Syllabus
Deep learning and GANs, 6hp

Issued by the WASP graduate school management group 2021 08 30.

Main field of study

AI/mlx

Course level

PhD student course

Course offered for

PhD Students in the WASP graduate school

Entry requirements

Basic machine learning, linear algebra, probability theory, basic optimization, programming (Python). The participants are assumed to have a background in mathematics corresponding to the contents of the WASP-course "Mathematics and Machine Learning".

Intended learning outcomes

After the course, students should be able to:

- explain the fundamental principles of supervised and unsupervised learning, including basic techniques like cross-validation to avoid overfitting,
- describe the standard cost functions optimized during supervised training and the standard solution techniques.
- explain how traditional feed-forward networks are constructed and why they can approximate “almost” and function (the universality theorem).
- summarize the key components in convolutional neural networks (CNNs) and their key advantages.
- argue for the benefits of transfer learning and data augmentation in situations when we have a limited amount of annotated/labelled data.
- train and apply CNNs to image applications.
- understand deep neural networks on a high level and know how to train them.
- understand generalization and why it is needed.
- understand why stochastic gradient descent has implicit regularization properties that help improve generalization in deep overparameterized neural networks.
- account for the theoretical background for probabilistic and generative deep learning techniques.
- implement methods based on recently published results for probabilistic or generative deep networks.

- understand the dynamics that govern common algorithms for training of deep neural networks
- reflect on different implicit regularization effects for different training algorithms

Course content

Module 1: Fundamentals of deep learning. Feedforward networks. Training procedures. Convolutional neural networks.

Module 2: Implicit regularization of stochastic gradient descent in overparameterized deep neural networks.

Module 3: Uncertainty Estimation. Out-of-Distribution Detection/Robustness. Deep Generative Modeling.

Teaching and working methods

Three modules. Each module includes preparatory assignments (including video lectures and quizzes), two days of intense teaching on site, as well as exercises and projects.

Examination

One hand-in assignments per module.

Grades

Fail or Pass