

Degree programme	Type	Course
Computer Vision	OB	1

Contact lecturer

Name : Maria Vanrell Martorell

Email : maria.vanrell@uab.cat

Teaching staff

Ramon Baldrich Caselles

Fernando Luis Vilariño Freire

Dimosthenis Karatzas

Pau Rodriguez Lopez

Carlos Boned Riera

Guillem Arias Bedmar

Lluís Gomez Bigorda

Teaching staff (external to UAB)

Xavier Suau

Group languages

You can consult this information at the [end](#) of the document.

Prerequisites

- Degree in Engineering, Maths, Physics or similar.
- Programming Skills in Python.

Objectives

Course Coordinator: Dr. Ramon Baldrich Caselles

The objective of this course is to introduce the Machine learning techniques for solving computer vision problems. Machine learning deals with the automatic analysis of large scale data. Nowadays it conforms the basics of many computer vision methods, specially those related to visual pattern recognition or classification,

where 'patterns' encompasses images of world objects, scenes and video sequences of human actions, to name a few.

This course presents the foundations and most important techniques for the classification of visual patterns, mainly focusing on supervised methods. Also, related topics like image descriptors and dimensionality reduction are addressed. As much as possible, all these techniques are tried and assessed on a practical project concerning scene description from pictures, together with the standard metrics and procedures for performance evaluation like precision-recall curves and k-fold cross-validation.

The learning outcomes are:

- (a) Distinguish the main types of ML techniques for computer vision: supervised vs. unsupervised, generative vs. discriminative, original feature space vs. feature vector kernelization.
- (b) Know the strong and weak points of the different methods, in part learned while solving a real pattern classification problem.
- (c) Being able to use existing method implementations and build them from scratch.

The course goes in depth in two main approaches to introduce ML into the image classification problem. Using: a) handcrafted image description, b) data driven image description. On the first case the Bag of Words is used, on the second one, the Deep Learning approach. The DL content is developed extensively providing both, theoretical basis of the different parts of modern Neural Networks architectures, and best practices to apply it on real applications.

Learning outcomes

- CA06 (Achieve the objectives of a project of vision carried out in a team.) Achieve the objectives of a project of vision carried out in a team.
- KA03 (Identify the computational learning methods that can be used based on the data to solve a problem of vision.) Identify the computational learning methods that can be used based on the data to solve a problem of vision.
- KA10 (Select the best experimentation procedures to be designed for computational learning from training to evaluation.) Select the best experimentation procedures to be designed for computational learning from training to evaluation.
- KA16 (Recognise the ethical, gender and environmental dimensions of systems of vision and their application.) Recognise the ethical, gender and environmental dimensions of systems of vision and their application.
- SA03 (Apply and evaluate computational learning techniques to solve a specific problem.) Apply and evaluate computational learning techniques to solve a specific problem.
- SA13 (Calculate the carbon footprint for any experiment that requires training a deep neural network.) Calculate the carbon footprint for any experiment that requires training a deep neural network.
- SA14 (Detect bias in learning data sets which allow the construction of systems that are socially discriminatory to be avoided.) Detect bias in learning data sets which allow the construction of systems that are socially discriminatory to be avoided.
- SA17 (Prepare oral presentations that allow debate of the results of a project of vision.) Prepare oral presentations that allow debate of the results of a project of vision.

Contents

1. Introduction to machine learning
2. Experimental Setup
3. Embeddings: SVM and Random Forest
4. Introduction to Neural Networks
5. Introduction to Deep Learning
6. Convolutional Neural Networks
7. Training: data pre-processing, initialization, gradient optimization
8. Image Classification
9. Understanding and visualizing CNNs
10. Efficient methods for Deep Learning

Learning activities and methodology

Title	Hours	ECTS	Learning outcomes
Homework	113	4.52	CA06, KA03, KA10, KA16, SA03, SA13, SA14, SA17
Project follow-up sessions	8	0.32	CA06, KA03, KA10, KA16, SA03, SA13, SA14, SA17
Lecture sessions	20	0.8	CA06, KA03, KA10, KA16, SA03, SA13, SA14, SA17

Supervised sessions: *(Some of these sessions could be Synchronous on-line sessions)*

- Lecture Sessions, where the lecturers will explain general contents about the topics. Some of them will be used to solve the problems.

Directed sessions:

- Project Sessions, where the problems and goals of the projects will be presented and discussed, students will interact with the project coordinator about problems and ideas on solving the project (approx. 1 hour/week)
- Presentation Session, where the students give an oral presentation about how they have solved the project and a demo of the results.
- Exam Session, where the students are evaluated individually. Knowledge achievements and problem-solving skills

Autonomous work:

- Student will autonomously study and work with the materials derived from the lectures.
- Student will work in groups to solve the problems of the projects with deliverables:
 - Code
 - Reports
 - Oral presentations

Annotation: within the schedule set by the centre or degree programme, 15 minutes of one class will be reserved for students to evaluate their lecturers and their courses or modules through questionnaires.

Assessment

Continuous assessment activities

Title	Weight	Hours	ECTS	Learning outcomes
Session attendance	0.05	0.5	0.02	CA06, KA03, KA10, KA16, SA03, SA13, SA14, SA17
Project	0.55	6	0.24	CA06, KA03, KA10, KA16, SA03, SA13, SA14, SA17
Exam	0.4	2.5	0.1	KA03, KA10, KA16, SA03

The final marks for this module will be computed with the following formula:

$$\text{Final Mark} = 0.4 \times \text{Exam} + 0.55 \times \text{Project} + 0.05 \times \text{Attendance}$$

where,

- Exam: is the mark obtained in the Module Exam (must be ≥ 3).
- Attendance: is the mark derived from the control of attendance at lectures (minimum 70%)
- Projects: is the mark provided by the project coordinator based on the weekly follow-up of the project and deliverables (must be ≥ 5). All accordingly with specific criteria such as:
 - Participation in discussion sessions and in team work (inter-member evaluations)
 - Delivery of mandatory and optional exercises.
 - Code development (style, comments, etc.)
 - Report (justification of the decisions in your project development)
 - Presentation (Talk and demonstrations on your project)

Only those students that fail (Final Mark < 5.0) can do a retake exam.

Bibliography

Journal papers:

1. Barber, D. "*Bayesian Reasoning and Machine Learning*". Cambridge University Press, 2012.
2. Yoshua Bengio. "*Learning Deep Architectures for AI*". Foundations and Trends in Machine Learning, Vol. 2, No. 1, 2009.
3. Christopher J. C. Burges. "*Dimension Reduction: A Guided Tour*". Foundations and Trends in Machine Learning, Vol. 2, No. 4, 2009.
4. Christoph H. Lampert. "*Kernel Methods in Computer Vision*". Foundations and Trends in Computer Graphics and Vision, Vol. 4, No. 3, 2008.
5. Tinne Tuytelaars and Krystian Mikolajczyk. "*Local Invariant Feature Detectors: A Survey*". Foundations and Trends in Computer Graphics and Vision, Vol. 3, No. 3, 2007.

Books:

1. Ian Goodfellow, Yoshua Bengio and Aaron Courville. "*Deep Learning*". 2016. Cambridge, MA, USA: The MIT Press. ISBN: 978-0262035613
2. Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar, "*Foundations of Machine Learning*" MIT Press, 2012. <http://www.cs.nyu.edu/~mohri/mlbook/>
3. Z.H. Zhou. "*Ensemble Methods: Foundations and Algorithms*". Chapman & Hall/CRC, 2012.

Reports:

1. Criminisi, A. and Shotton, J. and Konukoglu, E. "*Decision Forests for Classification, Regression, Density Estimation, Manifold Learning and Semi-Supervised Learning*". Technical report MSR-TR-2011-114. Microsoft Research, 2011.
http://research.microsoft.com/pubs/155552/decisionForests_MSR_TR_2011_114.pdf

Software

Tools for Python programming with special attention to most used Computer Vision and Deep Learning libraries

Course groups and languages

The information provided is provisional until November 30. After this date, you will be able to consult the language of each group through this [link](#). To access the information, you will need to enter the course CODE

Type of teaching	Group	Language	Semester	Shift
(TEm) Theory (master)	1	English	first semester	afternoon
(PLABm) Practical laboratories (master)	1	English	first semester	afternoon
(PLABm) Practical laboratories (master)	2	English	first semester	afternoon