

Deploying AI in Uncertain Environments: A Technical Limitation or a Human Characteristic?

Marc T.J Elliott

School of Electronics, Electrical Engineering and Computer Science
The Centre for Public Policy and Administration
Queen's University Belfast, United Kingdom
melliott22@qub.ac.uk

Introduction

Much of the recent literature in the wider area of AI ethics focuses on mitigating bias, ensuring transparency, or establishing accountability. Yet a foundational concept often underexamined is uncertainty: how is it conceptualised, modelled, and ultimately addressed in AI systems? From a computer science perspective, uncertainty is a technical condition, such as noisy data, where uncertainties exist that cause deviation from the intended actions and is thus treated as a quantifiable and solvable challenge. In contrast, social scientists and legal scholars frequently regard uncertainty as a constitutive feature of complex human systems, something to be interpreted, negotiated, or contested.

This divergence matters, AI systems are increasingly integrated into public decision-making domains, from healthcare to welfare, important decisions are being made with the help of algorithmically driven technologies. These environments demand not just accurate outputs, but interpretative flexibility, context sensitivity, and responsiveness to ambiguity. Designing AI systems for such domains requires moving beyond statistical notions of uncertainty and toward a richer account of its social and institutional roles.

My research therefore asks: *What does it mean for AI systems to engage with uncertainty in environments where ambiguity is not a flaw to be corrected, but a condition that requires meaningful navigation?* And: *Where and how can AI systems be designed to support, rather than displace, human interpretative capacities in such settings?*

Theoretical Foundations

My research is grounded in an interdisciplinary methodology, bringing together perspectives from human-computer interaction, psychology insights, and more recently drawing on legal theory. A central conceptual contribution of my research is distinction between the idea of *technical uncertainty* (e.g., incomplete or noisy data, applications where all

possible outcomes cannot be mapped, and model unpredictability) and *social uncertainty* (e.g., interpretative situations, ambiguity, pluralistic norms, and discretionary authority). While the former is often managed and treated through probabilistic reasoning and data handling techniques to reduce the impact, the latter tends to resist such simple resolutions.

My research draws on the work of Gerd Gigerenzer (2022, 2023), incorporating insights from the psychology domain to understand how society can be adapted to produce more effective AI that limits the social uncertainties impacting these systems. Specifically, my research expands on their work to further explore whether we should make adaptations that impact society, or if placing substantial time and resources to meaningfully integrate these systems in public settings risks eroding what makes society an inherently human environment. Furthermore, by utilising the legal theories of formalism and realism as part of the case study, I aim to show that uncertainty in human decision-making is not simply a lack of available information, extending beyond the common misleading proclamation that AI systems can always be made more effective with bigger and better data.

It is often a reflection of the competing values, irreducible ambiguity or systemic constraints that defy algorithmic modelling. For example, legal systems intentionally preserve a degree of vagueness to accommodate changing social norms and contextual interpretation. When AI systems are then deployed in such settings, such as the recent cases of *Snell v. United Specialty* (2024) and *United States v. Ross* (2025), which used LLMs to interpret ordinary meanings in legal contexts, it risks over-committing to formalistic, narrow, or rigid readings of rules (Posner & Saran 2025), effectively distorting the practices they are meant to augment.

Prior Work: The Stability Dilemma

In one line of my research, I recently examined this issue through what I call the *stability dilemma*: the tension between the public sector AI's need for stable, repeatable decisions and the contextual variability that real-world public applications and governance requires (Elliott & P. 2025).

This work established a new framework that builds upon Gigerenzer’s (2022) adapt-to-AI principle, providing a unique method for identifying uncertainty factors which exist for the use of AI across public sector applications. This examines public sector uncertainties across four dimensions: the intended system user, the primary system subject, the physical environment the system exists within, and the social environment in which the system is embedded.

My study finds that the greater challenge lies in determining what uncertainty factors should we adapt, or whether we should even make adaptations at all, in order to preserve society as an environment that should permit variation in interpretation and to positively view uncertainty as a mechanism that promotes human judgement.

I argue that current approaches to AI ethics, whether through fairness, bias mitigation, or transparency mechanisms, while important in the journey toward responsible AI still often fail to address this deeper tension. They primarily operate within the logic of technical uncertainty, assuming that better training data or algorithmic adjustments can resolve the ‘*problem*’. What is missing is a recognition that some uncertainties are, in fact, not only necessary, but desirable, especially within democratic institutions.

Uncertainty in LLMs for Legal Interpretation

The second research stream focuses on how LLMs are handling legal uncertainty. Following on from the conceptual framework established in *The Stability Dilemma*, this grounds my work within a legal scenario. Through experimental prompting of OpenAI’s GPT-4 across real-world legal cases where AI has previously played a role, I test their responses under conditions of interpretative ambiguity.

For example, in *Snell v. United Specialty* (2024) and *United States v. Ross* (2025), I analyse how LLMs respond when given different prompt structures that vary in legal framing and factual emphasis. A core contribution of this study works toward measuring the semantic similarity between LLM outputs to the case study responses, allowing for a clearer insight into the potential dangers of using LLMs for legal interpretative tasks. These specific cases were selected as they involve highly interpretative scenarios, in the case of *Snell v. United Specialty* (2024), Judge Newsom—out of frustration with the answers provided via traditional legal instruments—turned to ChatGPT to seek an answer on whether installing an in-ground trampoline might be considered as a form of landscaping. While this information ultimately did not play a role in the court’s outcome, it presents a new way in which the legal system is beginning to incorporate AI systems, and I believe this requires a detailed investigation into what circumstances might these applications enhance the legal system or whether it hinders the progressive evolution of law.

The preliminary results revealed that LLMs are often defaulting towards responses that have a greater alignment with legal formalism, favouring surface-level rule application over context-sensitive analysis. Posner and Saran (2025) identified this relationship between LLMs and legal formalists in their study. They demonstrated that the quality of LLM legal reasoning was closer to that of an undergraduate law student than that of a judge, whereby a judge would be more likely to consider the wider contextual significances and incorporate these into their decision-making even if it somewhat departs from the written rules of law.

While the LLMs interpretative responses to the questions posed in these case studies were initially impressive in fluency and reasonings for their interpretations, the models struggle to manage interpretative plurality. Furthermore, while infrequent, I found outputs diverge in answer when posed the same prompt numerous times and response quality varied based on prompt framing. These preliminary results suggest that LLMs lack stable internal reasoning when faced with tasks of legal uncertainty and simultaneously take a stance that aligns with a junior legal formalist, opposed to the realist perspective often taken by real-world judges.

Future Work

Going forward, I have two primary directions in which I intend to further contribute to the AI and society uncertainty discourse. First, following my current line of research that brings my computer science expertise to the legal domain, I aim to better formalise the concept of *legal entropy* as a tool for evaluating whether AI systems are suitable for particular legal tasks, and to extend this work to other domains such as education policy and environmental regulation, where somewhat similar tensions exist.

Secondly, a broader goal is to offer guidance on where to draw the line between appropriate adaptation—where AI systems absorb and mirror human judgment—and overreach, where they impose unwarranted certainty on domains that are inherently ambiguous. This boundary, I argue, cannot be determined by technical performance metrics alone; it must be informed by normative evaluation rooted in domain expertise, institutional purpose, and democratic accountability. I envision the construction of a framework that provides stakeholders with a general guide to establish whether adaptations to the public environment should or should not be made, and if they are, to require justification that demonstrates why this is a meaningful adaptation to make that does not risk eroding the humanness of society.

As public institutions move toward an “adapt first, react later” approach to AI, my research calls for a deliberate reckoning with uncertainty, not as a defect to be fixed, but as a signal of what makes human governance irreplaceable.

References

- Elliott, M. and D. P. 2025. The Stability Dilemma: Addressing Uncertainty in Public Sector AI. In In Proceedings of the IEEE International Symposium on Ethics in Engineering, Science, and Technology. Illinois: Institute of Electrical and Electronics Engineers. doi.org/0.1109/ETH-ICS65148.2025.11098414
- Gigerenzer, G. 2022. *How to Stay Smart in a Smart World: Why Human Intelligence Still Beats Algorithms*. Massachusetts MIT Press
- Gigerenzer, G. 2023. Psychological AI: Designing Algorithms Informed by Human Psychology. *Perspectives on Psychological Science* 19(5), 839-848. doi.org/10.1177/17456916231180597.
- Posner, A. and Saran, S. 2025. Judge AI: Assessing Large Language Models in Judicial Decision-Making. *University of Chicago Law School, Coase-Sandor Institute for Law & Economics Research Paper* 25(3). doi.org/10.2139/ssrn.5098708.
- Snell v. United Specialty Insurance Company. 2024. No. 22-12581. 11th Circuit, May 28, 2024. <https://media.ca11.uscourts.gov/opinions/pub/files/202212581.pdf>. Accessed: 2025-06-06
- United States v. Ross, 2025. No. 24-30761. 5th Circuit, June 3, 2025. <https://law.justia.com/cases/federal/appellate-courts/ca5/24-30761/24-30761-2025-06-03.html>. Accessed: 2025-07-16