Syllabus
Learning Feature Representations, 6hp

Issued by the WASP graduate school management group June 2026

Main field of study

AI/MLX

Course level

PhD student course

- AS track: elective
- AI track: elective
- Joint Curriculum: advanced

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 “Foundational Mathematics for Machine Learning”.

Module 1 and 3: Knowledge of calculus, linear algebra, and probability theory is very helpful. Basic understanding of machine learning. 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. Students will gain a principled understanding of how unsupervised and self-supervised methods learn compressed, informative and transferrable representations from data. They will be able to compare modern approaches and assess trade-offs in their theoretical foundations and practical implementation.

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 representing visual information.

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. Self-Supervised Representation Learning, Fredrik Lindsten
A central goal of machine learning is to learn how to extract meaningful, generic features from data that capture underlying structure and remain useful across a range of downstream tasks. In this module, we explore a variety of principled approaches to learning such representations in an unsupervised or self-supervised setting. We will examine how different model families and training objectives shape the representations that emerge, and what theoretical and practical trade-offs they involve. The emphasis will be on modern self-supervised approaches as well as the estimation and optimization techniques that make these methods tractable.

Module 2. Learning of discriminative models, Michael Felsberg
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 video analysis: object tracking and segmentation. We will consider various techniques such as correlation filter and vision transformers.

Module 3. Representations and 3D Geometry, Per-Erik Forssén
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. Lectures and seminars that cover terminology, theory, and algorithms.

Module 2. Lectures and seminars that cover terminology, theory, and algorithms. A group project to cover practical skills for implementing state-of-the-art algorithms.
Module 3. Lectures and seminars that cover terminology, theory, and algorithms.

In order to reduce travels, the course is arranged in two blocks that are to be done on site in Linköping. The first block introduces theory, and preparations for the second block, with respect to all three modules. The second block covers more advanced theory, and seminars covering the preparation tasks. At the end of the second block, the project hand-in task for module 2 is introduced. We will also describe the option to do an extended project (for 2hp extra credits) in connection with one of the WARAs.

Examination

Module 1. Active participation in the seminars. Preparatory questions on the seminar papers. Lecture attendance.

Module 2. Active participation in the seminars, handing-in of preparation tasks on the seminar papers, project with report.

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

Two extra hp can also be obtained by handing in an extended project report in January after the end of the course.

Grades

Fail or Pass