Syllabus
Learning Feature Representations, 6hp

Issued by the WASP graduate school management group *draft*

Main field of study

AI/MLX

Course level

PhD student course

Course offered for

PhD Students in the WASP graduate school

Entry requirements

The participants are assumed to have a background in mathematics corresponding to the contents of the WASP-course "Mathematics and Machine Learning".

Module 1 and 3: Knowledge of calculus, linear algebra and especially probability theory is very helpful. Basic understanding of machine learning is preferred. Programming skills in any language.

Module 2: Knowledge about advanced linear algebra, basics in machine learning, signal processing, and image analysis are required. Programming skills in Python+Numpy.

Intended learning outcomes

Module 1. Understand and implement several unsupervised approaches to feature learning including score matching, noise-contrastive estimation, latent variable models such as (deep) Boltzmann machines, variational auto-encoders, deep sparse coding and deep energy models.

Module 2. Be able to use concepts from computer vision learning such as generative and discriminative models, invariance and equivariance, and open-world problems in the design of algorithms. Implement state-of-the-art algorithms for visual object tracking.

Module 3. Recognize and explain many useful relations in 3D geometry and projective geometry and understand how they can be incorporated in deep neural networks.

Course content

Module 1. One of the main goals of representation learning is to learn how to extract generic features valid for a range of tasks from given data. In this module we focus on energy-based models for representation learning, which are closely linked to unsupervised learning. We will discuss restricted Boltzmann machine and higher-order variants. One particular emphasis is on how to estimate the parameters of energy-based models, since straightforward maximum likelihood estimation is not suitable for these models. We will show how proper scoring rules (specifically score matching and noise-contrastive estimation) can be used to estimate the parameters of energy-based models, and how they are connected with auto-encoders.

Module 2. Visual representations can be categorized into generative and discriminative models, depending on whether they are supposed to represent visual appearance explicitly or implicitly. An explicit representation is typically an image patch of a part of a feature map from a deep network. Implicit representations are dual to image patches or feature maps, in the sense that they are optimal for a discriminative task, such as localization, detection, or classification. In particular, we will look into the problem of visual object tracking, starting from classical least-squares approaches (Lucas-Kanade), continuing with online-learned discriminative correlation filters, and eventually ending with deep-feature-fusion-based state-of-the-art algorithms and the transition to video object segmentation.

Module 3. 3D geometry and projective geometry are essential aspects of real world perception for autonomous systems. In this module we will review results from projective geometry, such as plane-to-plane correspondence, epipolar, and oriented epipolar geometry, absolute pose estimation and more. We will put particular emphasis on how distances and errors are best defined, given geometry and probability theory. This is an important consideration when integrating geometric estimation in deep neural networks, and we will also look at how geometric optimization layers can be defined. We will also look at practical implications of the introduced theory for situations such as: learning to estimate absolute pose, and learning to perceive depth and 3D structure from video.

Teaching and working methods

Module 1. Two full-day lectures to cover terminology, theory, and algorithms. Group project to implement studied algorithms.

Module 2. Four lectures and two seminars to cover terminology, theory, and algorithms. A group project to cover practical skills for implementing state-of-the-art algorithms.

Module 3. Six lectures and two seminars to cover terminology, theory, and algorithms.

Examination

Module 1. Group project on learning deep energy with report

Module 2. Active participation in the two seminars, Handing in of preparation tasks on the seminar papers, Project on DCF tracking with report

Module 3. Active participation in the seminars. Preparatory questions on the seminar papers. Lecture attendance.

Grades

Fail or Pass