

Introduction to the project

Marcel Bollmann

Department of Computer and Information Science (IDA)

This introduction session

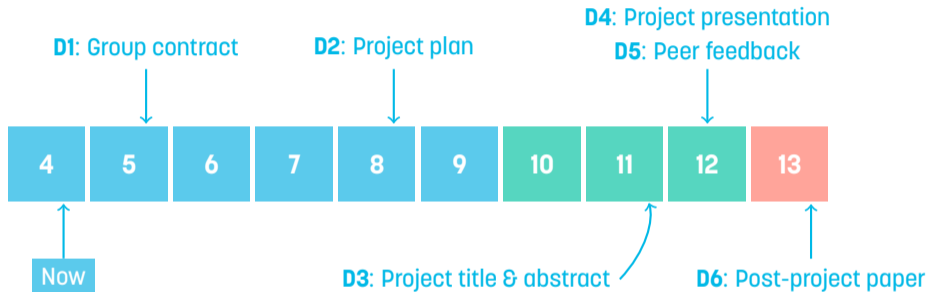
1. Logistics
2. Practical tips
3. Examples

Project logistics



The project's main purpose
is to let you assess
the difficulty and feasibility
of language technology applications.

Project timeline



Let's look at the [🔗 project assignments instructions](#) on the course website!

Practical tips



A common project structure

- 1 Choose an **application** of language technology.
- 2 Choose one or more **concrete systems** that implement this.
- 3 **Evaluate** the system(s) w.r.t. relevant properties.
- 4 Draw **conclusions** about the difficulty and feasibility.

Path 1: Start with the application

- **Pick an application** that you find interesting and want to know more about.
 - text classification
 - text generation, chatbots
 - question answering
 - ...
- Look into **software and datasets** that are available for that application.
- Find out how those systems are typically **evaluated**.

Path 2: Start with the data

- **Pick a dataset** that you find interesting and want to know more about.
 - [Kaggle](#), [HuggingFace Datasets](#), [Riksdagens öppna data](#), ...
 - Web scraping to gather your own dataset
- Look at the data and develop ideas **what you could do** with it.
 - Also consider: What have *others* done with it?

Implementation work

- **Coding** can be part of a project, but it is **not the main focus**.
 - Remember the learning objective!
- Don't start from scratch; **use existing libraries!**
 - [spaCy](#), [scikit-learn](#), [NLTK](#), [Gensim](#), [HuggingFace Transformers](#), ...
 - No requirement on the programming language!
- Anything related to language technology is fine; you **can go beyond** the course content!

Evaluate your approach

- **Intrinsic evaluation** using well-defined evaluation measures
 - *e.g.*, accuracy, precision, recall, perplexity, ...
- **Extrinsic evaluation** by embedding the application into a larger system
 - could involve doing a user study
- **Subjective evaluation** by manually analyzing the results
 - *e.g.* How well do the results fit known facts? How well do they fit a theory?

How to get help

- Pitch your project idea to me!
- I am offering **project group meetings** for you to get feedback.
 - Book a time slot through the link on the website.
- I am also available **via e-mail**.

Example projects from previous years



Book genre prediction

- Can we **predict a book's genre** (e.g. *horror, fantasy, romance*) from a summary of its plot?
 - Text classification problem
- Compare two different models & analyze the most predictive words per genre.
 - “kill” is predictive of horror, “power” is predictive of fantasy
- Also perform a small user study to see how difficult humans would find the task.

Cross-domain sentiment analysis using hotel reviews

- Can we **train a sentiment classification model** on one domain of texts, and later apply it to a different domain?
 - Sentiment analysis: predicting whether a text is e.g. “positive”, “negative”, or “neutral” towards something.
- Train on hotel reviews, test on product reviews.
 - Data from Tripadvisor & Amazon
- Compare different models, analyze the effect of review length.
 - SVM, BERT, VADER

Quantifying text emotiveness

- Can we detect **how emotionally engaged** someone was when producing a text?
- Consider psycholinguistic theories about how emotiveness can be measured in text.
 - Trager coefficient, aggressiveness coefficient, readiness to action
- Analyze the emotiveness of speeches over time, explain the results.

Lyrics generation

- How difficult is it to automatically **generate song lyrics**?
- Compare two advanced deep learning models.
 - LSTM model trained on song lyrics & ChatGPT
- Challenging to evaluate: conduct a survey with human participants.
 - “How much do you like the lyrics? How realistic do they sound?”

Change of language in Reddit over time

- Can we study **how language changed over time** in specific subreddits?
- Approximate the “change of language” by looking at different features.
 - sentiment analysis
 - average post and sentence length
 - verb/noun ratio
- Relate the findings to events that happened in the real world, and that may be discussed in these subreddits.

More project abstracts
from previous years
are on the course website!

