

Probabilistic Classifier

Let us fix predictions h_1, \dots, h_p .

Given a threshold $T \in [0, 1]$ we define

$$h_k^T = \begin{cases} 1 & \text{if } h_k \geq T \\ 0 & \text{if } h_k < T \end{cases}$$

For every T we can compute all the metrics (Precision, Recall, etc.)

Given a metric MET and a threshold T , we denote by $\text{MET}[T]$ the metric MET evaluated on h_1^T, \dots, h_p^T .

We obtain

$$\text{TP}[T] = |\{k \mid h_k^T = 1 \wedge c_k = 1\}|$$

Probabilistic Classifier

Let us fix predictions h_1, \dots, h_p .

Given a threshold $T \in [0, 1]$ we define

$$h_k^T = \begin{cases} 1 & \text{if } h_k \geq T \\ 0 & \text{if } h_k < T \end{cases}$$

For every T we can compute all the metrics (Precision, Recall, etc.)

Given a metric MET and a threshold T , we denote by MET[T] the metric MET evaluated on h_1^T, \dots, h_p^T .

We obtain

$$\text{TP}[T] = |\{k \mid h_k^T = 1 \wedge c_k = 1\}|$$

and

TN[T], FP[T], FN[T], Accuracy[T], Precision[T], Recall[T], F_1 [T], \dots

Probabilistic Classifier

Let us fix predictions h_1, \dots, h_p .

Given a threshold $T \in [0, 1]$ we define

$$h_k^T = \begin{cases} 1 & \text{if } h_k \geq T \\ 0 & \text{if } h_k < T \end{cases}$$

For every T we can compute all the metrics (Precision, Recall, etc.)

Given a metric MET and a threshold T , we denote by $\text{MET}[T]$ the metric MET evaluated on h_1^T, \dots, h_p^T .

We obtain

$$\text{TP}[T] = |\{k \mid h_k^T = 1 \wedge c_k = 1\}|$$

and

$\text{TN}[T], \text{FP}[T], \text{FN}[T], \text{Accuracy}[T], \text{Precision}[T], \text{Recall}[T], F_1[T], \dots$

However, all metrics are now functions of the threshold T .

Thresholded Classifier Metrics

Index	1	2	3	4	5	6	7	8	9	10	11	12
Actual	1	1	1	1	1	0	0	1	1	0	0	0
Predicted	.98	.95	.9	.86	.66	.48	.43	.42	.36	.15	.1	.05
T=0.5	TP	TP	TP	TP	TP	TN	TN	FN	FN	TN	TN	TN
T=0.42	TP	TP	TP	TP	TP	FP	FP	TP	FN	TN	TN	TN
T=0.1	TP	TP	TP	TP	TP	FP	FP	TP	TP	FP	FP	TN

Thresholded Classifier Metrics

Index	1	2	3	4	5	6	7	8	9	10	11	12
Actual	1	1	1	1	1	0	0	1	1	0	0	0
Predicted	.98	.95	.9	.86	.66	.48	.43	.42	.36	.15	.1	.05
T=0.5	TP	TP	TP	TP	TP	TN	TN	FN	FN	TN	TN	TN
T=0.42	TP	TP	TP	TP	TP	FP	FP	TP	FN	TN	TN	TN
T=0.1	TP	TP	TP	TP	TP	FP	FP	TP	TP	FP	FP	TN

For example, consider $T = 0.42$, then

Thresholded Classifier Metrics

Index	1	2	3	4	5	6	7	8	9	10	11	12
Actual	1	1	1	1	1	0	0	1	1	0	0	0
Predicted	.98	.95	.9	.86	.66	.48	.43	.42	.36	.15	.1	.05
T=0.5	TP	TP	TP	TP	TP	TN	TN	FN	FN	TN	TN	TN
T=0.42	TP	TP	TP	TP	TP	FP	FP	TP	FN	TN	TN	TN
T=0.1	TP	TP	TP	TP	TP	FP	FP	TP	TP	FP	FP	TN

For example, consider $T = 0.42$, then

$$TP[T] = 6 \quad FP[T] = 2 \quad FN[T] = 1 \quad TN[T] = 3$$

Thresholded Classifier Metrics

Index	1	2	3	4	5	6	7	8	9	10	11	12
Actual	1	1	1	1	1	0	0	1	1	0	0	0
Predicted	.98	.95	.9	.86	.66	.48	.43	.42	.36	.15	.1	.05
T=0.5	TP	TP	TP	TP	TP	TN	TN	FN	FN	TN	TN	TN
T=0.42	TP	TP	TP	TP	TP	FP	FP	TP	FN	TN	TN	TN
T=0.1	TP	TP	TP	TP	TP	FP	FP	TP	TP	FP	FP	TN

For example, consider $T = 0.42$, then

$$TP[T] = 6 \quad FP[T] = 2 \quad FN[T] = 1 \quad TN[T] = 3$$

$$Accuracy[T] = \frac{3+6}{12} \quad Precision[T] = \frac{6}{6+2} \quad Recall[T] = \frac{6}{6+1}$$

Thresholded Classifier Metrics

Index	1	2	3	4	5	6	7	8	9	10	11	12
Actual	1	1	1	1	1	0	0	1	1	0	0	0
Predicted	.98	.95	.9	.86	.66	.48	.43	.42	.36	.15	.1	.05
T=0.5	TP	TP	TP	TP	TP	TN	TN	FN	FN	TN	TN	TN
T=0.42	TP	TP	TP	TP	TP	FP	FP	TP	FN	TN	TN	TN
T=0.1	TP	TP	TP	TP	TP	FP	FP	TP	TP	FP	FP	TN

For example, consider $T = 0.42$, then

$$TP[T] = 6 \quad FP[T] = 2 \quad FN[T] = 1 \quad TN[T] = 3$$

$$Accuracy[T] = \frac{3+6}{12} \quad Precision[T] = \frac{6}{6+2} \quad Recall[T] = \frac{6}{6+1}$$

$$F_1[T] = \frac{2 \cdot 6/8 \cdot 6/7}{6/8 + 6/7} = 0.8$$

Receiver Operating Characteristic (ROC)

Consider two metrics for a given T :

$$\text{TPR}[T] = \frac{\text{TP}[T]}{\text{P}[T]} \quad (\text{True Positive Rate})$$

Receiver Operating Characteristic (ROC)

Consider two metrics for a given T :

$$\text{TPR}[T] = \frac{\text{TP}[T]}{P[T]} \quad (\text{True Positive Rate})$$

$$\text{FPR}[T] = \frac{\text{FP}[T]}{N[T]} \quad (\text{False Positive Rate})$$

Receiver Operating Characteristic (ROC)

Consider two metrics for a given T :

$$\text{TPR}[T] = \frac{\text{TP}[T]}{P[T]} \quad (\text{True Positive Rate})$$

$$\text{FPR}[T] = \frac{\text{FP}[T]}{N[T]} \quad (\text{False Positive Rate})$$

ROC curve is then a function $\text{ROC} : [0, 1] \rightarrow [0, 1]^2$ defined by

$$\text{ROC}(T) = (\text{TPR}[T], \text{FPR}[T])$$

Receiver Operating Characteristic (ROC)

Consider two metrics for a given T :

$$\text{TPR}[T] = \frac{\text{TP}[T]}{P[T]} \quad (\text{True Positive Rate})$$

$$\text{FPR}[T] = \frac{\text{FP}[T]}{N[T]} \quad (\text{False Positive Rate})$$

ROC curve is then a function $\text{ROC} : [0, 1] \rightarrow [0, 1]^2$ defined by

$$\text{ROC}(T) = (\text{TPR}[T], \text{FPR}[T])$$

Observe that

$$\text{ROC}(0) = (1, 1)$$

Because the classifier with $T = 0$ simply classifies everything as positive, i.e., into the class 1.

Both $\text{TPR}[T]$ and $\text{FPR}[T]$ are non-increasing in T .

Index	1	2	3	4	5	6	7	8	9	10	11	12
Actual	1	1	1	1	1	0	0	1	1	0	0	0
Predicted	.98	.95	.9	.86	.66	.48	.43	.42	.36	.15	.1	.05

Index	1	2	3	4	5	6	7	8	9	10	11	12
Actual	1	1	1	1	1	0	0	1	1	0	0	0
Predicted	.98	.95	.9	.86	.66	.48	.43	.42	.36	.15	.1	.05

► $0.00 \leq T \leq 0.05$: $\text{TPR} = 1$ and $\text{FPR} = 1$

Index	1	2	3	4	5	6	7	8	9	10	11	12
Actual	1	1	1	1	1	0	0	1	1	0	0	0
Predicted	.98	.95	.9	.86	.66	.48	.43	.42	.36	.15	.1	.05

- ▶ $0.00 \leq T \leq 0.05$: $\text{TPR} = 1$ and $\text{FPR} = 1$
- ▶ $0.05 < T \leq 0.10$: $\text{TPR} = 1$ and $\text{FPR} = 4/5$

Index	1	2	3	4	5	6	7	8	9	10	11	12
Actual	1	1	1	1	1	0	0	1	1	0	0	0
Predicted	.98	.95	.9	.86	.66	.48	.43	.42	.36	.15	.1	.05

- ▶ $0.00 \leq T \leq 0.05$: $\text{TPR} = 1$ and $\text{FPR} = 1$
- ▶ $0.05 < T \leq 0.10$: $\text{TPR} = 1$ and $\text{FPR} = 4/5$
- ▶ $0.10 < T \leq 0.15$: $\text{TPR} = 1$ and $\text{FPR} = 3/5$

Index	1	2	3	4	5	6	7	8	9	10	11	12
Actual	1	1	1	1	1	0	0	1	1	0	0	0
Predicted	.98	.95	.9	.86	.66	.48	.43	.42	.36	.15	.1	.05

- ▶ $0.00 \leq T \leq 0.05$: $\text{TPR} = 1$ and $\text{FPR} = 1$
- ▶ $0.05 < T \leq 0.10$: $\text{TPR} = 1$ and $\text{FPR} = 4/5$
- ▶ $0.10 < T \leq 0.15$: $\text{TPR} = 1$ and $\text{FPR} = 3/5$
- ▶ $0.15 < T \leq 0.36$: $\text{TPR} = 1$ and $\text{FPR} = 2/5$

Index	1	2	3	4	5	6	7	8	9	10	11	12
Actual	1	1	1	1	1	0	0	1	1	0	0	0
Predicted	.98	.95	.9	.86	.66	.48	.43	.42	.36	.15	.1	.05

- ▶ $0.00 \leq T \leq 0.05$: $\text{TPR} = 1$ and $\text{FPR} = 1$
- ▶ $0.05 < T \leq 0.10$: $\text{TPR} = 1$ and $\text{FPR} = 4/5$
- ▶ $0.10 < T \leq 0.15$: $\text{TPR} = 1$ and $\text{FPR} = 3/5$
- ▶ $0.15 < T \leq 0.36$: $\text{TPR} = 1$ and $\text{FPR} = 2/5$
- ▶ $0.36 < T \leq 0.42$: $\text{TPR} = 6/7$ and $\text{FPR} = 2/5$

Index	1	2	3	4	5	6	7	8	9	10	11	12
Actual	1	1	1	1	1	0	0	1	1	0	0	0
Predicted	.98	.95	.9	.86	.66	.48	.43	.42	.36	.15	.1	.05

- ▶ $0.00 \leq T \leq 0.05$: $\text{TPR} = 1$ and $\text{FPR} = 1$
- ▶ $0.05 < T \leq 0.10$: $\text{TPR} = 1$ and $\text{FPR} = 4/5$
- ▶ $0.10 < T \leq 0.15$: $\text{TPR} = 1$ and $\text{FPR} = 3/5$
- ▶ $0.15 < T \leq 0.36$: $\text{TPR} = 1$ and $\text{FPR} = 2/5$
- ▶ $0.36 < T \leq 0.42$: $\text{TPR} = 6/7$ and $\text{FPR} = 2/5$
- ▶ $0.42 < T \leq 0.43$: $\text{TPR} = 5/7$ and $\text{FPR} = 2/5$

Index	1	2	3	4	5	6	7	8	9	10	11	12
Actual	1	1	1	1	1	0	0	1	1	0	0	0
Predicted	.98	.95	.9	.86	.66	.48	.43	.42	.36	.15	.1	.05

- ▶ $0.00 \leq T \leq 0.05$: $\text{TPR} = 1$ and $\text{FPR} = 1$
- ▶ $0.05 < T \leq 0.10$: $\text{TPR} = 1$ and $\text{FPR} = 4/5$
- ▶ $0.10 < T \leq 0.15$: $\text{TPR} = 1$ and $\text{FPR} = 3/5$
- ▶ $0.15 < T \leq 0.36$: $\text{TPR} = 1$ and $\text{FPR} = 2/5$
- ▶ $0.36 < T \leq 0.42$: $\text{TPR} = 6/7$ and $\text{FPR} = 2/5$
- ▶ $0.42 < T \leq 0.43$: $\text{TPR} = 5/7$ and $\text{FPR} = 2/5$
- ▶ $0.43 < T \leq 0.48$: $\text{TPR} = 5/7$ and $\text{FPR} = 1/5$

Index	1	2	3	4	5	6	7	8	9	10	11	12
Actual	1	1	1	1	1	0	0	1	1	0	0	0
Predicted	.98	.95	.9	.86	.66	.48	.43	.42	.36	.15	.1	.05

- ▶ $0.00 \leq T \leq 0.05$: TPR = 1 and FPR = 1
- ▶ $0.05 < T \leq 0.10$: TPR = 1 and FPR = 4/5
- ▶ $0.10 < T \leq 0.15$: TPR = 1 and FPR = 3/5
- ▶ $0.15 < T \leq 0.36$: TPR = 1 and FPR = 2/5
- ▶ $0.36 < T \leq 0.42$: TPR = 6/7 and FPR = 2/5
- ▶ $0.42 < T \leq 0.43$: TPR = 5/7 and FPR = 2/5
- ▶ $0.43 < T \leq 0.48$: TPR = 5/7 and FPR = 1/5
- ▶ $0.48 < T \leq 0.66$: TPR = 5/7 and FPR = 0

Index	1	2	3	4	5	6	7	8	9	10	11	12
Actual	1	1	1	1	1	0	0	1	1	0	0	0
Predicted	.98	.95	.9	.86	.66	.48	.43	.42	.36	.15	.1	.05

- ▶ $0.00 \leq T \leq 0.05$: TPR = 1 and FPR = 1
- ▶ $0.05 < T \leq 0.10$: TPR = 1 and FPR = 4/5
- ▶ $0.10 < T \leq 0.15$: TPR = 1 and FPR = 3/5
- ▶ $0.15 < T \leq 0.36$: TPR = 1 and FPR = 2/5
- ▶ $0.36 < T \leq 0.42$: TPR = 6/7 and FPR = 2/5
- ▶ $0.42 < T \leq 0.43$: TPR = 5/7 and FPR = 2/5
- ▶ $0.43 < T \leq 0.48$: TPR = 5/7 and FPR = 1/5
- ▶ $0.48 < T \leq 0.66$: TPR = 5/7 and FPR = 0
- ▶ $0.66 < T \leq 0.86$: TPR = 4/7 and FPR = 0

Index	1	2	3	4	5	6	7	8	9	10	11	12
Actual	1	1	1	1	1	0	0	1	1	0	0	0
Predicted	.98	.95	.9	.86	.66	.48	.43	.42	.36	.15	.1	.05

- ▶ $0.00 \leq T \leq 0.05$: TPR = 1 and FPR = 1
- ▶ $0.05 < T \leq 0.10$: TPR = 1 and FPR = 4/5
- ▶ $0.10 < T \leq 0.15$: TPR = 1 and FPR = 3/5
- ▶ $0.15 < T \leq 0.36$: TPR = 1 and FPR = 2/5
- ▶ $0.36 < T \leq 0.42$: TPR = 6/7 and FPR = 2/5
- ▶ $0.42 < T \leq 0.43$: TPR = 5/7 and FPR = 2/5
- ▶ $0.43 < T \leq 0.48$: TPR = 5/7 and FPR = 1/5
- ▶ $0.48 < T \leq 0.66$: TPR = 5/7 and FPR = 0
- ▶ $0.66 < T \leq 0.86$: TPR = 4/7 and FPR = 0
- ▶ $0.86 < T \leq 0.90$: TPR = 3/7 and FPR = 0

Index	1	2	3	4	5	6	7	8	9	10	11	12
Actual	1	1	1	1	1	0	0	1	1	0	0	0
Predicted	.98	.95	.9	.86	.66	.48	.43	.42	.36	.15	.1	.05

- ▶ $0.00 \leq T \leq 0.05$: TPR = 1 and FPR = 1
- ▶ $0.05 < T \leq 0.10$: TPR = 1 and FPR = 4/5
- ▶ $0.10 < T \leq 0.15$: TPR = 1 and FPR = 3/5
- ▶ $0.15 < T \leq 0.36$: TPR = 1 and FPR = 2/5
- ▶ $0.36 < T \leq 0.42$: TPR = 6/7 and FPR = 2/5
- ▶ $0.42 < T \leq 0.43$: TPR = 5/7 and FPR = 2/5
- ▶ $0.43 < T \leq 0.48$: TPR = 5/7 and FPR = 1/5
- ▶ $0.48 < T \leq 0.66$: TPR = 5/7 and FPR = 0
- ▶ $0.66 < T \leq 0.86$: TPR = 4/7 and FPR = 0
- ▶ $0.86 < T \leq 0.90$: TPR = 3/7 and FPR = 0
- ▶ $0.90 < T \leq 0.95$: TPR = 2/7 and FPR = 0

Index	1	2	3	4	5	6	7	8	9	10	11	12
Actual	1	1	1	1	1	0	0	1	1	0	0	0
Predicted	.98	.95	.9	.86	.66	.48	.43	.42	.36	.15	.1	.05

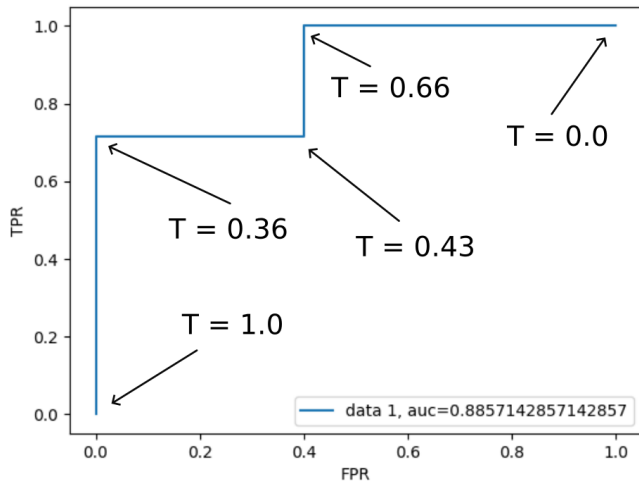
- ▶ $0.00 \leq T \leq 0.05$: TPR = 1 and FPR = 1
- ▶ $0.05 < T \leq 0.10$: TPR = 1 and FPR = 4/5
- ▶ $0.10 < T \leq 0.15$: TPR = 1 and FPR = 3/5
- ▶ $0.15 < T \leq 0.36$: TPR = 1 and FPR = 2/5
- ▶ $0.36 < T \leq 0.42$: TPR = 6/7 and FPR = 2/5
- ▶ $0.42 < T \leq 0.43$: TPR = 5/7 and FPR = 2/5
- ▶ $0.43 < T \leq 0.48$: TPR = 5/7 and FPR = 1/5
- ▶ $0.48 < T \leq 0.66$: TPR = 5/7 and FPR = 0
- ▶ $0.66 < T \leq 0.86$: TPR = 4/7 and FPR = 0
- ▶ $0.86 < T \leq 0.90$: TPR = 3/7 and FPR = 0
- ▶ $0.90 < T \leq 0.95$: TPR = 2/7 and FPR = 0
- ▶ $0.95 < T \leq 0.98$: TPR = 1/7 and FPR = 0

Index	1	2	3	4	5	6	7	8	9	10	11	12
Actual	1	1	1	1	1	0	0	1	1	0	0	0
Predicted	.98	.95	.9	.86	.66	.48	.43	.42	.36	.15	.1	.05

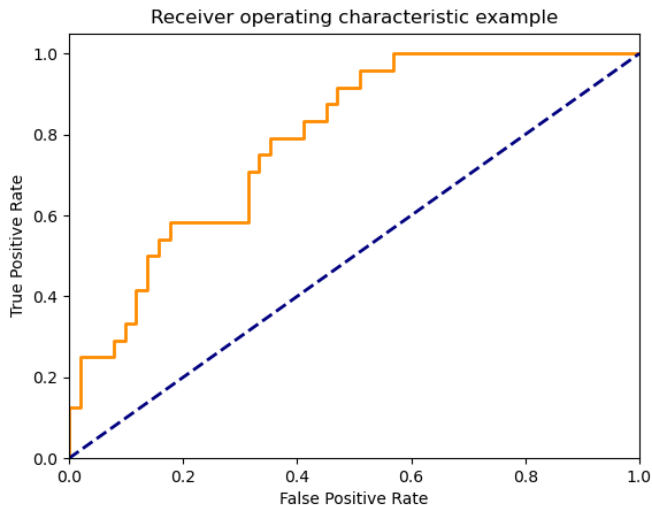
- ▶ $0.00 \leq T \leq 0.05$: TPR = 1 and FPR = 1
- ▶ $0.05 < T \leq 0.10$: TPR = 1 and FPR = 4/5
- ▶ $0.10 < T \leq 0.15$: TPR = 1 and FPR = 3/5
- ▶ $0.15 < T \leq 0.36$: TPR = 1 and FPR = 2/5
- ▶ $0.36 < T \leq 0.42$: TPR = 6/7 and FPR = 2/5
- ▶ $0.42 < T \leq 0.43$: TPR = 5/7 and FPR = 2/5
- ▶ $0.43 < T \leq 0.48$: TPR = 5/7 and FPR = 1/5
- ▶ $0.48 < T \leq 0.66$: TPR = 5/7 and FPR = 0
- ▶ $0.66 < T \leq 0.86$: TPR = 4/7 and FPR = 0
- ▶ $0.86 < T \leq 0.90$: TPR = 3/7 and FPR = 0
- ▶ $0.90 < T \leq 0.95$: TPR = 2/7 and FPR = 0
- ▶ $0.95 < T \leq 0.98$: TPR = 1/7 and FPR = 0
- ▶ $0.98 < T \leq 1.00$: TPR = 0 and FPR = 0

ROC

Index	1	2	3	4	5	6	7	8	9	10	11	12
Actual	1	1	1	1	1	0	0	1	1	0	0	0
Predicted	.98	.95	.9	.86	.66	.48	.43	.42	.36	.15	.1	.05

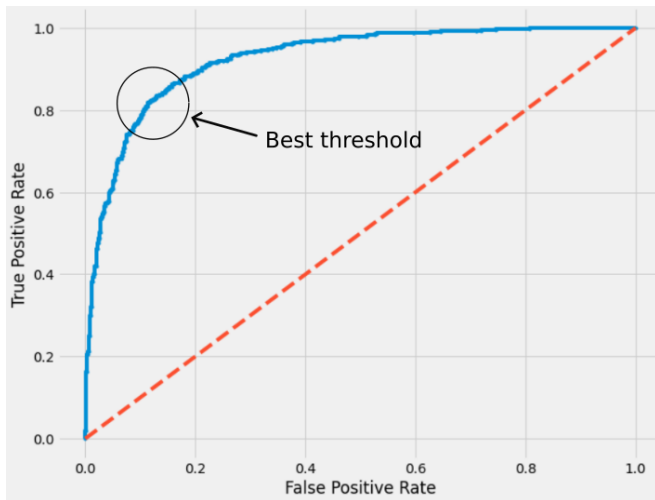


Iris Dataset - A Classifier



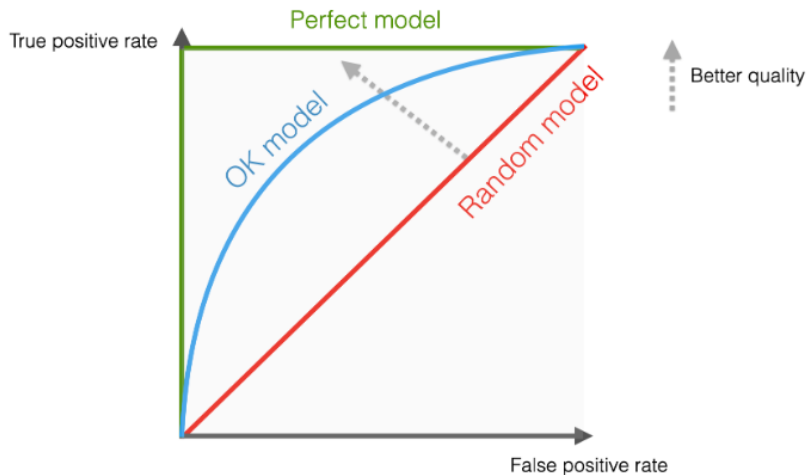
Example from the scikit-learn manual - SVM classifier trained in Iris

Using ROC and Threshold



Search for the best threshold at the elbow of the ROC curve.

ROC - Explanation

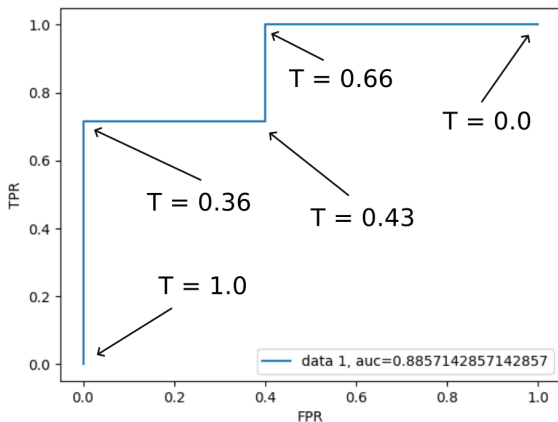


The larger the *area under the ROC curve (ROC-AUC)*, the better.

ROC-AUC ranges from 0 to 1. $\text{ROC-AUC} \approx 0.5$ indicates random guessing.

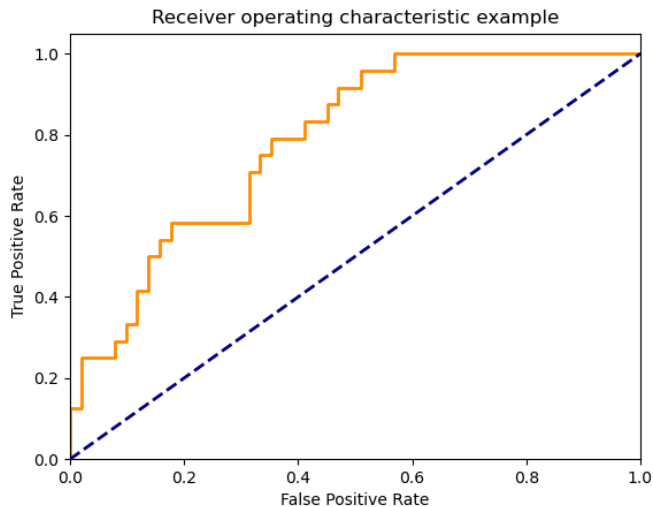
ROC-AUC

Index	1	2	3	4	5	6	7	8	9	10	11	12
Actual	1	1	1	1	1	0	0	1	1	0	0	0
Predicted	.98	.95	.9	.86	.66	.48	.43	.42	.36	.15	.1	.05



ROC-AUC = 0.8857

Iris - ROC-AUC



ROC-AUC = 0.79

ROC-AUC - Probabilistic Interpretation

How is the ROC-AUC connected with the samples?

ROC-AUC - Probabilistic Interpretation

How is the ROC-AUC connected with the samples?

Consider our cancer detection example:

Index	1	2	3	4	5	6	7	8	9	10	11	12
Actual	1	1	1	1	1	0	0	1	1	0	0	0
Predicted	.98	.95	.9	.86	.66	.48	.43	.42	.36	.15	.1	.05

ROC-AUC - Probabilistic Interpretation

How is the ROC-AUC connected with the samples?

Consider our cancer detection example:

Index	1	2	3	4	5	6	7	8	9	10	11	12
Actual	1	1	1	1	1	0	0	1	1	0	0	0
Predicted	.98	.95	.9	.86	.66	.48	.43	.42	.36	.15	.1	.05

AUC has a probabilistic explanation:

Consider the following experiment:

ROC-AUC - Probabilistic Interpretation

How is the ROC-AUC connected with the samples?

Consider our cancer detection example:

Index	1	2	3	4	5	6	7	8	9	10	11	12
Actual	1	1	1	1	1	0	0	1	1	0	0	0
Predicted	.98	.95	.9	.86	.66	.48	.43	.42	.36	.15	.1	.05

AUC has a probabilistic explanation:

Consider the following experiment:

- ▶ Choose randomly a patient i from positive patients
Each positive patient has the same probability of being chosen.

ROC-AUC - Probabilistic Interpretation

How is the ROC-AUC connected with the samples?

Consider our cancer detection example:

Index	1	2	3	4	5	6	7	8	9	10	11	12
Actual	1	1	1	1	1	0	0	1	1	0	0	0
Predicted	.98	.95	.9	.86	.66	.48	.43	.42	.36	.15	.1	.05

AUC has a probabilistic explanation:

Consider the following experiment:

- ▶ Choose randomly a patient i from positive patients
Each positive patient has the same probability of being chosen.
- ▶ Choose randomly a patient j from negative patients
Each negative patient has the same probability of being chosen.

ROC-AUC - Probabilistic Interpretation

How is the ROC-AUC connected with the samples?

Consider our cancer detection example:

Index	1	2	3	4	5	6	7	8	9	10	11	12
Actual	1	1	1	1	1	0	0	1	1	0	0	0
Predicted	.98	.95	.9	.86	.66	.48	.43	.42	.36	.15	.1	.05

AUC has a probabilistic explanation:

Consider the following experiment:

- ▶ Choose randomly a patient i from positive patients
Each positive patient has the same probability of being chosen.
- ▶ Choose randomly a patient j from negative patients
Each negative patient has the same probability of being chosen.
- ▶ Check if $h_i > h_j$.

ROC-AUC - Probabilistic Interpretation

How is the ROC-AUC connected with the samples?

Consider our cancer detection example:

Index	1	2	3	4	5	6	7	8	9	10	11	12
Actual	1	1	1	1	1	0	0	1	1	0	0	0
Predicted	.98	.95	.9	.86	.66	.48	.43	.42	.36	.15	.1	.05

AUC has a probabilistic explanation:

Consider the following experiment:

- ▶ Choose randomly a patient i from positive patients
Each positive patient has the same probability of being chosen.
- ▶ Choose randomly a patient j from negative patients
Each negative patient has the same probability of being chosen.
- ▶ Check if $h_i > h_j$.

The ROC-AUC is the probability of succeeding in the $h_i > h_j$ test.

Summary

We have discussed various metrics that can be used to evaluate the quality of a classifier.

The metrics summarize the results of evaluation on a given dataset.

Summary

We have discussed various metrics that can be used to evaluate the quality of a classifier.

The metrics summarize the results of evaluation on a given dataset.

We have discussed metrics for evaluating

- ▶ binary classifiers,
Accuracy, Precision, Recall, F_1 , and few more
- ▶ multi-class classifiers,
Accuracy, Precision, Recall, F_1
- ▶ probabilistic classifiers,
parametrized metrics, ROC-AUC

Summary

We have discussed various metrics that can be used to evaluate the quality of a classifier.

The metrics summarize the results of evaluation on a given dataset.

We have discussed metrics for evaluating

- ▶ binary classifiers,
Accuracy, Precision, Recall, F_1 , and few more
- ▶ multi-class classifiers,
Accuracy, Precision, Recall, F_1
- ▶ probabilistic classifiers,
parametrized metrics, ROC-AUC

There are still several questions unanswered:

- ▶ When to use the metrics.
- ▶ How to estimate the influence of sampling the dataset.

Use of Evaluation Metrics

In our case, the following scenarios are typical:

- ▶ **Final test:** Evaluate the model on the test set (separated at the beginning of training) and then compute the metrics. May inform the user about the quality of the model.

Use of Evaluation Metrics

In our case, the following scenarios are typical:

- ▶ **Final test:** Evaluate the model on the test set (separated at the beginning of training) and then compute the metrics. May inform the user about the quality of the model.
- ▶ **Validation:** Evaluate models on a separate validation set and use the metrics to compare models.

There are (at least) two scenarios in which this happens:

- ▶ Hyperparameter fine-tuning.
- ▶ Comparison of different models (e.g., KNN and decision trees).

Use of Evaluation Metrics

In our case, the following scenarios are typical:

- ▶ **Final test:** Evaluate the model on the test set (separated at the beginning of training) and then compute the metrics. May inform the user about the quality of the model.
- ▶ **Validation:** Evaluate models on a separate validation set and use the metrics to compare models.

There are (at least) two scenarios in which this happens:

- ▶ Hyperparameter fine-tuning.
- ▶ Comparison of different models (e.g., KNN and decision trees).

Keep in mind that the metrics are artificial, and the results of the model are roughly summarized.

It would be best if you always strived to test the proper functionality of your model in as natural conditions as possible.

For example, a model for medical diagnosis should be evaluated by medical doctors who may observe many features of its behavior that are difficult to express quantitatively.

How to Estimate Significance

Machine learning models are typically trained on (pseudo) random samples of data objects.

For example, a set of patients treated by the concrete hospital.

How to Estimate Significance

Machine learning models are typically trained on (pseudo) random samples of data objects.

For example, a set of patients treated by the concrete hospital.

However, the purpose of testing/evaluation is to get information about the whole population (i.e., all possible patients).

How to Estimate Significance

Machine learning models are typically trained on (pseudo) random samples of data objects.

For example, a set of patients treated by the concrete hospital.

However, the purpose of testing/evaluation is to get information about the whole population (i.e., all possible patients).

How do we estimate how much specific properties of the given sample influence our model?

This is a challenging question; methods of inferential statistics are needed to get the answer.

How to Estimate Significance

Machine learning models are typically trained on (pseudo) random samples of data objects.

For example, a set of patients treated by the concrete hospital.

However, the purpose of testing/evaluation is to get information about the whole population (i.e., all possible patients).

How do we estimate how much specific properties of the given sample influence our model?

This is a challenging question; methods of inferential statistics are needed to get the answer.

We will consider these issues in some later lecture. Concretely,

- ▶ *Bias-variance* tradeoff
- ▶ *Statistical tests* for testing
 - ▶ significance of the metrics values,
 - ▶ paired t-tests for comparing models.

How to Compare Classifiers

Let us consider two classifiers. How do you compare them?

How to Compare Classifiers

Let us consider two classifiers. How do you compare them?

Accuracies and F_1 scores can be compared easily (they are just numbers).

How to Compare Classifiers

Let us consider two classifiers. How do you compare them?

Accuracies and F_1 scores can be compared easily (they are just numbers).

How to compare $(\text{Precision}_1, \text{Recall}_1)$ of the first classifier with $(\text{Precision}_2, \text{Recall}_2)$ of the second classifier?

How to Compare Classifiers

Let us consider two classifiers. How do you compare them?

Accuracies and F_1 scores can be compared easily (they are just numbers).

How to compare $(\text{Precision}_1, \text{Recall}_1)$ of the first classifier with $(\text{Precision}_2, \text{Recall}_2)$ of the second classifier?

Thresholding

- ▶ Introduce a threshold $0 \leq t \leq 1$
- ▶ Demand, one of the two metrics (typically the Recall), to be at least t . That is

$$\text{Recall}_1 \geq t \quad \text{Recall}_2 \geq t$$

- ▶ Compare the values of the other metric numerically. In our case, decide whether

$$\text{Precision}_1 \geq \text{Precision}_2$$

(Still need to be concerned about the statistical significance.)

Example

Actual condition	Predicted condition	
	Canc.	Non-canc.
Cancer	6	2
Non-canc.	1	3
Total	8 + 4 = 12	

Actual condition	Predicted condition	
	Canc.	Non-canc.
Cancer	5	3
Non-canc.	0	4
Total	8 + 4 = 12	

$$\text{Precision}_1 = \frac{6}{7} \quad \text{Recall}_1 = \frac{6}{8}$$

$$\text{Precision}_2 = \frac{5}{5} = 1 \quad \text{Recall}_2 = \frac{5}{8}$$

Consider a threshold t on the Recall.

The second classifier is better if the threshold t is $5/8$.

If the threshold t is $6/8$, then the second classifier is unacceptable.