

Toward a Neuro-Symbolic Approach to Structure and Process International Tax Law

Zarja Hude¹

¹ London School of Economics and Political Science, London WC2A 2AE, UK
z.hude@lse.ac.uk

Abstract. Large Language Models (LLMs) have demonstrated remarkable promise in performing legal tasks. Yet, trust in their use remains limited due to non-determinism and opacity—qualities that conflict with the transparent, interpretable, and auditable nature of legal reasoning. This PhD project explores how neuro-symbolic methods, which combine neural networks with logic-based systems, can bridge this divide. Drawing on legal positivism and particularly the legacy of H. L. A. Hart, the study frames symbolic AI as aligned with structured reasoning, and views neural models as instruments of interpretive flexibility, comparable to judicial discretion in "hard cases." The project proposes a framework that combines structure and discretion, mirroring Hart's theory of adjudication. It argues that symbolic AI is best suited for handling "easy cases," where legal rules are clear and determinate, while neural AI can fill the interpretative gaps where the rules "run out." Through this novel neuro-symbolic framework, the project recontextualises Hart's legal theory for the AI era and shows how subsequent critiques can guide future research in this theoretically-grounded approach to structuring and processing law and legal reasoning.

Keywords: Legal AI, Neuro-Symbolic AI, Legal Reasoning.

1 Motivation and Research Questions

LLMs have demonstrated remarkable capabilities in performing legal tasks.¹ Yet, despite their achievements, LLMs have yet to gain widespread trust, primarily due to two critical limitations: non-determinism and opacity (Barbiero et al. 2023). These limitations are not incidental but inherent to the data-driven, associative nature of LLMs. As a subset of neural AI, LLMs use probabilistic modelling, making them non-deterministic under standard settings—from the same input they can produce different outputs. This inconsistency poses significant challenges in legal contexts, where reliability is essential. Opacity, meanwhile, arises from the way decisions are made. Because the decision-making process is distributed across billions of parameters rather than explicit rules, LLMs lack the transparency and interpretability needed to earn trust in high-stakes domains such as law.

Before the rise of LLMs, legal AI was dominated by symbolic methods that encoded legal rules and reasoning into structured, logic-based frameworks to mirror human legal

¹ See, for instance, surveys conducted by Lai et al. (2024) and Katz et al. (2023).

problem-solving.² During this symbolic AI era, the challenges of non-determinism and opacity were largely absent. Logic-based systems operate with deterministic precision, allowing decisions to be traced back to their source. Conclusions follow explicitly defined reasoning paths, providing transparency, auditability, and interpretability. Given identical inputs, symbolic systems consistently produce the same outputs—delivering a reliability that modern LLMs struggle to match.

Yet, knowledge-based systems face challenges of their own. Developing them demands significant effort to define explicit rules and construct logical pathways, resulting in rigid architectures that are difficult to scale. Unlike human reasoning, which fluidly adapts to nuance and interpretive complexity, symbolic AI struggles with the context-dependent nature of law, which resists being reduced to explicit rules and logical deductions. Consequently, many attempts at creating legal knowledge-based systems failed to gain widespread adoption or scale effectively. Here, LLMs offer a complement—their flexibility and scalability show promise in overcoming longstanding limitations of symbolic systems.³

Since neural and symbolic approaches have complementary strengths that can offset each other's limitations, their combination presents a promising direction for structuring and processing law and legal reasoning. This PhD examines how to best combine the two approaches to reflect legal thinking processes, using tax law as a case study. It further explores whether this combined approach could enhance reasoning systems beyond the legal domain. The investigation unfolds through two main research questions.

- **RQ1:** How can symbolic and neural methods be effectively combined to structure and process international tax law?
 - **RQ1.1:** Can the feasibility of the proposed combination be validated through a case study focused on individual tax residency tests?
- **RQ2:** Can this neuro-symbolic combination be generalised beyond the legal domain, specifically to inform the development of explainability frameworks for general-purpose AI systems?

The outcome of RQ1 will be a theoretical proposal, which will be tested on individual tax residency tests across various jurisdictions through RQ1.1. After validation, RQ2 will examine whether the framework can extend beyond tax law—first to legal reasoning broadly, and ultimately to general-purpose AI systems.

The thematic link between RQ1 and RQ2 lies in explainability—a critical requirement in both legal reasoning and AI systems. If the proposed neuro-symbolic approach (RQ1) successfully captures legal reasoning's transparency and clarity (RQ1.1), we can then explore its potential contribution to general-purpose AI systems (RQ2), where explainability remains essential yet largely unachieved.

² Knowledge-based systems focusing on tax law include McCarty (1977), Hellawell (1980), Michaelsen (1984), Sanders (2001), Merigoux et al. (2021), and Lawsky (2022).

³ See, for instance, Kant et al. (2024) and Kant et al. (2025).

2 Related Work and Research Challenges

Over the last five decades, the field of legal AI research has undergone a dramatic evolution. What began as conceptual explorations of AI's potential in the legal domain⁴ has progressed through legal expert systems,⁵ legal analytics,⁶ legal NLP and machine learning models,⁷ and domain-specific LLMs.⁸

This brief review examines two main trajectories in this evolution: the "symbolic track," characterised by logic-based systems, and the "neural and LLM track," driven by neural architectures and statistical learning. Though these tracks developed largely in parallel, there is now growing interest in combining them through neuro-symbolic AI.

2.1 Symbolic AI Track

Symbolic AI in law focuses on formal logic and structured knowledge representation in legal contexts. Efforts include formalising and coding legal sources,⁹ developing structured argumentation and reasoning frameworks¹⁰ and processing legal information in a structured way.

While these systems show potential, they struggle with the size, complexity and nuance of law. Creating rules demands extensive manual effort, and purely logical structures struggle to capture deeper semantic legal meanings. These limitations have motivated the shift toward data-driven, neural and LLM approaches in legal AI.

2.2 Neural and LLM Track

In 2010s, machine learning and NLP enabled models to learn from large legal datasets. This opened new possibilities for machine summarisation, data preprocessing, and classification, information retrieval and extraction, and text generation.¹¹ Advances in deep learning further expanded the capabilities of legal NLP, allowing models to engage with legal texts more effectively.

Despite their strengths, LLMs are yet to earn full trust in legal settings. Their reliability is limited by their non-deterministic nature—they can generate different outputs for the same input under standard settings. Additionally, their internal reasoning processes remain largely opaque, making it impossible to trace how decisions are reached. These features at odds with the ideals of transparency and traceability that

⁴ See, for instance, Gardner (1984).

⁵ See, for instance, Sergot et al. (1986).

⁶ See, for instance, Katz (2012).

⁷ See, for instance, Aletras et al. (2016).

⁸ See, for instance, Colombo et al. (2024).

⁹ See, for instance, Sergot et al. (1986), Hoekstra et al. (2007), Satoh et al. (2011), Merigoux, Chataing, and Protzenko (2021), Listenmaa et al. (2021), Lawsky (2022).

¹⁰ See, for instance, Ashley (1991), Verheij (1999), Bench-Capon and Sartor (2003), Bex et al. (2003), Gordon and Walton (2006), Zurek and Araszkiwicz (2013), Al-Abdulkarim et al., Prakken et al. (2015).

¹¹ Taxonomy of engineering tasks in legal NLP is taken from Katz et al. (2023).

underpin legal reasoning. In reverse, symbolic AI systems are deterministic, showing reasoning paths to original sources and ensuring consistent outputs for identical inputs.

2.3 Neuro-Symbolic Track

Combining neural and symbolic AI provides a promising way to overcome their separate limitations. The key question becomes: how can the two approaches be best combined in the legal context?

Several legal neuro-symbolic architectures have already been explored. They can be compared using Kautz's (2022) framework, with three of his six categories being particularly relevant:

1. **Neuro | Symbolic architecture.** Machine learning techniques are first applied to structure information, followed by logic-based methods that reason with this structured data. Examples of this approach can be found in LUIMA in Ashley (2017), Branting (2020), Mumford et al. (2023), Nguyen et al. (2023), Kant et al. (2024), and Kant et al. (2025).
2. **Neuro[Symbolic] architecture.** Symbolic reasoning is integrated directly within a neural system. Drawing on Kahneman (2011) dual-system theory, this design mimics the interplay between fast, intuitive processing (System 1) and slower, deliberate reasoning (System 2). One of the options outlined in Nguyen et al. (2023) exemplifies this approach.
3. **Symbolic[Neuro] architecture.** Neural pattern recognition is employed as a subroutine within a broader symbolic reasoning framework. While Zeleznikow et al. (1995) provides the closest reference to this model, this PhD takes a different approach. Instead of using neural and symbolic components separately, we propose combining them into a unified system where they complement each other.

Unlike prior work, this project proposes a dual neuro-symbolic approach, employing a Neuro | Symbolic architecture for legal data structuring and a Symbolic[Neuro] architecture for legal data processing. In the *structuring* phase, LLMs can help create structured knowledge-based systems. During the *processing* phase, the LLMs are proposed to be invoked only when encoded rules prove insufficient to resolve the issue at hand. This design choice is grounded in legal philosophy—specifically, Hart's theory of adjudication. In Hart's terms, LLMs are proposed to fill the interpretative gap when the rules run out, mirroring the judicial discretion in hard cases.

3 Philosophical Foundation

To understand the best way to combine neuro-symbolic AI in the legal domain, one must reflect on the nature of law itself: What role do norms play in the legal system? What defines a legal system? And ultimately, what is the very concept of law? This section explores these fundamental questions through Hart's theory of adjudication (1961). Rooted in legal positivism, Hart's framework for understanding law has taken on fresh significance in the age of AI.

3.1 Legal Positivism and H. L. A. Hart

Hart was a proponent of legal positivism, a philosophical approach that views law as *is*—expressed through explicitly written rules enacted through formal processes—not as it *should be*. In this view, legal rules exist independently of ethics and morality.

Hart's philosophy was later critiqued as incomplete by philosophers such as Dworkin, Alexy, MacCormick, among others. In short, they argued that legal adjudication requires more than just legal rules, stressing the role of legal principles and moral considerations.

Nonetheless, these philosophical developments do not invalidate the importance of legal or diminish the value of structuring legal rules alone as a starting point. Rather, they suggest that systems should eventually expand to include additional layers of normative constraints, such as values and principles. The focused scope on legal rules and their structured application through Hartian philosophy thus acknowledges potential later developments while maintaining a focused analytical scope.

Focusing solely on explicit rules also helps avoid critical risks that come with alternative approaches. Attempting to capture unwritten or subjective elements would require the system to determine what the law *should be*—a task that lies beyond this PhD's purpose and introduces significant technical and bias risks.

Though the model remains grounded in legal positivism and its inherent limitations, it can still offer valuable insights to support human adjudication. Moreover, it remains open to later extensions that might incorporate modern philosophical approaches to additional normative constraints in the future.

3.2 Hart's Theory of Adjudication

Hart, writing in the mid-20th century, set out to answer a fundamental question in legal theory: how do judges decide cases? At the time, legal scholarship was dominated by two opposing views. *Legal formalists* claimed that judges ought to reason by mechanically applying legal rules to facts, as if legal reasoning were a closed logical system. On the other end, *legal realists* argued that judges actually make decisions based on their own preferences or broader social and political considerations, and only later rationalise those outcomes through legal reasoning.

This tension between formalism and realism offers an interesting parallel to the modern divide between symbolic and neural approaches in AI. Formalists viewed the law as a system of explicit, deductive rules—akin to *symbolic AI*, which operates through structured representations and logical inference. Crucially, the degree of formalisation is flexible: it can stay high-level to permit flexibility, or detail constraints to limit interpretation. In contrast, realists emphasised the intuitive and context-sensitive aspects of legal reasoning—comparable to the behaviour *neural AI*, relying on data-driven learning and statistical associations without pre-defined logic.

Hart found both perspectives inadequate. In his view, formalism failed to capture the flexibility required in actual legal reasoning, whilst realism overstated the extent to which law is indeterminate and judicial decision-making unconstrained. In response, Hart developed his own theory of adjudication.

He acknowledged that in most cases, judges can apply existing legal rules to resolve disputes in a relatively straightforward way. These are "easy cases", where the relevant facts clearly fall within the established meaning of a rule. However, legal rules may not

always yield definite outcomes. The language of existing rules may be vague, contested, or silent on the specific issue. In such “hard cases,” judges must exercise a degree of discretion to fill the interpretative gap.

By exercising discretion, judges effectively establish a new point of reference for future cases. In this sense, they are not merely applying existing law but, in a limited sense, creating new law. This interpretative function keeps the legal system both adaptable and responsive to changes in a constantly evolving society.

Crucially, judicial discretion is neither absolute nor arbitrary but constrained by established legal norms. Discretion, then, is best understood as a constrained form of lawmaking—an interpretive function that allows judges to extend the law to address novel or ambiguous circumstances while remaining anchored in the broader structure of the existing legal system.

3.3 Hart and Neuro-Symbolic AI

Hart's balanced view of adjudication offers a valuable framework for reconciling these perspectives through the lens of neuro-symbolic AI. This project hypothesises that knowledge-based symbolic AI is well-suited to clear-cut situations, where the input directly maps to the knowledge the system has. When there is no direct match, neural models provide a degree of adaptability that mirrors the judicial discretion required in “hard cases.”

Just as Hart acknowledged that judicial discretion is constrained by the legal framework—and is not a license for unconstrained decision-making—so too can neural components function within well-defined boundaries. They can be activated only when symbolic reasoning, grounded in encoded rules, reaches its interpretative limits and the rules “run out.” This neuro-symbolic approach is grounded in the authority and structure of rules yet remains flexible enough to navigate the uncertainties and complexities that inevitably arise in legal practice.

4 The Proposed Neuro-Symbolic Approach

Building on the complementary strengths of neural and symbolic AI components and the Hartian framework, this section proposes a method to combine both approaches for structuring and processing legal data, using international tax law as a case study. Since this project is in its early stages, the technical implementation relies on several hypotheses that will be tested as the project develops.

4.1 Symbolic[Neuro] Approach for Legal Data Processing

This section bridges Hart's theory of adjudication with neuro-symbolic AI system design, revealing a crucial divide: while judges engage in *adjudication*, LLMs only make *predictions*. This distinction suggests that the neuro-symbolic AI system can, at most, serve to inform adjudicative processes but not replace them.

Bridging Hart and Neuro-Symbolic AI. As discussed, Hart’s framework distinguishes between “easy cases”, where facts of the case clearly fall within the core

meaning of rules, and "hard cases," where questions lie in a margin of uncertainty. In hard cases, judicial discretion bridges the interpretative gap between the rule and the issue at hand. This discretion is not arbitrary—it remains bound by existing rules and the legal framework.

In computational terms, "easy cases" have legal rules that map directly to the input, enabling reasoning through formal knowledge-based approaches alone. However, not all legal questions are this straightforward. Rules can "run out," leaving their application ambiguous or incomplete. In these "hard cases," where rules do not directly address the issue at hand, interpretation becomes necessary. To fill this interpretative gap, we propose invoking LLMs to predict the most likely resolution.

This neuro-symbolic approach mimics how judges, as per Hart, exercise rule-bound discretion when confronting situations that do not fit neatly within the established rule meanings. Because symbolic logic is used at the primary structure and LLMs are invoked only for interpretative gaps, the architecture falls under the Symbolic[Neuro] category in Kautz’s (2022) categorisation.

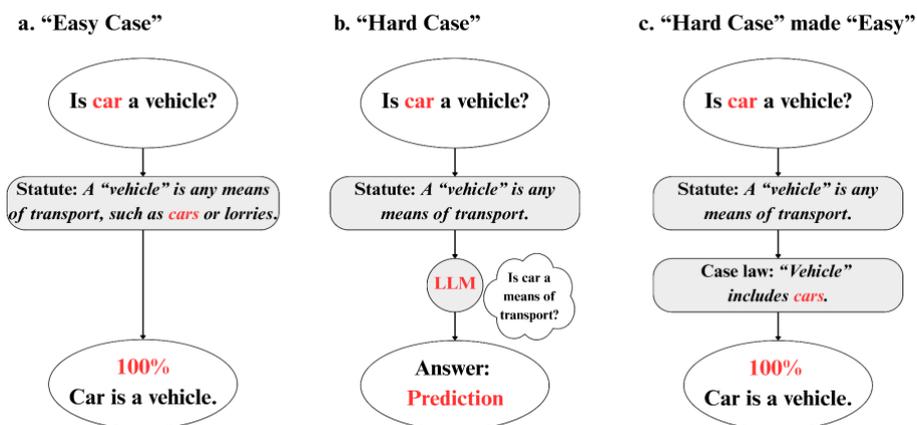


Fig 1: "Easy" and "hard" cases in a Symbolic[Neuro] Approach in the Legal Data Processing Stage.

The Divide: Hart’s Adjudication and LLM Prediction. It is crucial, however, to distinguish between Hart's view of *adjudication* and what LLMs actually do: *prediction*. For Hart, adjudication is a normative, institutional process—how judges ought to decide, especially when rules are unclear. It involves authority and responsibility within a legal system. LLMs, by contrast, generate statistically likely text based on patterns in data. They provide no authority, no normative judgment, and no accountability—just prediction.

While this indicates that neuro-symbolic AI systems cannot replace adjudicative processes, their predictive outputs remain valuable for informing decisions. Especially outside the courtrooms, these systems can help to navigate legal information, explore plausible interpretations, and understand possible legal outcomes. This potential forms both the motivation and orientation of this PhD project.

4.2 Neuro | Symbolic Approach for Legal Data Structuring

While Hart's theory corresponds to legal data processing, modern computational systems require data to be structured first. The project hypothesises that a neuro-symbolic approach could enhance this initial structuring phase as well.

Law, particularly from a legal positivist perspective, is fundamentally grounded in legal texts. Legal texts typically exist in natural language in, at best, semi-structured formats. Raw text in this format has limited utility and poses significant challenges for computational processing. To derive meaningful insights, legal text must be organised into data structures. Several data structures have been explored to organise legal information, including tables, databases, trees, and graphs.¹²

The project hypothesises that graph-based data structures offer the best approach to structuring law. This approach requires careful organisation to capture both the context and complexity of interconnected legal concepts. With proper contextualisation, these graph structures can effectively transform raw legal data into actionable knowledge.

The challenges of manually structuring legal data (*see limitations in the “Symbolic AI Track” section*) and LLMs' demonstrated potential to address these longstanding issues, make a case for LLM-assisted legal data structuring—which falls under the Neuro | Symbolic category in Kautz's (2022) classification. While trust-related issues persist during the structuring stage (*see limitations in the “Neural and LLM Track” section*), these risks can be mitigated through human oversight in identifying and structuring legal concepts. Unlike legal data processing, the structuring stage works with *finite* sets of legal texts, suggesting that the need for human intervention will gradually decrease over time. This reduction in manual involvement should enhance the approach's scalability.

The author also acknowledges limitations of graph data structures themselves, particularly their computational intensity and limited scalability when compared to, for instance, relational databases or tables. While the foundation is promising, its practical value remains to be proven as the project unfolds.

4.3 Tax Law Case Study

The choice to evaluate the proposed neuro-symbolic approach using individual tax residency tests is deliberate.

1. **Complexity and Scale.** Tax law is notoriously complex due to its technical structure, frequent changes, and vast, diverse sources of data (Lawsky 2017). It is typically the longest law that a country has (Harris and Oliver 2010). While this complexity, which grows more rapidly than in other fields (Katz et al. 2020), challenges even experts, it also creates opportunities for computational systems to analyse and apply tax law more effectively by leveraging its rich and abundant regulatory data.
2. **Clarity, Predictability and Algorithmic Nature.** Tax law stands out for its emphasis on clarity and predictability (Nay et al. 2023), making it easier to model,

¹² See, for instance, Flood and Goodenough (2014), Koniaris et al. (2018), Yu et al. (2020), Coupette et al. (2021), Hanauer et al. (2023).

structure and automate. Its reliance on logic, rules, and numerical calculations further enhances its compatibility with computational systems. This structured nature has already led to widespread adoption of tax software by both governments and private sector (Lawsky 2020).

3. **Social Fairness.** The complexity of tax law allows those that can acquire or afford specialised legal expertise to exploit loopholes while leaving others at a disadvantage (Lawsky 2024). This contributes to unequal access to legal certainty and increases the overall burden on both less-resourced taxpayers and tax administrations.
4. **Democratic Legitimacy.** The complexity of tax law can weaken democratic legitimacy, as unresolved ambiguities shift interpretive authority from elected lawmakers to unelected administrators or judges. This transfer of power undermines the principle that no taxation should occur without representation.
5. **Individual tax residency tests.** A case study of individual tax residency tests—based on the OECD Model Convention and its national adaptations—provides a realistic benchmark. The setup combining shared criteria and jurisdiction-specific variations allows for meaningful comparisons, evaluation of cross-jurisdictional portability, and collection of empirical data on performance, transparency, and error patterns—three metrics that current literature rarely examines in combination.

5 Contributions

Rather than introducing an entirely new computational paradigm, this project builds on well-established methodologies from both law and data science. Its primary contribution lies in demonstrating—through a unified framework—that symbolic rule-encoding and neural language modelling can operate side by side, each compensating for the other's limitations within a legally significant task.

Another key innovation is grounding neuro-symbolic system design in legal philosophy. The choice of legal positivism and Hart's theory of adjudication offers a framework that is both feasible to implement and adaptable. The approach allows for future extensions to incorporate additional constraints proposed by Hart's critics like Dworkin, Alexy, and MacCormick.

Moreover, the framework spans both legal data structuring *and* processing—two stages that previous work has often treated in isolation. Narrower task optimisation has led to specialised, standalone solutions that have proven difficult to adapt or extend within broader legal workflows, limiting their utility beyond their original scope. In contrast, the proposed architecture is built with the entire legal-AI pipeline in mind—from data structuring and representation to structured reasoning, LLM integration, and output generation. This end-to-end perspective ensures that design choices are made with a view toward the complete workflow, and ambition to provide a reusable foundation for future legal AI projects.

Furthermore, testing a neuro-symbolic framework on real individual tax residency tests across various jurisdictions aims to provide a realistic and practical benchmark for evaluating its performance. The developed method could contribute to legal AI evaluation—an area that remains largely underexplored.

Finally, this thesis explores how a neuro-symbolic framework inspired by legal reasoning methods could advance research on explainability in general-purpose AI systems. This intersection remains largely unexplored, likely due to the scarcity of expertise spanning both relevant domains.

6 Conclusion

This PhD project investigates how neuro-symbolic methods can be combined to structure and process legal knowledge—particularly in international tax law—and whether such a framework can inform research in the explainability of general-purpose AI systems. The motivation stems from two major limitations of current LLMs: their non-deterministic nature and lack of interpretability. These flaws are especially problematic in law, where transparency, traceability, and normative justification are essential.

Legal AI has evolved from symbolic systems based on formal logic to data-driven approaches advanced by neural networks and machine learning techniques, including LLMs. While symbolic methods offer precision and clarity, they lack flexibility and scalability. Neural models, on the other hand, can process vast legal corpora and adapt to new inputs, but they are often opaque and unaccountable. These trade-offs have sparked interest in hybrid approaches that combine both methods to overcome their individual limitations—though finding the right balance remains a challenge.

Philosophically, the project revisits Hart's theory of adjudication. Hart viewed legal rules as having an "open texture"—while they have clear core meanings, there are edge cases where interpretation becomes necessary and judges must exercise discretion. His view of judicial discretion in these "hard cases" is paralleled with the potential role of neural models in handling legal indeterminacy. Hart's notion of "discretion within structure" resonates with how symbolic constraints can guide neural flexibility—a balance this project aims to implement computationally and test through a case study on individual tax residency tests.

This research explores the emerging field of neuro-symbolic AI for law, grounding the research in both computational and jurisprudential theory. It aims to contribute a conceptual and technical foundation for a neuro-symbolic framework in law for both legal data structuring and processing, tested on a real-world case study of individual tax residency tests. The project aims to develop a structured yet adaptable model of legal reasoning that can be generalised across legal and broader contexts, opening new pathways for explainable AI in law and beyond.

Acknowledgments. This work has greatly benefited from the supervision, guidance, and support of Andy Summers (LSE Law), Eduardo Baistrocchi (LSE Law), and Milan Vojnović (LSE Statistics). I gratefully acknowledge the LSE Studentship, Ad Futura Studentship, Law Brainer, and Hack the Law by King's Entrepreneurship Lab for their support. I also thank Pedro Pereira de Morais Pacheco and Juan Kostelec for their valuable feedback and contributions.

Disclosure of Interests. The author declares no competing interests.

References

- Al-Abdulkarim, Latifa, Katie Atkinson, and Trevor Bench-Capon. 2014. 'Abstract Dialectical Frameworks for Legal Reasoning'. In *Legal Knowledge and Information Systems*, IOS Press, 61–70. doi:10.3233/978-1-61499-468-8-61.
- Aletras, Nikolaos, Dimitrios Tsarapatsanis, Daniel Preoțiu-Pietro, and Vasileios Lampos. 2016. 'Predicting Judicial Decisions of the European Court of Human Rights: A Natural Language Processing Perspective'. *PeerJ Computer Science* 2: e93. doi:10.7717/peerj-cs.93.
- Ashley, Kevin D. 1991. *Modeling Legal Arguments: Reasoning with Cases and Hypotheticals*. Cambridge, MA, USA: MIT Press.
- Ashley, Kevin D. 2017. *Artificial Intelligence and Legal Analytics: New Tools for Law Practice in the Digital Age*. Cambridge: Cambridge University Press. doi:10.1017/9781316761380.
- Barbiero, Pietro, Gabriele Ciravegna, Francesco Giannini, Mateo Espinosa Zarlenga, Lucie Charlotte Magister, Alberto Tonda, Pietro Lio', et al. 2023. 'Interpretable Neural-Symbolic Concept Reasoning'. doi:10.48550/ARXIV.2304.14068.
- Bench-Capon, Trevor, and Giovanni Sartor. 2003. 'A Model of Legal Reasoning with Cases Incorporating Theories and Values'. *Artificial Intelligence* 150(1): 97–143. doi:10.1016/S0004-3702(03)00108-5.
- Bex, Floris, Henry Prakken, Chris Reed, and Douglas Walton. 2003. 'Towards a Formal Account of Reasoning About Evidence: Argumentation Schemes and Generalisations'. *Artificial Intelligence and Law* 11(2–3): 125–65. doi:10.1023/b:arti.0000046007.11806.9a.
- Branting, L Karl. 'Explanation in Hybrid, Two-Stage Models of Legal Prediction'.
- Colombo, Pierre, Telmo Pessoa Pires, Malik Boudiaf, Dominic Culver, Rui Melo, Caio Corro, Andre F. T. Martins, et al. 2024. 'SaulLM-7B: A Pioneering Large Language Model for Law'. doi:10.48550/arXiv.2403.03883.
- Coupette, Corinna, Janis Beckedorf, Dirk Hartung, Michael Bommarito, and Daniel Martin Katz. 2021. 'Measuring Law Over Time: A Network Analytical Framework with an Application to Statutes and Regulations in the United States and Germany'. *Frontiers in Physics* 9: 658463. doi:10.3389/fphy.2021.658463.
- Flood, Mark D., and Oliver R. Goodenough. 2014. 'Contract as Automaton: The Computational Representation of Financial Agreements'. *SSRN Electronic Journal*. doi:10.2139/ssrn.2538224.
- Gardner, Anne Von Der Lieth. 1984. 'An Artificial Intelligence Approach to Legal Reasoning'. phd. Stanford University.

- Gordon, Thomas F., and Douglas Walton. 2006. 'The Carneades Argumentation Framework - Using Presumptions and Exceptions to Model Critical Questions'. In *Computational Models of Argument*, IOS Press, 195–207. <https://ebooks.iospress.nl/publication/2947> (May 15, 2025).
- Hanauer, Kathrin, Novotná, Tereza, and Matteo Pascucci. 2023. 'Assisted Normative Reasoning with Aristotelian Diagrams'. In *Legal Knowledge and Information Systems*, IOS Press, 89–94. doi:10.3233/FAIA230949.
- Harris, Peter, and David Oliver. 2010. *International Commercial Tax*. Cambridge: Cambridge University Press. doi:10.1017/CBO9780511777028.
- Hart, Hla. 1961. *The Concept of Law*. Oxford, United Kingdom: Oxford University Press UK.
- Hellawell, Robert. 1980. 'A Computer Program for Legal Planning and Analysis: Taxation of Stock Redemptions'. *Columbia Law Review* 80(7): 1363. doi:10.2307/1122169.
- Hoekstra, R., J. Breuker, M. Bello, and A. Boer. 2007. 'The LKIF Core Ontology of Basic Legal Concepts'. <https://www.semanticscholar.org/paper/The-LKIF-Core-Ontology-of-Basic-Legal-Concepts-Hoekstra-Breuker/029b9f01cd61944135c957baa23de3ef3d3c40fd> (May 15, 2025).
- Kahneman, Daniel. 2011. *Thinking, Fast and Slow*. New York, NY, US: Farrar, Straus and Giroux.
- Kant, Manuj, Manav Kant, Marzieh Nabi, Preston Carlson, and Megan Ma. 2024. 'Equitable Access to Justice: Logical LLMs Show Promise'. doi:10.48550/arXiv.2410.09904.
- Kant, Manuj, Sareh Nabi, Manav Kant, Roland Scharrer, Megan Ma, and Marzieh Nabi. 2025. 'Towards Robust Legal Reasoning: Harnessing Logical LLMs in Law'. doi:10.48550/arXiv.2502.17638.
- Katz, Daniel Martin. 2012. 'Quantitative Legal Prediction – or – How I Learned to Stop Worrying and Start Preparing for the Data Driven Future of the Legal Services Industry'. <https://papers.ssrn.com/abstract=2187752> (May 15, 2025).
- Katz, Daniel Martin, Corinna Coupette, Janis Beckedorf, and Dirk Hartung. 2020. 'Complex Societies and the Growth of the Law'. *Scientific Reports* 10(1): 18737. doi:10.1038/s41598-020-73623-x.
- Katz, Daniel Martin, Dirk Hartung, Lauritz Gerlach, Abhik Jana, and Michael J. Bommarito II. 2023. 'Natural Language Processing in the Legal Domain'. doi:10.48550/arXiv.2302.12039.
- Kautz, Henry A. 2022. 'The Third AI Summer: AAAI Robert S. Engelmore Memorial Lecture'. *AI Magazine* 43(1): 105–25. doi:10.1002/aaai.12036.
- Koniaris, Marios, Ioannis Anagnostopoulos, and Yannis Vassiliou. 2018. 'Network Analysis in the Legal Domain: A Complex Model for European Union Legal Sources'. *Journal of Complex Networks* 6(2): 243–68. doi:10.1093/comnet/cnx029.

- Lai, Jinqi, Wensheng Gan, Jiayang Wu, Zhenlian Qi, and Philip S. Yu. 2024. 'Large Language Models in Law: A Survey'. *AI Open* 5: 181–96. doi:10.1016/j.aiopen.2024.09.002.
- Lawsky, Sarah B. 2017. 'Formalizing the Code'. doi:10.2139/ssrn.2987065.
- Lawsky, Sarah B. 2020. 'Form as Formalization'. <https://papers.ssrn.com/abstract=3587576> (December 15, 2024).
- Lawsky, Sarah B. 2022. 'Coding the Code: Catala and Computationally Accessible Tax Law'. <https://papers.ssrn.com/abstract=4291177> (December 15, 2024).
- Lawsky, Sarah B. 2024. 'Tax Law and Flexible Formalizations'. doi:10.2139/ssrn.5026598.
- Listenmaa, Inari, Maryam Hanafiah, Regina Cheong, and Andreas Källberg. 2021. 'Towards CNL-Based Verbalization of Computational Contracts'. In *Proceedings of the Seventh International Workshop on Controlled Natural Language (CNL 2020/21)*, eds. Tobias Kuhn, Silvie Spreeuwenberg, Stijn Hoppenbrouwers, and Norbert E. Fuchs. Amsterdam, Netherlands: Special Interest Group on Controlled Natural Language. <https://aclanthology.org/2021.cnl-1.10/> (May 15, 2025).
- McCarty, L. Thorne. 1977. 'Reflections on "Taxman": An Experiment in Artificial Intelligence and Legal Reasoning'. *Harvard Law Review* 90(5): 837–93. doi:10.2307/1340132.
- Merigoux, Denis, Nicolas Chataing, and Jonathan Protzenko. 2021. 'Catala: A Programming Language for the Law'. *Proceedings of the ACM on Programming Languages* 5(ICFP): 1–29. doi:10.1145/3473582.
- Michaelsen, Robert H. 1984. 'An Expert System for Federal Tax Planning'. *Expert Systems* 1(2): 149–67. doi:10.1111/j.1468-0394.1984.tb00440.x.
- Mumford, Jack, Katie Atkinson, and Trevor Bench-Capon. 2023. 'Combining a Legal Knowledge Model with Machine Learning for Reasoning with Legal Cases'. In *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law, ICAIL '23*, New York, NY, USA: Association for Computing Machinery, 167–76. doi:10.1145/3594536.3595158.
- Nay, John J., David Karamardian, Sarah B. Lawsky, Wenting Tao, Meghana Bhat, Raghav Jain, Aaron Travis Lee, Jonathan H. Choi, and Jungo Kasai. 2023. 'Large Language Models as Tax Attorneys: A Case Study in Legal Capabilities Emergence'. doi:10.48550/arXiv.2306.07075.
- Nguyen, Ha-Thanh, Francesca Toni, Kostas Stathis, and Ken Satoh. 2023. 'Beyond Logic Programming for Legal Reasoning'. doi:10.48550/arXiv.2306.16632.
- Prakken, Henry, Adam Wyner, Trevor Bench-Capon, and Katie Atkinson. 2015. 'A Formalization of Argumentation Schemes for Legal Case-Based Reasoning in ASPIC+'. *Journal of Logic and Computation* 25(5): 1141–66. doi:10.1093/logcom/ext010.

- Sanders, Kathryn E. 2001. 'CHIRON: Planning in an Open-Textured Domain'. *Artificial Intelligence and Law* 9(4): 225–69. doi:10.1023/A:1013824413224.
- Satoh, Ken, Kento Asai, Takamune Kogawa, Masahiro Kubota, Megumi Nakamura, Yoshiaki Nishigai, Kei Shirakawa, and Chiaki Takano. 2011. 'PROLEG: An Implementation of the Presupposed Ultimate Fact Theory of Japanese Civil Code by PROLOG Technology'. In *New Frontiers in Artificial Intelligence*, Lecture Notes in Computer Science, eds. Takashi Onada, Daisuke Bekki, and Elin McCready. Berlin, Heidelberg: Springer Berlin Heidelberg, 153–64. doi:10.1007/978-3-642-25655-4_14.
- Sergot, M. J., F. Sadri, R. A. Kowalski, F. Kriwaczek, P. Hammond, and H. T. Cory. 1986. 'The British Nationality Act as a Logic Program'. *Communications of the ACM* 29(5): 370–86. doi:10.1145/5689.5920.
- Verheij, Bart. 1999. 'Two Approaches to Dialectical Argumentation: Admissible Sets and Argumentation Stages'. <https://www.semanticscholar.org/paper/Two-Approaches-to-Dialectical-Argumentation%3A-Sets-Verheij/dae1e05a96a0b76d6f932f2e8f60e1f3686c1f71> (May 15, 2025).
- Yu, Jay, Kevin McCluskey, and Saikat Mukherjee. 2020. 'Tax Knowledge Graph for a Smarter and More Personalized TurboTax'. doi:10.48550/arXiv.2009.06103.
- Zeleznikow, John, Andrew Stranieri, and Mark Gawler. 1995. 'Project Report: Split-Up — A Legal Expert System Which Determines Property Division upon Divorce'. *Artificial Intelligence and Law* 3(4): 267–75. doi:10.1007/BF00871852.
- Zurek, Tomasz, and Michał Araszkiewicz. 2013. 'Modeling Teleological Interpretation'. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Law*, ICAIL '13, New York, NY, USA: Association for Computing Machinery, 160–68. doi:10.1145/2514601.2514619.

EDUCATION

London School of Economics and Political Science, LSE Law School

London, UK

PhD Research: "A semantic framework for personal income tax regime: A computational legal analysis."

Sep'24 - Sep'28

Researching how neuro-symbolic methods can be combined to model and process law in highly regulated domains such as tax law. Supervised by Andy Summers and Eduardo Baistrocchi (both LSE Law School) and Milan Vojnovic (LSE Department of Statistics).

University of Cambridge, Faculty of Law

Cambridge, UK

Master of Corporate Law • First class honours • Clare Hall College Titular Scholar

Oct'22 - Jul'23

Deals Course (legal and economic structuring of corporate transactions) • Corporate Finance Law • Corporate law and the Digital Economy • Corporate Taxation (comparative analysis of UK, US, German tax regimes) • Shareholder Activism • Law Firm as a Business

Activities and societies: Cambridge Blockchain Society (Committee Member) • Cambridge European Society (Committee Member) • Cambridge University Triathlon Club (Varsity duathlon, Cambridge half-marathon) • Clare Hall College Boat Club (W1 rowing crew member)

Bucerius Law School

Hamburg, Germany

Legal Technology and Operations (Summer School)

Jul'23

University of Ljubljana

Ljubljana, Slovenia

Bachelor of Laws (LL.B.)

Oct'17 - Jan'22

GPA: 91.1/10 • top 5% • Dean's award for outstanding academic success • Dean's award for extracurricular activities (Willem C. Vis Moot).

WORK EXPERIENCE / RESPONSIBILITIES

[King's Entrepreneurship Lab](#)

Cambridge, UK

Senior Associate (promoted from Associate)

Oct'23 - present

Co-founded [Weekly Workshops with Brunches](#) and [Large Language Models \(LLM\) x Law Hackathon](#).

[AixLaw Conference](#) (two-day interdisciplinary conference with 30+ speakers, attracting nearly 200 attendees)

Ljubljana, Slovenia

Organising Committee

Oct'23 - present

Shaped programme, secured funding, coordinated speakers, oversaw marketing, presented, arranged publications, post-event mentoring.

[EY People Advisory Services \(Technology Centre of Excellence\)](#)

London, UK

Project manager (promoted from Intern)

Jul'21 - Aug'24

Oversaw application of generative AI to law and policies in the immigration field. Served as a bridge between users, business and engineers. Focused on legal data preprocessing, legal knowledge representation, system design, interface, user experience, feedback implementation.

Awards: CESA Ideation Lab 2.0 winner (2022; among 136 contributions from 29 countries, shortlisted by innovation community, country managing partners, regional leadership); EMEA #OnePAS Recognition Award (2023; ~100 countries, awarded to one employee per month).

[Cambridge Center for Finance, Technology and Regulation](#)

Cambridge, UK

Research Affiliate

Jul'23 - present

Exploring the application of network theory combined with sentiment analysis to understand community dynamics.

[University of Ljubljana](#) x [EUTOPIA European University](#) x [EY Slovenia](#)

Ljubljana, Slovenia

Project Lead (promoted from Research Assistant)

Jul'21 - Oct'23

Initiated a project filtering out relevant legal information (migration, social security, tax) for cross-border working researchers. Led the project to grow and cover 10 jurisdictions with a team of 9 people. Successfully secured Horizon Europe funding (€123K).

[Blockchain Think Tank Slovenia](#)

Ljubljana, Slovenia

Paralegal

Nov'21 - Oct'22

Worked closely with developers to analyse and advise on legal issues, solutions, and policy trends emerging from Web3 technologies.

[Faculty of Mathematics and Physics, University of Ljubljana](#) x [Abelium LL.C.](#)

Ljubljana, Slovenia

Research Assistant

Feb'20 - Aug'20

Worked closely with mathematicians, physicists and computer scientists to structure regulatory data in the AML/CFT field.

[Attorney Miran Hude](#)

Kranj, Slovenia

Paralegal

Oct'17 - Oct'22

Learnt a lot about legal practice by assisting at a family law firm on an ad hoc basis throughout my undergraduate studies.

PUBLICATIONS, LECTURES, CONFERENCES

- Hude, Z.; Paulik, T.: "LLMs x Law: Applications, Challenges, Opportunities", at Cambridge Immersion Programme: AI and the Law, directed by Professor Felix Steffek, organised by Møller Institute, Churchill College, Cambridge (2024).
- Hude, Z.; Igljar, M., Sanya-Mondoh B.: "DAOs: Introducing a New Era of Governance", in Ljubljana Law Review, 83 (2023), pp. 125-149
- Hude, Z.: "Unilateral Termination of Employment of Pensionable Workers" in: Ljubljana Law Review, 82 (2022), pp. 381-391.
- Hude, Z.: "Economic Analysis of the Proposed Additional 6th Tax Bracket in Slovenian Personal Income Tax Scale", in: DENAR Magazine, professional journal on finance and taxation, Vol. 515 (2020)