



Laboratory Course Artificial Intelligence - Deep Learning Lab (English)

Examiner: Prof. Dr. Steffen Staab · Supervision and contact: Daniel Frank, Lluís Pastor, Tim Schneider (Analytic Computing, Institute for Artificial Intelligence)

Deep Learning Lab gives students hands-on research experience. In teams of up to three, students take an open research question and work through the full cycle: framing it, reviewing literature, choosing a methodology, implementing it, and evaluating it (including ablations and error analysis). Topics change yearly, drawn from the Analytic Computing group's current work in scientific machine learning and generative AI.

Structure. Staff present topics in weeks 1–2; students then form teams and pick a project. For the first six weeks, each student also prepares and presents a notebook on an assigned topic outside their project, building shared foundations. Class time splits evenly between notebook work and the group project; most project work happens outside class.

Prerequisites:

- Solid ML theory (e.g., Prof. Staab's Machine Learning course or equivalent)
- Strong Python skills and experience with a deep learning framework (e.g., PyTorch)
- Comfort with linear algebra, calculus, and probability
- Ability to read papers and work independently on open-ended problems

Format & grading. A semester-long, independent group project (teams of ≤ 3), assessed on: documented/reproducible code, a short research-paper-style report, and a final presentation. The focus is scientific rigor and clear communication, not polish.

Who it's for. Ideal for students curious about real research who thrive on open-ended problems. It's demanding and self-directed — not an ML intro — so those lacking a strong ML background or willingness to grapple with unsolved problems may find it frustrating. Strong projects can lead to theses or further work with the group, a good stepping stone toward a research-oriented Master's or PhD.

