

3.36pt

Classification

Stefano Gualandi

Università di Pavia, Dipartimento di Matematica

email: stefano.gualandi@unipv.it
twitter: @famo2spaghi
blog: <http://stegua.github.com>
web: <http://matematica.unipv.it/gualandi/opt4ml>

Supervised Learning

Definition 1 (Supervised Learning)

Supervised Learning is the task of learning (inferring) a function f that maps input vectors to their corresponding target vectors, by using a dataset containing a given set of pairs of (*input*, *output*) samples. Examples:






- **REGRESSION**: the output vectors take one or more continuous values.
- **CLASSIFICATION**: the output vectors take one value of a finite number of discrete categories. Special case: binary classification.

Supervised Learning: Classification

Let us take $Y = \{-1, 1\}$, with $y = 1$ means **YES**, and $y = -1$ means **NO**.


Supervised Learning: Classification

Let us take $Y = \{-1, 1\}$, with $y = 1$ means **YES**, and $y = -1$ means **NO**.

$x_1 =$		$x_2 =$		$x_3 =$		$x_4 =$		$x_5 =$	
	$y_1 = 1$		$y_2 = 1$		$y_3 = -1$		$y_4 = 1$		$y_5 = -1$

Supervised Learning: Classification

Let us take $Y = \{-1, 1\}$, with $y = 1$ means **YES**, and $y = -1$ means **NO**.




① Use the TRAINING SET \Rightarrow run learning algorithm \Rightarrow get

$$f : X \rightarrow Y \quad \text{where } f := f_m \circ f_{m-1} \circ \dots \circ f_1$$

Supervised Learning: Classification

Let us take $Y = \{-1, 1\}$, with $y = 1$ means **YES**, and $y = -1$ means **NO**.



- ① Use the TRAINING SET \Rightarrow run learning algorithm \Rightarrow get

$$f : X \rightarrow Y \quad \text{where } f := f_m \circ f_{m-1} \circ \dots \circ f_1$$

- ② Use a TEST SET to validate f :

$$f(x_6 = \text{blueberry muffin}) = -1 \quad f(x_7 = \text{medium brown Chihuahua puppy}) = 1$$

Classification models

A **classification model**, or **classifier**, is used to label an example as belonging to one of a finite set of categories, called **classes**.

Classification models

A **classification model**, or **classifier**, is used to label an example as belonging to one of a finite set of categories, called **classes**.

One-class Learning: the training set contains examples drawn from only one class. Example: *anomaly detection*.

Classification models

A **classification model**, or **classifier**, is used to label an example as belonging to one of a finite set of categories, called **classes**.

One-class Learning: the training set contains examples drawn from only one class. Example: *anomaly detection*.

Two-class Learning (Binary Classification): the training set contains examples drawn from exactly two classes (positive and negative), and the objective is to find the boundary that separates the two classes.

Classification models

A **classification model**, or **classifier**, is used to label an example as belonging to one of a finite set of categories, called **classes**.

One-class Learning: the training set contains examples drawn from only one class. Example: *anomaly detection*.

Two-class Learning (Binary Classification): the training set contains examples drawn from exactly two classes (positive and negative), and the objective is to find the boundary that separates the two classes.


Multi-class Learning: involves finding the boundaries that separates more than two classes from each other.

Classification (basic) Methods

- 1 **k -Nearest Neighbours:** New examples are assigned a class based on how **similar**, using a **METRIC (!)**, they are to examples in the training data. The *learned* model is the training examples itself.



Classification (basic) Methods

- 1 **k -Nearest Neighbours:** New examples are assigned a class based on how **similar**, using a **METRIC (!)**, they are to examples in the training data. The *learned* model is the training examples itself.
- 2 **Linear Regression:** We can train a linear regression model for each class, and given a new example, assign it to the linear relationship that better predict the example.



Classification (basic) Methods

- 1 ***k*-Nearest Neighbours**: New examples are assigned a class based on how **similar**, using a **METRIC (!)**, they are to examples in the training data. The *learned* model is the training examples itself.
- 2 **Linear Regression**: We can train a linear regression model for each class, and given a new example, assign it to the linear relationship that better predict the example.



- 3 **Logistic Regression**: We can train a logistic regression model (see, Chap. 4.4 in Hastie et al.)

Evaluating Classifiers

Given a data set, we divide the data into:

Training Set: it is used to learn a model, and contains usually 80% of the data.

Evaluating Classifiers

Given a data set, we divide the data into:

Training Set: it is used to learn a model, and contains usually 80% of the data.

Test Set: it is used to evaluate the learned model, and it contains usually 20% of the data.

Evaluating Classifiers

Given a data set, we divide the data into:

Training Set: it is used to learn a model, and contains usually 80% of the data.

Test Set: it is used to evaluate the learned model, and it contains usually 20% of the data.

Given a data set, to increase the performance, it is better to fit the model by iteratively changing the Training and Test set, and by recording the model that yields the **BEST RESULTS** (see n -fold cross validation).

Evaluating Classifiers

Given a data set, we divide the data into:

Training Set: it is used to learn a model, and contains usually 80% of the data.

Test Set: it is used to evaluate the learned model, and it contains usually 20% of the data.

Given a data set, to increase the performance, it is better to fit the model by iteratively changing the Training and Test set, and by recording the model that yields the **BEST RESULTS** (see n -fold cross validation).


... but **BEST RESULTS** with respect to which METRIC?

Binary Classification: Confusion Matrix

		Condition (as determined by " Gold standard ")		
		Condition Positive	Condition Negative	
Test Outcome	Test Outcome Positive	True Positive	False Positive (Type I error)	Positive predictive value = $\frac{\Sigma \text{ True Positive}}{\Sigma \text{ Test Outcome Positive}}$
	Test Outcome Negative	False Negative (Type II error)	True Negative	Negative predictive value = $\frac{\Sigma \text{ True Negative}}{\Sigma \text{ Test Outcome Negative}}$
		Sensitivity = $\frac{\Sigma \text{ True Positive}}{\Sigma \text{ Condition Positive}}$	Specificity = $\frac{\Sigma \text{ True Negative}}{\Sigma \text{ Condition Negative}}$	

Binary Classification: Confusion Matrix

Type I error
(false positive)



Type II error
(false negative)



Accuracy

$$\text{ACCURACY} := \frac{\#TP + \#TN}{\#TP + \#TN + \#FP + \#FN}$$

Accuracy

$$\text{ACCURACY} := \frac{\#TP + \#TN}{\#TP + \#TN + \#FP + \#FN}$$

TP: True Positive

TN: True Negative

FP: False Positive

FN: False Negative

Accuracy

$$\text{ACCURACY} := \frac{\#TP + \#TN}{\#TP + \#TN + \#FP + \#FN}$$

TP: True Positive

TN: True Negative

FP: False Positive

FN: False Negative

QUESTION: When Accuracy make sense for evaluating a binary classifier? Example?

Evaluating Classifiers with Imbalanced Classes

		Condition (as determined by "class classifier")		
		Condition Positive	Condition Negative	
Test Outcome	Test Positive	True Positive (True Alarm)	False Positive (False Alarm)	Positive predictive value = # True Positive / # Test Outcome Positive
	Test Negative	False Negative (Miss)	True Negative	Negative predictive value = # True Negative / # Test Outcome Negative
		Sensitivity = # True Positive / # Condition Positive	Specificity = # True Negative / # Condition Negative	

$$\text{SENSITIVITY (RECALL)} := \frac{\#TP}{\#TP + \#FN}$$

$$\text{SPECIFICITY (PRECISION)} := \frac{\#TN}{\#TN + \#FP}$$

$$\text{POSITIVE PREDICTIVE VALUE} := \frac{\#TP}{\#TP + \#FP}$$

$$\text{NEGATIVE PREDICTIVE VALUE} := \frac{\#TN}{\#TN + \#FN}$$

Evaluating a Classifiers: Loss Functions

① Mean Square Error/Quadratic Loss/L2 Loss:

$$MSE := \frac{1}{n} \sum_{i=1}^n (y_i - f(x_i; w))^2$$

Evaluating a Classifiers: Loss Functions

① Mean Square Error/Quadratic Loss/L2 Loss:

$$MSE := \frac{1}{n} \sum_{i=1}^n (y_i - f(x_i; w))^2$$

② Mean Absolute Error/L1 Loss:

$$MAE := \frac{1}{n} \sum_{i=1}^n |y_i - f(x_i; w)|$$

Evaluating a Classifiers: Loss Functions

- ① Mean Square Error/Quadratic Loss/L2 Loss:

$$MSE := \frac{1}{n} \sum_{i=1}^n (y_i - f(x_i; w))^2$$

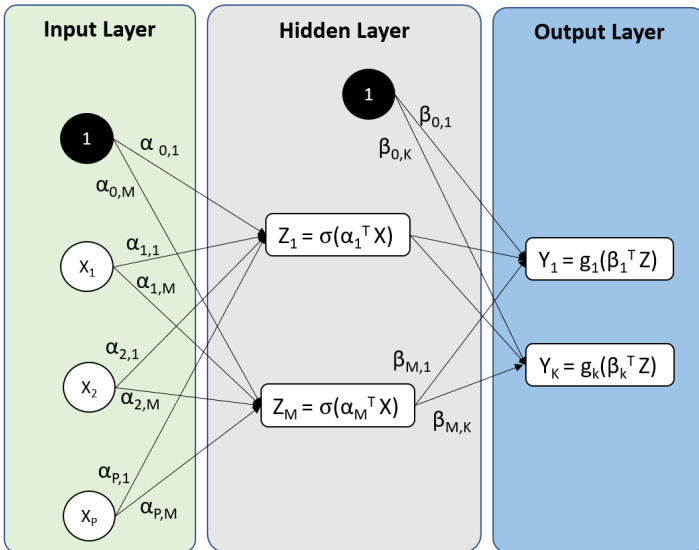
- ② Mean Absolute Error/L1 Loss:

$$MAE := \frac{1}{n} \sum_{i=1}^n |y_i - f(x_i; w)|$$

- ③ Cross Entropy Loss/Negative Log Likelihood:

$$CrossEntropy := \sum_{i=1}^n - (y_i \log(f(x_i; w)) + (1 - y_i) \log(1 - f(x_i; w)))$$

Classification with Multilayer Neural Networks



$$\alpha \in D^p \times M$$

$$\beta \in D^M \times K$$

Generative Adversarial Attacks

Original image




Dermatoscopic image of a benign melanocytic nevus, along with the diagnostic probability computed by a deep neural network.



+ 0.04 ×

Adversarial noise



Perturbation computed by a common adversarial attack technique. See (7) for details.

=

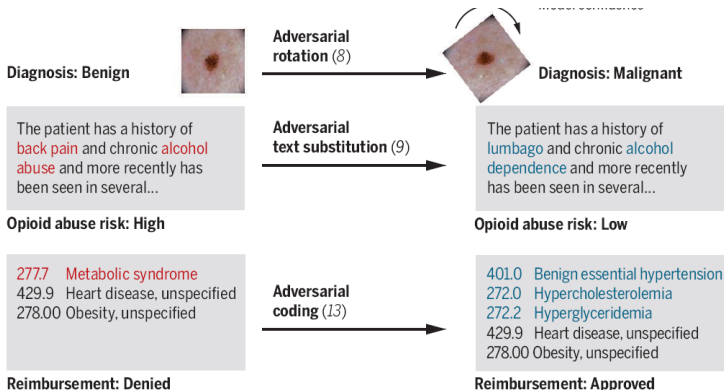
Adversarial example



Combined image of nevus and attack perturbation and the diagnostic probabilities from the same deep neural network.



Generative Adversarial Attacks



Convolutional NN

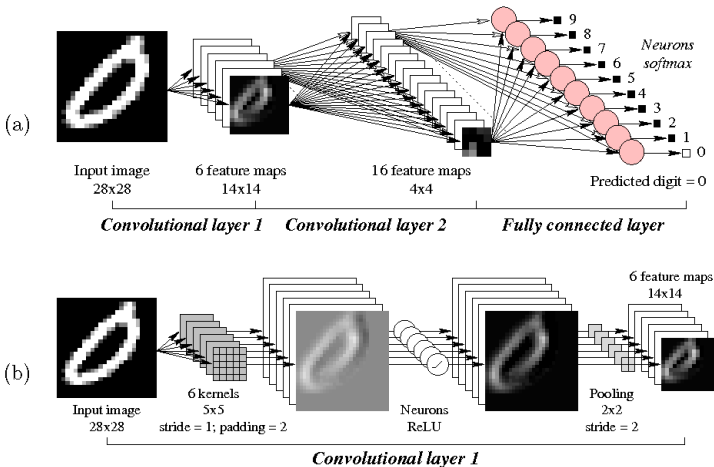


Figure 1: A convolutional neural network for the MNIST problem: global architecture (a) and detailed view of the first convolutional layer (b).

Recurrent NN




Figure 1: **Deep recurrent neural network prediction architecture.** The circles represent network layers, the solid lines represent weighted connections and the dashed lines represent predictions.

Elegant Machine Learning: Flux