Legal Science and Computer Science: A Preliminary Discussion on How to Represent the "Penumbra" Cone with AI

Angela Condello[1], Giorgio Maria Di Nunzio[2]

[1] Department of Law, University of Messina, Italy
[2] Department of Information Engineering, University of Padua, Italy
[1]Piazza Pugliatti, 1 - 98122 Messina - Italy
[2]Via Gradenigo, 6/a - 35131 Padova - Italy
angela.condello@unime.it, giorgiomaria.dinunzio@unipd.it

Abstract

Legal science encounters significant challenges with the widespread integration of AI software across various legal operations. The distinction between signs, senses, and references from a linguistic point of view, as drawn by Gottlob Frege at the end of the 19th century, underscores the complexity of legal language, especially in multilingual contexts like the European Union. In this paper, we describe the problems of legal terminology, examining the "penumbra" problem through Herbert Hart's legal theory of meaning. We also analyze the feasibility of training automatic systems to handle conflicts between different interpretations of legal norms, particularly in multilingual legal systems. By examining the transformative impact of Artificial Intelligence on traditional legal practices, this research contributes to the theoretical discussion about the exploration of innovative methodologies for simplifying complex terminologies without compromising meaning.

Keywords: Legal terminology, Linguistic Sign, Terminology

1. Background

In this paper, we explore the multifaceted challenges facing legal science considering the widespread adoption of AI software across various legal operations, such as verification, drafting, risk analysis, and prediction,[1] with specific reference to the potential confusion between signs, senses, and their reference. Such distinction, as it is widely known, was drawn by Gottlob Frege in a renowned paper published in 1892 (Frege, 1892). Frege thereby defines it in such a kind that to the sign there corresponds a definite sense and to that - in turn - a definite reference, while to a given reference (an object) there does not belong only one single sign. From this perspective, the same sense (e.g., equality) can have different expressions in different languages and realms, and even in the same language.

Within our interdisciplinary context, between legal philosophy and computer engineering, we aim at narrowing down onto a pivotal issue: the evolving dynamics of legal terminology due to the pervasive and ever-increasing use of AI software by legal professionals, including lawyers, judges, and notaries (Rissland et al., 2003), with specific reference to the potential confusion between signs, senses and references caused by such use of AI software for legal professionals. How can we prevent the blurring of this fundamental differentiation in philosophy of language, a differentiation that is extremely delicate in legal science?

Central to our investigation is Herbert Hart's theoretical framework (Hart and Leslie, 2012), which posits that legal concepts, mediated through the terms that indicate them, exhibit a dual nature. While, in theory, they possess a core of settled meaning, they are also surrounded by a "penumbra" of debatable cases, known as "hard cases" (Dworkin, 1975), wherein the application of words is neither evidently applicable nor categorically ruled out. As (Rissland et al., 2003) explain, legal rules derive their dynamic nature in part through the dynamic, open-textured nature of the terms used in the rules. As new situations arise, interpretation of the meaning of these terms changes as well. Such background is complexified in realms like the European Union, where translation of legal concepts is per se a very problematic issue both for lawyers and linguists.

Against this background, the integration of AI software in legal practice raises critical questions that we want to explore; in particular, whether it is conceivable to anticipate the potential emergence of hard cases and, subsequently, prepare legal software to navigate the intricate core-penumbra problem inherent in legal meaning. In addition, the increasing sophistication of these technologies and their availability have generated two divergent narratives about their potential implications, as described by (Whalen, 2022). These narratives alternately express excitement

---

[1] https://joinup.ec.europa.eu/collection/justice-law-and-security/solution/leos-open-source-software-editing-legislation/discussion/smart-leos-which-new-functionalities-should-be-implemented-next-and-what-can-be-learnt-corrigenda

about legal technology's potential to make the law more efficient and improve access to justice, or concern about the ways in which it may exacerbate existing biases or otherwise systematically harm justice.

Our research extends beyond the theoretical questions and addresses practical considerations tied to the intersection of AI and legal semantics. In this context, one main issue arises: can automatic systems be trained to foresee the contours of hard cases and adapt to the subtle distinctions of legal meaning? Can we measure the uncertainty of legal concepts and argumentation to handle the conflicts between different interpretations of norms (da Costa Pereira et al. 2017)? Can we foresee, by working with interdisciplinary methodology, the potential confusion between words and the concepts they refer to, especially in multilingual legal systems? Could it be too risky to use AI systems also in multilingual realms?

This challenge involves understanding how AI systems can effectively discern the relevance of contextual complications as well as societal changes, a task that is very important in the context of Hart's legal theory of meaning.

As we examine the implications of AI software on legal terminology, our analysis recognizes the transformative impact on traditional legal practices. The diffusion of AI technologies introduces a paradigm shift, necessitating a reevaluation of established legal methodologies. We explore the potential repercussions of this shift on the interpretation of legal documents and the inherent stability (or instability) of legal concepts. The balance between settled meanings and the penumbra of hard cases becomes increasingly important in a field where AI contributes to legal decision-making processes. By examining the core-penumbra problem through the lens of Hart's legal theory of meaning, we shed light on the challenges and opportunities posed by the integration of AI in legal science. Through this interdisciplinary analysis, we contribute to the ongoing discourse on the evolving nature of legal semantics in an era marked by the influence of augmented intelligence.

2. State of the Art: Penumbra and Simplification of Legal Texts

In this section, we present an overview of the state-of-the-art of the recent research papers that have been dealing with the issue of penumbra and text simplification in legal texts. It is important to highlight also a recent comprehensive systematic review in legal natural language processing (Quevedo et al., 2023) that complements this overview from an NLP point of view. In our analysis, we have searched Google Scholar to create a sort of systematic review of the topic with two queries: 1) "legal text" and "penumbra", 2) "legal text" and "text simplification"; then, we have kept only research papers that were published in the last two years in conferences or journals related to computer science/engineering or to interdisciplinary fields. We also kept the most recent articles that have been made available through the arXiv platform.

In (Stathis et al., 2024), the authors introduce Intelligent Contracts (iContracts), a new field blending AI and law, facing challenges like data quality. The focus is on Proactive Control Data (PCD) to enhance iContracts, a novel area in research. The extent to which Proactive Data impacts an Intelligent Contract depends on their quantitative identification and qualitative assessment, as well as the use of relevant technologies that integrate the risk assessment data when a contract is generated. Results include successful PCD generation, significant impact on contract drafting, and methods for assessing PCD quality.

The work presented by (Jiang et al., 2024) discusses leveraging large language models (LLMs) to enhance legal education for non-experts by employing storytelling. Since Law is, by nature, a sensitive domain, and computational tools must be designed responsibly, it is critical to design comprehension tools in ways that do not oversimplify or overgeneralize the nuances of legal jargon. In this context, the authors introduce a new dataset called LEGALSTORIES, comprising complex legal doctrines explained through stories and multiple-choice questions generated by LLMs. The idea is that storytelling aids in relating legal concepts to personal experiences and exhibits higher retention rates among non-native speakers.

In (Engel and McAdams, 2024), the authors test Large Language Models (LLM), ChatGPT in particular, to generate evidence on the ordinary meaning of statutory terms taking into account that many terms qualify as penumbral, and the legislative context often has some influence. The authors emphasize the importance of considering a distribution of replies rather than solely relying on the "best" reply identified by ChatGPT. Using Chat 3.5, the setting of these experiments defines prompts and refine them given contextual factors and historical periods. These experiments represent the first attempt to use GPT for empirical data on statutory term meanings, indicating potential for improving legal interpretation despite the need for caution.

The study presented in (Dixit et al., 2024) explores the effectiveness of extractive text summarization for condensing legal documents while retaining crucial aspects. In particular, the proposed approach of decreasing any lexical or

syntactic intricacy related to the text without modifying the substance of the text is carried out. It is the pre-processing stage that finally results in the selection of a useful sentence. This work evaluates different classification models using the ROUGE scores. Extractive summarization selects relevant content chunks, ensuring well-structured summaries with all legal elements intact. These methods are favored in legal documents for their preservation of original content. The study advocates for comprehensive legal summaries covering all aspects.

In (Westermann et al., 2023), the author addresses the challenge laypeople face in understanding legal opportunities and remedies due to difficulty in assessing legal issues from factual descriptions. Understanding which legal opportunities or remedies are available to laypeople requires an analysis of which legal issues are raised by these facts, which may be difficult for laypeople to assess. This gap can cause laypeople to miss out on benefits or be unable to resolve their disputes. This research proposes an automated approach to analyze layperson-provided descriptions and map them to relevant legal issues, enhancing access to justice. The findings offer insights for legal professionals and developers to bridge the gap between layperson language and legal issues, potentially improving access to justice through legal decision support systems.

The study presented by (Kiliroor et al., 2023) addresses the challenge of understanding lengthy and complex legal documents, highlighting the importance of accessibility for impartiality. It proposes a text simplification method tailored to the legal domain, aiming to make legal text more comprehensible. The model identifies complex words and substitutes them with simpler alternatives using a word embedding model and sentiment analysis model. Trained on a dataset combining Indian Legal Documents Corpus, the approach successfully detects and replaces complex words with simpler ones, maintaining the original sentiment. This method has the potential to enhance accessibility to legal texts, saving time for individuals navigating legal documents while promoting impartiality.

In (Billi et al., 2023), the authors promote the integration of Large Language Models into rule-based legal systems to enhance accessibility, usability, and transparency, aligning with democratic principles in legal technology. This paper introduces a methodology to translate rule-based system explanations into natural language, enabling clearer and faster interactions for all users. Additionally, it empowers laypeople to perform complex legal tasks independently through a chain of prompts, facilitating autonomous legal comparisons. This approach aims to democratize legal technology, making it more inclusive and comprehensible for users, while also promoting transparency and stakeholder involvement in the legal decision-making process.

3. Representing "Penumbra" in Machine Learning

In this section, we briefly list some possibilities to represent the concept of penumbra from a machine learning point of view and we try to clarify how such representation becomes more complex in multi-linguistic contexts such as the European Union's courts (ECJ, ECHR). Particularly, in the context of natural language processing (NLP) and decision-making algorithms, the penumbra can be linked to areas of uncertainty or ambiguity where the model's predictions may not be unequivocal.

3.1 Uncertainty in Model Predictions:

In machine learning models, especially those based on probabilistic frameworks like Bayesian models, predictions are often associated with a degree of uncertainty (Neil et al., 2019). The model may provide a probability distribution over possible outcomes rather than a definitive answer. This uncertainty may reflect the penumbral aspect, where certain instances may fall into a gray area, making it challenging for the model to make a clear-cut decision.

3.2 Boundary Cases:

Much like legal penumbra involving hard cases, machine learning models may struggle with boundary cases (Atkinson and Bench-Capon, 2019). These are instances that lie on the edge of the decision boundary, where small changes in input features can lead to different predictions. These boundary cases represent situations where the model's confidence is lower, and decisions may be less straightforward.

3.3 Context Sensitivity

In legal terms, the interpretation of a term may vary based on the context in which it is used. Similarly, machine learning models, specifically NLP models, often rely on context to make accurate predictions (Sosa Andrés, 2023). The model's understanding of certain terms or features may exhibit variability based on the surrounding context, introducing a level of interpretation flexibility analogous to the legal penumbra.

3.4 Language Sensitivity

Language is also a big part of the analysis and interpretation of legal terminology (Kalinina and Kudryashova, 2022). The linguistic nature of the term, the specific characteristic of a legal concept, the discrepancies between the state legal systems, the socio-cultural content of legal terms in different languages are only a few examples of the issues concerned with language. Therefore, the combination of language and legal

knowledge, as well as culture understanding, is necessary in understanding the content and translating it into another language functionally and in accordance with the target group.

3.5 Interpretable Machine Learning:

Interpretable machine learning models aim to provide transparency into how decisions are made (Farayola et al., 2023). Despite efforts to achieve interpretability, there may still be instances where the model's reasoning is not entirely clear. This lack of clarity aligns with the penumbral nature, where certain cases may defy straightforward interpretation.

In essence, the concept of penumbra in legal science, with its shades of interpretation ambiguity, can find in machine learning models dealing with uncertainty, boundary cases, context sensitivity, and adaptability.

4. Conclusions and Future Perspectives

In this paper, we started a discussion about research efforts that can explore the intersection of machine learning and legal theory to develop novel approaches for representing the penumbra in legal texts. Drawing upon theoretical frameworks such as Hart's legal theory of meaning, Machine Learning researchers can develop computational models that capture the dynamic nature of legal concepts and their surrounding penumbra. By integrating legal theory into machine-learning algorithms, researchers can create more sophisticated representations of legal ambiguity and uncertainty, facilitating more accurate and contextually appropriate legal interpretations.

We highlighted some ideas of possible research into representing the penumbra concept in legal texts through machine learning which may hold significant promise for advancing the understanding and application of legal semantics in AI-driven legal systems. One important aspect of this investigation lies in refining machine learning models to effectively capture the semantic differences of legal ambiguity inherent in the penumbra. This entails developing algorithms capable of identifying the boundaries of uncertainty within legal texts in order to distinguishing between settled meanings and hard cases, as proposed by Hart. By incorporating probabilistic frameworks such as Bayesian models, researchers could explore how uncertainty in model predictions reflects the penumbral aspect of legal interpretation, providing a more nuanced understanding of complex legal concepts.

Furthermore, research could focus on enhancing machine learning models' sensitivity to contextual variations in legal texts. Just as legal interpretations may vary depending on the context in which terms are used, NLP models must be trained to recognize and adapt to fine-grained contexts within legal documents. Techniques such as contextual embeddings and attention mechanisms can help capture the subtle shifts in meaning that occur within different legal contexts, thereby improving the models' ability to navigate the penumbra of legal interpretation.

Additionally, investigating the impact of language sensitivity on machine learning representations of legal texts is essential, especially in multilingual legal systems like those found in the European Union. Understanding how linguistic differences influence legal interpretation can inform the development of more robust machine learning models capable of handling diverse linguistic and cultural nuances. More specifically, we aim to link the problem of the penumbra in other interdisciplinary areas, such as Digital Humanities. On one hand, the description of the uncertainty of concepts can be used to store and index automatically non-catalogued and unprocessed material, which has to be, not only preserved, but also described, shared and made accessible (Grbac, 2021). On the other hand, the methodology to represent uncertainty can be included in machine translation systems where we have difficulties in the translation processes required in a multilingual and multicultural environment such as that of international cooperation (Vezzani et al., 2022). Finally, the same methodology can be included in other experimental research that deals with semantic phenomena involved in the process of determining a conceptual expression such as synonymy, polysemy, and elliptical segments (Pulizzotto et al., 2018).

In conclusion, research into representing the penumbra concept in legal texts through machine learning offers a rich and multifaceted landscape of opportunities. By refining machine learning models' ability to capture legal ambiguity, sensitivity to contextual variations, and interpretability, researchers can develop more robust and trustworthy AI-driven legal systems. Moreover, integrating legal theory into machine-learning algorithms and exploring innovative learning techniques can further advance our understanding of legal semantics and pave the way for more effective and equitable legal decision-making processes.

5. Bibliographical References

Atkinson, Katie, and Bench-Capon, Trevor (2019). Reasoning with Legal Cases: Analogy or Rule Application? In *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Law*, 12–21. ICAIL '19. New York, NY, USA: Association for Computing

Machinery. https://doi.org/10.1145/3322640.3326695.

Billi, Marco, Parenti, Alessandro, Pisano, Giuseppe, and Sanchi, Marco (2023). Large Language Models and Explainable Law: A Hybrid Methodology. arXiv.Org. November 20, 2023. https://arxiv.org/abs/2311.11811v1.

da Costa Pereira, Célia, Tettamanzi, Andrea G. B., Liao, Beishui, Malerba, Alessandra, Rotolo, Antonino, and van der Torre, Leendert (2017). Combining Fuzzy Logic and Formal Argumentation for Legal Interpretation. In *Proceedings of the 16th Edition of the International Conference on Articial Intelligence and Law*, 49–58. ICAIL '17. New York, NY, USA: Association for Computing Machinery. https://doi.org/10.1145/3086512.3086532.

Dixit, Utkarsh, Gupta, Sonam, Yadav, Arun Kumar, and Yadav, Divakar (2024). Analyzing the Impact of Extractive Summarization Techniques on Legal Text. In *Proceedings of Data Analytics and Management*, edited by Abhishek Swaroop, Zdzislaw Polkowski, Sérgio Duarte Correia, and Bal Virdee, 585–602. Lecture Notes in Networks and Systems. Singapore: Springer Nature. https://doi.org/10.1007/978-981-99-6544-1_44.

Dworkin, Ronald (1975). Hard Cases. *Harvard Law Review*, vol. 88, n. 6, pp. 1057-1109 https://doi.org/10.2307/1340249

Engel, Christoph and McAdams, Richard H. (2024) Asking GPT for the Ordinary Meaning of Statutory Terms. *MPI Collective Goods Discussion Paper*, No. 2024/5, U of Chicago, Public Law Working Paper No. 848, http://dx.doi.org/10.2139/ssrn.4718347

Farayola, Michael Mayowa, Irina Tal, Regina Connolly, Takfarinas Saber, and Malika Bendechache (2023). Ethics and Trustworthiness of AI for Predicting the Risk of Recidivism: A Systematic Literature Review. *Information* 14 (8): 426. https://doi.org/10.3390/info14080426.

Frege, Gottlob (1892). On Sense and Reference ["Über Sinn und Bedeutung"]. Zeitschrift für Philosophie und philosophische Kritik, vol. 100, pp. 25–50.

Grbac, Deborah (2021). The United Nations Depository Libraries System as an "open community": The ongoing evolution from a knowledge base to a knowledge network. Umanistica Digitale 199–216. https://doi.org/10.6092/issn.2532-8816/13676

Hart, Herbert. L. A. and Green, Leslie (2012). *The Concept of Law*. Third Edition, Third Edition. Clarendon Law Series. Oxford, New York: Oxford University Press.

Jiang, Hang, Zhang, Xiajie, Mahari, Robert, Kessler, Daniel, Ma, Eric, August, Tal, Li, Irene, Pentland, Alex, Kim, Yoon, Kabbara, Jad, Roy, Deb (2024). Leveraging Large Language Models for Learning Complex Legal Concepts through Storytelling. arXiv. https://doi.org/10.48550/arXiv.2402.17019.

Kalinina, Marina G., and Kudryashova, Sofya V. (2022). French, Spanish And German Terminology In Legal Discourse –Problematic Aspects In Translation. In *European Proceedings of Social and Behavioural Sciences* State and Law in the Context of Modern Challenges). https://doi.org/10.15405/epsbs.2022.01.45.

Kiliroor, Cinu C., Sagar, Som and Sundara Didde, Swani (2023). Sentiment-Based Simplification of Legal Text. In *Proceedings of the 4th International Conference on Communication, Devices and Computing*, edited by Dilip Kumar Sarkar, Pradip Kumar Sadhu, Sunandan Bhunia, Jagannath Samanta, and Suman Paul, 463–75. Lecture Notes in Electrical Engineering. Singapore: Springer Nature. https://doi.org/10.1007/978-981-99-2710-4_38.

Neil, Martin, Fenton, Norman, Lagnado, David, and Gill, Richard David (2019). Modelling Competing Legal Arguments Using Bayesian Model Comparison and Averaging. *Artificial Intelligence and Law* 27 (4): 403–30. https://doi.org/10.1007/s10506-019-09250-3.

Pulizzotto, Davide, Chartier, Jean-François, Lareau, Francis, Meunier, Jean-Guy, Chartrand, Louis, 2018. Conceptual Analysis in a computer-assisted framework: mind in Peirce. Umanistica Digitale https://doi.org/10.6092/issn.2532-8816/7305

Quevedo, Ernesto, Cerny, Tomas, Rodriguez, Alejandro, Rivas, Pablo, Yero, Jorge, Sooksatra, Korn, Zhakubayev, Alibek, and Taibi, Davide (2023). Legal Natural Language Processing from 2015-2022: A Comprehensive Systematic Mapping Study of Advances and Applications. *IEEE Access*, 1–1. https://doi.org/10.1109/ACCESS.2023.3333946.

Rissland, Edwina L., Ashley, Kevin D., and Loui, Ronald P. (2003). AI and Law: A Fruitful Synergy. *Artificial Intelligence*, AI and Law, 150 (1): 1–15. https://doi.org/10.1016/S0004-3702(03)00122-X.

Sosa Andrés, Maximiliano (2023). Legal Uncertainty and Its Consequences: A Natural Language Processing Approach. https://urn.kb.se/resolve?urn=urn:nbn:se:uu:diva-495521.

Stathis, Georgios, Biagioni, Giulia, de Graaf, Klaas Andries, Trantas, Athanasios, and van den Herik, Jaap (2024). The Value of Proactive Data for Intelligent Contracts. In *Intelligent Sustainable Systems*, pp. 107–25. Lecture Notes in Networks and Systems. Singapore: Springer Nature. https://doi.org/10.1007/978-981-99-7569-3_10.

Vezzani, Federica, Di Nunzio, Giorgio Maria, Silecchia, Sara (2022). La fraseologia dei trattati internazionali di disarmo: la risorsa terminologica DITTO. Umanistica Digitale 91–

117. https://doi.org/10.6092/issn.2532-8816/14796

Westermann, Hannes, Meeùs, Sébastien, Godet, Mia, Troussel, Aurore, Tan, Jinzhe, Savelka, Jaromir, and Benyekhlef, Karim (2023). Bridging the Gap: Mapping Layperson Narratives to Legal Issues with Language Models. In *Proceedings of the 6th Workshop on Automated Semantic Analysis of Information in Legal Text*, CEUR Workshop Proceedings. Braga, Portugal: CEUR. https://ceur-ws.org/Vol-3441/#paper5.

Whalen, Ryan (2022). Defining Legal Technology and Its Implications. *International Journal of Law and Information Technology* 30 (1): 47–67. https://doi.org/10.1093/ijlit/eaac005.