

Understanding and Building Large Language Models, 6.0 credits

Understanding and Building Large Language Models, 6.0 hp

Third-cycle education course

6FIDA19

Department of Computer and Information Science

Valid from: First half-year 2026

Approved by
Head of Department

Approved
2026-02-03

Registration number

Entry requirements

Students are expected to have enough CS and AI background to be able to follow technical descriptions of deep learning methods. Basic understanding of deep learning is expected. Students are also expected to be able to implement deep learning solutions using Python and PyTorch.

Learning outcomes

After completing the course, the student should be able to:

- * Explain the technical underpinnings of large language models.
- * Explain the processes involved in training a large language model.
- * Implement and train a basic large language model from scratch in PyTorch.
- * Read and comprehend recent, academic papers on LLMs and have knowledge of the common terms used in them (alignment, scaling laws, RLHF, prompt engineering, instruction tuning, etc.).
- * Understand and discuss concepts and terminology of state-of-the-art LLMs.
- * Develop an ability to distinguish fact from fantasy in this fast-moving field.

Contents

The aim of the course is to explain the methods and techniques used by large generative AI models such as large language models work and explore how to build them. The focus is on the technical aspects such as methods and techniques. The course is thus more about machine learning than about natural language processing.

Educational methods

The course consists of 10 lectures where the course material is presented. The students are expected to complete lab assignments on their own, either individually or in pairs.

Lectures are expected to be given Monday afternoons 15-17 in a hybrid mode, either on site in Linköping or online through Zoom. The lectures will be recorded and made available afterwards.

Examination

Develop an LLM from scratch and conduct at least two experiments related to the lectures and write a short report on it. The lab is expected to have four parts:

1. Develop a simple data pre-processing pipeline
2. Pre-train a GPT-style LLM
3. Fine-tune the LLM
4. Evaluate the LLM

Grading

Two-grade scale

General information

Interest registration: Emilie Lind: forskarladok@ida.liu.se.