Teaching Natural Language Processing in Law School

Daniel Braun
University of Twente
Department of High-tech Business
and Entrepreneurship
d.braun@utwente.nl

Abstract

Fuelled by technical advances, the interest in Natural Language Processing in the legal domain has rapidly increased over the last months and years. The design, usage, and testing of domain-specific systems, but also assessing these systems from a legal perspective, needs competencies at the intersection of law and Natural Language Processing. While the demand for such competencies is high among students, only a few law schools, particularly in Europe, teach such competencies. In this paper, we present the design for a Natural Language Processing course for postgraduate law students that is based on the principle of constructive alignment and has proven to be successful over the last three years.

1 Introduction

Like most fields of occupation, the legal profession is undergoing a drastic change due to the introduction of Artificial Intelligence (AI) and particularly Natural Language Processing (NLP) in the workplace. While research on applications of AI in the legal domain has a long history going back to the 1970s (Rissland et al., 2003), the legal practice was, for a long and arguably more so than other fields, hesitant to adopt AI technology and the pace of adoption started to increase only very recently (Oskamp and Lauritsen, 2002; Araszkiewicz et al., 2022). While there have been many good reasons for this hesitance, including strict regulations and personal liability of legal professionals (Vladika et al., 2024), both the acceptance and use of AI technology are rapidly increasing within the legal profession in recent years (Weinstein, 2022).

This increase in usage also results in an increased need for legal professionals with technological expertise, both in traditional roles but also in specialised emerging roles like legal engineer or legal technologist. According to the "Future Ready Lawyer Report" by Wolters Kluwer, only 24% of lawyers say that they understand "transformational technologies" like AI and big data (Weinstein, 2022). These numbers highlight that there is a significant demand for teaching technological essentials to (future) legal professionals. As language is the most important tool of lawyers and other legal professionals, among AI technologies, NLP technologies are arguably most relevant for them with a number of practical applications including contract generation, document management, legal decision-making, anonymization, and many other tasks (Vladika et al., 2024). Additionally, competencies in both NLP and law are not only needed for the application of NLP to legal problems but also for the application of law to NLP. With new regulations like the European Union's AI Act, there will be an increased need for lawyers with a good understanding of AI technology who are able to counsel clients who develop, distribute, or use such systems.

While many US law schools have already included NLP in their curriculum (Johnson, 2023), in Europe, such offerings are still rather rare. Partially, this is caused by the fact that law degrees in most European countries are undergraduate degrees that often follow a highly regulated four to five year curriculum, unlike in the US, where law schools are part of postgraduate education. However, as consecutive Master of Laws (LL.M.) degrees are becoming more popular in Europe, law schools have started to offer such degrees particularly focused on Legal Technology ("Legal Tech"). This paper outlines the design of an introductory NLP course that is designed for this setting, i.e. as an elective course in postgraduate education for law students without prior knowledge of AI or computer science. The course has been taught in a very similar fashion for three years at an LL.M. program at a public German university. The aim of this design is to inform and inspire the development of similar courses.

Proceedings of the Sixth Workshop on Teaching NLP, pages 85–90

2 Skills and Intended Learning Outcomes

Following the principle of constructive alignment (Biggs, 1996), the course was designed around the intended learning outcomes (ILOs) and more fundamentally the skills that are relevant for legal professionals. Given the context for which the course was designed, no relevant previous knowledge in AI or computer science can be expected. Therefore, an NLP course in a law school can clearly not go as deep into technological details as a traditional NLP course. After all, a law graduate with a specialisation in legal technology is, most likely, also not going to work as a programmer. Therefore, the most important question in designing an NLP course for law students is: *What are the relevant skills that can help future legal professionals to work at the intersection of law and NLP?*

It is worth emphasising that the course design outlined in this paper is not meant to be part of the general curriculum every law student has to go through, but rather as an elective for those that want to work at the intersection of law and NLP where a deepened interdisciplinary understanding is needed. We identified three different areas in which such interdisciplinary skills are mainly needed:

- **Law practice**: NLP tools will be part of the daily work routine of many legal professionals in the future. Understanding the fundamental principles these tools operate on will help legal professionals to make informed decisions on how and when to use which technology or tool.
- **Technology development and implementation**: Legal professionals will be involved in the development and implementation of legal tech tools in many different roles in which they will need a solid understanding of the underlying technology, like product manager or implementation consultant.
- **Research**: Whether as a legal data scientist or empirical legal scholar, legal research, both in academia and practice, will include many different NLP methods. Especially in interdisciplinary research projects, annotating legal corpora and preparing them to be used as training data will require legal expertise and an understanding of the requirements for training data for ML models.

Based on these areas, we derived ILOs that law graduates would need to achieve to be prepared for such roles. The six derived ILOs are shown in Table 1. Each ILO is classified according to Bloom's taxonomy (Anderson and Krathwohl, 2001). Since the goal of the course is to provide participants with practical NLP skills that are relevant in the legal domain, most of the ILOs are on the higher levels of Bloom's taxonomy (create, evaluate, analyse, apply), i.e. go beyond acquiring just theoretical knowledge.

3 Course Plan

Based on the ILOs, we developed a plan for a course that runs over the span of a semester (∼15 weeks) and has a workload of 140 hours (or 5 points in the European Credit Transfer and Accumulation System, ECTS). Of those 140 hours, 56 are contact hours: 30 hours (or 2 hours per week) of lectures and 26 hours of tutorials (13 x 2 hours). The remaining 84 hours are reserved for self-study and project work (see Section 4.1). Particularly for the achievement of the ILOs that are related to the higher levels of Bloom's taxonomy, practical experience, as mediated through tutorials and project work, is key to successful learning. Table 2 shows an overview of all educational activities and their content. While the lectures mainly focus on theoretical knowledge, the tutorials are designed as hands-on sessions, in which participants interact with different tools. Some of the course materials, particularly for the practical sessions, are available on GitHub[1]. It is important to stress that the designed course does not attempt to replace a computer science curriculum. Therefore, the lectures mainly stay on a conceptual level, without, e.g., going into the mathematical details or optimisation algorithms in ML. Similarly, in the tutorials, participants will not write complete programs but rather are provided with Jupyter notebooks that they can configure in a low-code fashion, so that they get a basic understanding of the process and structure that underlies the training of ML models, without necessarily having the learn how to program. Weeks 12 to 14 are reserved for guest lectures from practitioners. The idea is to invite representatives of different organisations that have a need for professionals with skills at the intersection of law and NLP to show the broad range of potential roles, not only within traditional law firms but

[1] https://github.com/DaBr01/Teaching-NLP-in-Law-School

Table 1: Intended Learning Outcomes (ILOs) and their classification according to Bloom's taxonomy

#	ILO	Level
1	Students can explain terminology, methods, theories, and fundamental concepts of artificial intelligence, machine learning, and natural language processing.	understand
2	Students can explain the relevance of artificial intelligence and natural language processing for the legal domain and discuss their implications on different legal professions.	understand
3	Students can use standard software for the automation of process steps in legal knowledge work (e.g. no-code platforms and document generation tools).	apply
4	Students can analyse problem formulations from the legal practice and identify suitable NLP methods for support and automation, as well as potential risks that may arise from them (like biases and lack of accountability or explainability).	analyse, evaluate
5	Students can design annotation guidelines for legal data sets and generate new corpora through annotation.	create
6	Students can develop small legal applications for specific problem formulations based on standard tools and libraries.	create

also within established companies or organisations and start-ups.

4 Assessment

The assessment of the course consists of an exam at the end of the course (that can be either written or orally), and a project on which students work in groups of two to three. Both assessment forms will determine 50% of the final grade each. While the exam is purely a formative assessment, the project contains summative and formative assessment moments. The goal of formative assessment is to assess the level of knowledge a student has at a certain point in time. In the context of the course, formative assessment is used to assess, at the end of the course, whether the ILOs have been achieved and to derive a grade. Formative assessment, on the other hand, is part of the teaching process and aims to provide the students, but also the teacher, with information about the current progress of the students, in order to allow for interventions and assure that, in the end, the ILOs will be achieved. (Garrison and Ehringhaus, 2007)

4.1 Project

For the project, each team will be provided with a corpus and a problem description. Each team will have to solve the following tasks based on their corpus and problem description:

1. Develop an annotation guideline.
2. Annotate the corpus based on the developed guideline.
3. Train a model that is able to extract information that is relevant in the context of the given problem description.
4. Write a project report of max. 4 pages that describes the results of your project.
5. Present the project results to your peers.

The summative assessment element of the project will be the final project grade that is based on the annotation guidelines (30%), the trained model (30%), the project report (30%), and the presentation (10%). In addition, participants hand in their annotation guidelines in week 4 for a formative assessment in which they receive feedback but also give feedback to other groups as part of a peer feedback process.

5 Literature

While there are many textbooks available on NLP (e.g. by Jurafsky and Martin (2008) or Hapke et al. (2019)), almost all of them are much more technical than is suited for the outlined course. The textbooks by Biemann et al. (2022) and Ignatow and Mihalcea (2017) present an introduction to NLP particularly targeted at the social sciences. As such, they include introductions to fundamental concepts that are also suitable for the context of the course but do not address any particularities of the processing of legal language or domain-specific tasks and tools. Particularly because of this lack of literature, we hope that the course design presented in this paper can inform and inspire the design of NLP courses in the context of law schools.

Table 2: Course Plan

	Type	Title	Content
Week 1	Lecture	Introduction to AI and ML	Definition and history of AI; Fundamentals of ML; Decision Trees; k-nearest Neighbours; Linear Regression; k-means Clustering; Critical reflection on correlation and causality
	Tutorial	No-code platforms and legal expert systems	Creation of a simple legal expert system with a no-code platform (e.g. Bryter, Open Decision, or Mioto)
Week 2	Lecture	Text annotation	Fundamentals of legal data annotation; Binary / multi-label / multi-class annotations; Document / section / sentence / word / sequence annotation; Disagreement between annotators; Annotation guidelines; Annotation tools; Quality of annotations
	Tutorial	Annotation of court decisions	Introduction to an annotation tools (e.g. Doccano); Writing annotation guidelines; Annotation of a small corpus of court decisions
Week 3	Lecture	Introduction to Natural Language Processing	NLP definitions and tasks; Challenges of automated text processing (compared to images or numbers); Specifics of legal texts; Pipeline architectures; Pre-processing; Content extraction; Sentence segmentation; Stemming and lemmatization; Lexica
	Tutorial	Pre-Processing	Implementation of a pre-processing pipeline for privacy policies based on a Jupyter notebook template
Week 4	Lecture	Information Extraction	Named entity recognition; Regular expressions; Citation networks; Information extraction; Relation extraction; Argument mining; Evaluation (precision / recall / F1)
	Tutorial	Information Extraction	Regular expression exercise; Implementation of NER and IE for the annotated corpus of week 2 based on a Juypter notebook template
	Project	Deadline	Annotation guidelines
Week 5	Lecture	Text classification and Topic Modelling	Text classification; Topic Modelling; Contract analysis; Natural Language Understanding services
	Tutorial	Contract Analysis	Analysis of a contract corpus with three different means (contract analysis tool (e.g. Summize), NLU cloud service (e.g. Azure AI), ChatGPT) and comparison of results
Week 6	Lecture	Language Models and Vector Representations	Vector Representations; Distributional Hypothesis; Word Embeddings; Neural Networks; Large Language Models; Fine-Tuning
	Tutorial	Vector Representations	Comparison of IE performance with different vector representations (bag of words, tf-idf, word embeddings) as input with a provided Jupyter notebook template; Visualisation and interpretation of Word Embeddings
	Project	Deadline	Peer-feedback annotation guidelines
Week 7	Lecture	Similarity Analysis	Jaccard coefficients; Word Mover's Distance; Cosine similarity; Document vectors; Detection of AI-generated texts
	Tutorial	Similarity Analysis	Comparison of different similarity metrics on the court decision corpus

	Type	Title	Content
Week 8	Lecture	Introduction to Natural Language Generation	Data2Text / Image2Text / Text2Text; Templates; Rule-based NLG; Stochastic NLG; Hallucination; Abstractive and extractive summarisation; Evaluation and metrics (BLEU, ROUGE)
	Tutorial	Text Generation	Rule-based text generation with SimpleNLG; Text generation with GPT-2; Exercise on hallucination
Week 9	Lecture	Document Processing and Automation	eDiscovery; Document management, Legal document generation tools; Process automation
	Tutorial	Document generation	Document generation with Word; Document generation with no-code platform from week 1; Document generation with specialised tool
Week 10	Lecture	Legal Case Outcome Prediction	State of the research in legal case outcome prediction; Relevant features (particularly non-legal features like time, ...); Bias; Reflection about the difference between predicting the court decision and predicting the "right" decision
	Tutorial	Legal Case Outcome Prediction	Training of a simple model for the prediction of legal case outcomes based on a provided Jupyter notebook template
	Project	Deadline	Model and Jupyter notebook
Week 11	Lecture	Explainable AI	Explainability; Interpretability; Adversarial attacks; Taxonomy of explainability (post hoc/ante hoc, global/local, model specific/model agnostics); Example-based approaches; Critical reflection of "explanations" generated by LLMs
	Tutorial	LIME	Introduction of the LIME library for the generation of explanations
Week 12	Lecture	Guest lecture established organisation	Show possible job profiles in established organisations, e.g. legal operations officer from large companies, publishing houses, software providers, NGOs, government
	Tutorial	Project Support	Before the final project submission, participants have time to ask questions and get support
Week 13	Lecture	Guest lecture startup	Show possible job profiles in the startup world (e.g. legal tech startups, fintech, ...)
	Tutorial	Startup software	Introduction to the product of the startup
	Project	Deadline	Report
Week 14	Lecture	Guest lecture law firm	Job profiles in classical law forms
Week 15	Lecture	Project Presentation	Teams present the results of their project work

References

Lorin W Anderson and David R Krathwohl. 2001. *A taxonomy for learning, teaching, and assessing: A revision of Bloom's taxonomy of educational objectives: complete edition*. Addison Wesley Longman, Inc.

Michał Araszkiewicz, Trevor Bench-Capon, Enrico Francesconi, Marc Lauritsen, and Antonino Rotolo. 2022. Thirty years of artificial intelligence and law: overviews. *Artificial Intelligence and Law*, 30(4):593–610.

Chris Biemann, Gerhard Heyer, and Uwe Quasthoff. 2022. *Wissensrohstoff Text*. Springer.

John Biggs. 1996. Enhancing teaching through constructive alignment. *Higher education*, 32(3):347–364.

Catherine Garrison and Michael Ehringhaus. 2007. Formative and summative assessments in the classroom.

Hannes Max Hapke, Hobson Lane, and Cole Howard. 2019. Natural language processing in action.

Gabe Ignatow and Rada Mihalcea. 2017. *An introduc-*

tion to text mining: Research design, data collection, and analysis. Sage Publications.

Julia Johnson. 2023. Innovative law schools: Artificial intelligence. Last accessed 2024-05-06.

Daniel Jurafsky and James H Martin. 2008. *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition*. Pearson.

Anja Oskamp and Marc Lauritsen. 2002. Ai in law practice? so far, not much. *Artificial Intelligence and Law*, 10(4):227–236.

Edwina L. Rissland, Kevin D. Ashley, and R.P. Loui. 2003. Ai and law: A fruitful synergy. *Artificial Intelligence*, 150(1):1–15. AI and Law.

Juraj Vladika, Stephen Meisenbacher, Martina Preis, Alexandra Klymenko, and Florian Matthes. 2024. Towards a structured overview of use cases for natural language processing in the legal domain: A german perspective. *arXiv preprint arXiv:2404.18759*.

Stuart Weinstein. 2022. *Lawyers' Perceptions on the Use of AI*, pages 413–432. T.M.C. Asser Press, The Hague.