

Human-Centered Computing in Legal NLP

An Application to Refugee Status Determination

Claire Barale

The University of Edinburgh
School of Informatics
Edinburgh, Scotland
claire.barale@ed.ac.uk

Abstract

This paper proposes an approach to the design of an ethical human-AI reasoning support system for decision makers in refugee law. In the context of refugee status determination, practitioners mostly rely on text data. We therefore investigate human-AI cooperation in legal natural language processing. Specifically, we want to determine which design methods can be transposed to legal text analytics. Although little work has been done so far on human-centered design methods applicable to the legal domain, we assume that introducing iterative cooperation and user engagement in the design process is (1) a method to reduce technical limitations of an NLP system and (2) that it will help design more ethical and effective applications by taking users' preferences and feedback into account. The proposed methodology is based on three main design steps: cognitive process formalization in models understandable by both humans and computers, speculative design of prototypes, and semi-directed interviews with a sample of potential users.

1 Scope of the research proposal

At the core of the global refugee crisis is the legal procedure of Refugee Status Determination (RSD), i.e. the decision of granting refugee status or not. Refugee adjudication is a high-stakes, life-altering decision that impacts vulnerable people. Our project aims at helping and supporting all parties involved in refugee status adjudications to make better decisions by using data-driven intelligence. It looks at building an ethical human-AI decision support system and focuses on augmenting human legal reasoning through the use of machine learning models. The aim is neither to output a decision nor to recommend one, as we think refugee status determination should ultimately be made by human experts.

Potential users of the system are stakeholders in the legal decision process such as a lawyer, counsel, judge, civil servant, or case worker. Although not the direct users, asylum-seekers are essential interested parties as they should directly benefit from improvements in the procedure.

Text data in refugee law includes cases and decisions, country reports, international conventions and local refugee status regulations. Our work is based on a data set containing the text of first instance decisions rendered in Canada over the past 25 years (approx. 20,000 decisions). Given the importance of text and language, its interpretations and levels of meaning in law, we want to explore the application of state-of-the-art natural language processing (NLP) methods to extract and organize information from past decisions.

We hypothesize that human-centered computing (HCC), design and human-computer interaction (HCI) methods can be exploited in legal NLP systems to enhance trust and overall performance by providing easier access to information and reducing risks associated to the use of AI in the legal field. Trust in our system is not immediate for users and we will need to provide rational guarantees and good evidence of safety, understood as effective avoidance of risks and harms. Precisely, we assume that trust can be warranted by modeling features of interpersonal trust, by ensuring usefulness of the system and its functionalities and demonstrating its benefits. As a starting point, we assess potential risks and describe them as well as potential unwanted events or consequences.

While there is little specific literature on human-centered computing and human-computer interaction in law, we build on general HCC and HCI literature for high-stakes decision making. Given the above stated hypothesis, this document aims at exploring relevant methodologies and design

processes that can support the conception of our system.

2 Legal NLP background

Legal AI focuses on building AI-powered tools for the legal domain. Much of it specifically relies on NLP methods to help accomplish legal tasks (Zhong et al., 2020; Dale, 2019; Branting et al., 2018). Here, common functionalities include information retrieval (Undavia et al., 2018), database management (Refworld), similar case matching (Morris, 2019; Trappey et al., 2020; Undavia et al., 2018), legal prediction (Katz et al., 2017; Chen and Eagel, 2017; Medvedeva et al., 2020), text summarization, legal advice, contract and document automation and review. Work on legal design also looks at legal procedure and systems with the aim of developing user-centered methodologies and designs approaches (Hagan, 2020). Although it does not necessarily imply the use of AI systems, automation and text analysis is a major field of investigation and LegalTech has recently received a lot of attention.

As it is arguably difficult for a machine learning-powered system to capture qualitative data, any textual representation that can be processed with NLP and text analytics tools will be partial and subject to errors. Text analytics is limited when it comes to capturing meaning, context and legal arguments, and is only able to try and generalize knowledge based on past decisions and historical data that were contained in the training data set (Ashley, 2017).

3 Risks and obstacles

This section identifies some risks that we anticipate to arise from this project. Risks associated with the design of our system are both technical constraints and ethical considerations, especially in terms of impact on the users. We specifically assessed how the design of our system could negatively impact individuals whether legal practitioners or claimants. We will link this approach of ethics and impact assessment with human-centered computing and try to combine human and AI learning and reasoning.

A literature review and preliminary research has highlighted the following risks and limitations. The first risk concerns asylum seekers needs throughout their application process: risk of unjustified decision as to the determination of their status, risk to refuse refugee status to someone who would be

granted the status had our tool not been used, lack of support and information and risk that the application process becomes more painful for the asylum seeker. Other potential risks include: narrow AI in law and need for manual engineering, combining human and machine legal reasoning, accuracy bias, fairness and interpretability, accountability, privacy concerns, impact of the use of AI on the legal process and the law.

From this assessment, we chose to gather risks in four categories that represent clear requirements to work on the design of the system. Since each one of these concerns user requirements, it is worth noting that different users may have different requirements for each one of these risks and that design should facilitate tradeoffs. We conclude that the main challenges to design our system will be to guarantee *trust, usefulness, usability*, and provide *benefits* for refugees.

4 Human-AI cooperation

Human-centered computing is commonly defined as the use of computing technologies centered on human experiences (Amershi et al., 2019; Shneiderman, 2020).

Human-AI cooperation is the proposed way to mitigate the risks listed above by combining benefits from AI systems such as computational power with human abilities including intuition and context-aware reasoning. Based on our review of the literature, we find that human-AI interaction may provide an interesting way to try and mitigate the uncertainty of legal procedure while also addressing some limitations of AI algorithms, which will hopefully lead to higher acceptance from legal practitioners. Users indeed need guarantees to use the system, which would require several qualities such as transparency and justification, but also improved user experience and design.

It is assumed that involving the user through interaction and cooperation with the application naturally generates more trust. This approach is also called “mutualism” (Siddarth et al., 2021), cognitive computing (Ashley, 2017; Zatarain, 2018), interactive machine learning (Dudley and Kristensson, 2018) and has the advantage of reducing the need for comparison or even competition between humans and AI, lack of accountability, and to mitigate the problem of control over an AI application.

This method typically involves trade-offs, leading us to think in terms of the balance of cooper-

ation between an AI application and its user. The question is to find what methods can be used to translate this theoretical approach in terms of design of the model, functionalities and user interface.

5 Effective human-centered design

This section aims at analyzing functionalities of our system in sight of ways of working, procedures and design approaches across the three domains involved in our research: NLP, refugee law, and human-centered design. Table 1 is not meant to be exhaustive and displays a preliminary analysis. It aims at determining shared features between domains that are conflicting and will require further attention when building our system. The table is based on principles of human-centered research and “legal design” – defined as the convergence of legal theory and frameworks and HCI approaches (Hagan, 2020). From this, we expect to be able to better translate principles into users’ specific needs. We want to make sure that benefits toward asylum seekers are at the core of the methodology.

Table 1 highlights a number of key issues:

1. HCC and HCI rely upon adaptation of a system to its users for a positive outcome and are experimental while legal procedure and frameworks are fixed and not flexible by principle. Specifically, refugee law is rule-based and outlines precise categories of reasons for which refugee status can be granted. Legal compliance is of course an important requirement of the system that will have to be prioritized.
2. The second conflicting point is uncertainty, as we know that NLP-based methods will not reach 100% accuracy, especially given the sparse data available in refugee law. On the other hand, we don’t want legal procedure and decision to reflect any uncertainty. For instance, while summaries of applications can reduce the work load, they should be very carefully reviewed so that no important element of a case is missed. For this reason, we also need to include other evaluation criteria besides accuracy-based ones.
3. Understanding NLP functionalities relies on a technical understanding, which may prove difficult in practice and limit the integration of such functionalities into legal procedures and reasoning. In the same way that legal

reasoning should be explainable and able to justify decisions, our system should be able to give clear reasons as to its approach and outputs.

4. HCC aims at involving all stakeholders and their specific requirements, when legal procedure is restricted to specific individuals directly involved in the procedure. Therefore it is worth noting that different stakeholders may have different requirements.
5. Since we want to capture human legal intuition and thinking accurately, we want to design the functionalities of the systems based on the process of legal reasoning as it is practiced by human beings. For instance, similarity analysis reproduces legal reasoning by analogy and precedent.

6 Proposed design methodology

The general idea that underpins our approach is that we aim to develop algorithms that not only learn from data, but also through exposure to human practices and interactions with human experts. To achieve this, we will employ methods of participatory design, value-sensitive design and rapid prototyping.

Building on table 1, we propose the following methodology to translate the mapping into a human-centered design process. The methodology is summarized visually in figure 1.

6.1 Step 1: Understanding and formalizing cognitive processes

As our research looks at “augmenting” legal human reasoning by using NLP tools, it would first require breaking down the human decision-making process into machine understandable steps. A main difficulty will likely be to divide a human reasoning into logical steps, to link elements between them, and, ideally, to identify inference steps and causal links, which are of course not always apparent in human thought. This is true when designing our application, but also in the users’ understanding of the application outcome and in explaining the steps followed by the system.

We want to make sure that our design reflects the cognitive process of the legal decision-maker for two reasons. First, because the closer our system will reflect human cognitive processes, the better it

NLP system functionalities and methods	Legal decision making and procedure	HCC-HCI methods and design
Information retrieval and text analysis: keywords analysis and argument mining whether based on a query by document or by question typed by the user	Based on legal frameworks (international convention on refugee status (UNHCR, 1951)) and country reports, using legal data bases (Refworld)	Participatory design, value-sensitive design and iterative process by successive prototypes
Similarity analysis: retrieving similar past cases	Decisions rendered by text leading to a positive or negative outcome decided by a country jurisdiction (facts, application of the legal framework and procedure explanation)	Use of systems is experience dependent and guided by users' intuition
Text summarizing: summarizing a case with some relevant predefined features and summary of the facts	Facts and refugee story gathered by interviews (conducted by civil servants) and hearings	Importance of user interface and visualization of the data, process and outputs of the system
Accuracy and performance of the model	Legal expertise (lawyers, counsel, judges)	Cognitive process and intuition in using a legal AI tool
Feature analysis and comparison with country reports (factual) information	Procedures and procedural fairness	Support function of the design in guiding changes in legal procedures and ways of deciding
Data and model possible biases	Cognitive biases, impact of non-legal and non-factual parameters	Design biases

Table 1: Mapping for design guidelines and effective cooperation

will be understood and intuitive to use both in terms of functionalities and interface usability. Second, because it will help dividing tasks into machine-understandable processes for which we can design effective algorithms.

6.2 Step 2: Prototyping

Work on this project will proceed iteratively in designing and testing a series of prototypes design. Each prototype will be followed by an evaluation step as described in section 7 below. Our first prototype will propose various functionalities relying on legal text analytics, as described in the first column of table 1.

6.3 Step 3: Understanding users' preferences and requirements

We will present each prototype and results obtained with it to selected legal professionals (refugee lawyers, counsels, judges). We hope to get feedback on the system from its potential users as well as from legal scholars. The core of this work will be to understand users' requirements, their views on the use of AI in the target domain, the poten-

tial usage they can envision for machine learning systems, and to investigate their levels of trust and acceptability toward AI in the context of refugee law. We specifically expect to test the usefulness of the proposed functionalities of the system in terms of benefits for the decision-making process. We also want to observe and test the usability of the interface.

To this end, we will meet with a sample of legal professionals involved in international law (about 10 interviewees). We expect to recruit both judges and lawyers or case worker submitting the applications. This will require developing guidelines for meeting topics so that we can effectively compare answers across stakeholder engagement activities, which will take the form of semi-directed interviews and workshops. Relevant questions to ask would be for instance: what are precise users' requirements, what functionalities are the most helpful, what is the tasks that takes the longest and can cause delay in processing a claim.

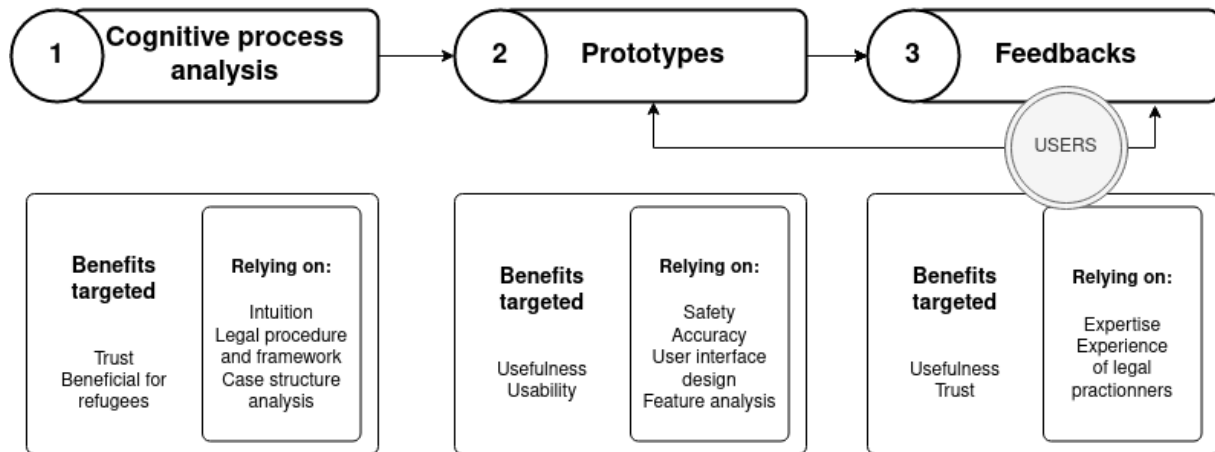


Figure 1: Workflow

7 Evaluation methods

To evaluate our system, we plan to use both quantitative metrics and qualitative analysis and to expand the scope of our evaluation beyond accuracy-based measures. Metrics should be three-fold:

- From NLP, we will evaluate accuracy, quality and performance of the model.
- From HCI, we will evaluate users' speed of comprehension, positive user experience, ease of use of the proposed interface
- From the legal point of view, we will evaluate legal accuracy (accordance to procedures, frameworks and laws), legal relevance of highlighted information, administrative burden (Hagan, 2020), and relevance of propose functionalities.

As our system aims at benefiting refugees, we want to add an additional evaluation metric in the form of "design for dignity" (Almohamed and Vyas, 2016) that accounts for the beneficial use of AI and its positive inputs toward a specifically vulnerable population as refugees.

8 Conclusion and future work

This document highlights some solutions and methods for designing an NLP-powered decision support system aiming at providing additional insight to the refugee status determination process. It should be treated as a starting point towards exploring how NLP tools could be beneficial to asylum seekers and help understand reasons and steps leading to a decision outcome. In the future, we plan to

test empirically this methodology and implement the above listed functionalities.

References

- Asam Almohamed and Dhaval Vyas. 2016. Designing for the marginalized: A step towards understanding the lives of refugees and asylum seekers. In *Proceedings of the 2016 acm conference companion publication on designing interactive systems*, pages 165–168.
- Saleema Amershi, Dan Weld, Mihaela Vorvoreanu, Adam Fourney, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, Paul N. Bennett, Kori Inkpen, Jaime Teevan, Ruth Kikin-Gil, and Eric Horvitz. 2019. *Guidelines for Human-AI Interaction*. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pages 1–13, Glasgow Scotland Uk. ACM.
- Kevin D Ashley. 2017. *Artificial intelligence and legal analytics: new tools for law practice in the digital age*. Cambridge University Press.
- L. Karl Branting, Alexander Yeh, Brandy Weiss, Elizabeth Merkhofer, and Bradford Brown. 2018. *Inducing Predictive Models for Decision Support in Administrative Adjudication*. *AI Approaches to the Complexity of Legal Systems*, 10791:465–477. Series Title: Lecture Notes in Computer Science.
- Daniel L. Chen and Jess Eagel. 2017. *Can machine learning help predict the outcome of asylum adjudications?* In *Proceedings of the 16th edition of the International Conference on Artificial Intelligence and Law*, pages 237–240, London United Kingdom. ACM.
- Robert Dale. 2019. *Law and Word Order: NLP in Legal Tech*. *Natural Language Engineering*, 25(1):211–217.

- John J. Dudley and Per Ola Kristensson. 2018. [A Review of User Interface Design for Interactive Machine Learning](#). *ACM Transactions on Interactive Intelligent Systems*, 8(2):1–37.
- Margaret Hagan. 2020. [Legal Design as a Thing: A Theory of Change and a Set of Methods to Craft a Human-Centered Legal System](#). *Design Issues*, 36(3):3–15.
- Daniel Martin Katz, Michael J. Bommarito, and Josh Blackman. 2017. [A general approach for predicting the behavior of the Supreme Court of the United States](#). *Plos one*, 12(4):e0174698.
- Masha Medvedeva, Michel Vols, and Martijn Wieling. 2020. [Using machine learning to predict decisions of the european court of human rights](#). *Artificial Intelligence and Law*, 28(2):237–266.
- Jason Morris. 2019. [User-Friendly Open-Source Case-Based Legal Reasoning](#). In *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Law*, pages 270–271, Montreal QC Canada. ACM.
- UNHCR Refworld. [Refworld | Country Reports](#).
- Ben Shneiderman. 2020. [Human-centered artificial intelligence: Reliable, safe & trustworthy](#). *International Journal of Human-Computer Interaction*, 36(6):495–504.
- Divya Siddarth, Daron Acemoglu, Danielle Allen, Kate Crawford, James Evans, Michael Jordan, and E. Glen Weyl. 2021. [How AI fails us](#).
- Charles V. Trappey, Amy J.C. Trappey, and Bo-Hung Liu. 2020. [Identify trademark legal case precedents - Using machine learning to enable semantic analysis of judgments](#). *World Patent Information*, 62:101980.
- Samir Undavia, Adam Meyers, and John Ortega. 2018. [A Comparative Study of Classifying Legal Documents with Neural Networks](#). pages 515–522.
- UNHCR. 1951. [Convention and Protocol Relating to the Status of Refugees](#).
- Jesus Manuel Niebla Zatarain. 2018. [Artificial Intelligence and Legal Analytics: New Tools for Law Practice in the Digital Age](#). *SCRIPT-ed*, 15(1):156–161.
- Haoxi Zhong, Chaojun Xiao, Cunchao Tu, Tianyang Zhang, Zhiyuan Liu, and Maosong Sun. 2020. [How does NLP benefit legal system: A summary of legal artificial intelligence](#). *CoRR*, abs/2004.12158.