Changes to the course in 2026

Marcel Bollmann

This document informs about changes to the course *Language Technology* compared to the previous course instance. The course is co-taught as `729G86` and `TDP030`.

1 Previous course evaluation

The course iteration in 2025 received an average rating (i.e. answer to the question *"What is your overall evaluation of the course?"*) of 4.17 for `729G86` and 3.17 for `TDP030`. 18 out of 91 students responded to the course evaluation.

Many students had positive feedback about the course content and the lectures. A common criticism was that the instructions for the labs were not clear enough, which we tried to address for 2026 *(see below)*.

2 Basis for the changes

- **The field of language technology has seen a period of rapid development**, from the emergence of powerful neural encoder models (starting with the initial release of BERT in 2018) to the ongoing trend of more and more powerful large language models (LLMs). These technologies have had a big impact on the wider society (e.g. through the emergence of AI assistants), but also on language technology as a research field, as they have changed what kind of tasks and models are considered relevant. **These developments have not been adequately reflected in the course contents so far**.
- The free-text comments in previous course evaluations expressed that **the labs should have clearer instructions**. This sentiment was echoed by several students who responded to the evaluation. Other criticisms that we have occasionally seen in course evaluations throughout the years include that there is **not enough time** for the labs and that **advanced labs are too difficult/time-consuming**.

3 Changes being carried out

The course syllabi have been updated to **add an additional learning goal** (ILO 5), which states that students should be able to *"reason about the capabilities and limitations of large language models for solving language technology problems"*. This new ILO 5 makes explicit that LLMs are now a core part of the course content.

The **course contents have been entirely reworked**. In particular:

- The **course topics** have been adjusted to include more content related to modern neural network approaches, such as masked language models and generative large language models.
- The **number of lectures** has been increased from five to eight (not counting the introduction week), while the **number of labs** has been reduced by one. Previously, the course content was structured into five topics with one lecture and one lab assignment each; now, the course is structured into four topics with two lectures and one lab assignment each. The motivation for these changes is as follows:
  - Lectures were very content-dense and didn't have much time for questions or exam practice; in 2026, lectures will include more time for questions, discussions, and going through exercises similar to those that will be on the exam.
  - Students in previous years occasionally felt that they didn't have enough time for the labs; while the submission deadlines for the labs remain the same as before, we hope that reducing the total number of labs by one will help with the overall workload, as well as compensate for the extra number of scheduled lectures.
- The **content of the labs** has been reworked to be in line with the new content of the lectures. Particular attention has been given to improving the clarity of instructions within the lab notebooks. The advanced labs have been made a bit easier.
- The **group project** instructions have been changed to require the use of an LLM as part of the project, in line with the new ILO 5. We now suggest a "default project" that is specifically tailored to the updated learning goals, but students can still formulate an individual project like before, as that was a part of the course that was regularly praised by students. Since the default project already gives students more guidance than before, the separate "project introduction" session was removed.