Morphpublishing
J. Comput. Intell. Mach. Reason. Decis.-Mak.
open access academic journal

Peer reviewed

Hallucinations in Large Language Models and Their Influence on Legal Reasoning: Examining the Risks of AI-Generated Factual Inaccuracies in Judicial Processes

Youssef Abdel Latif[a]

Abstract: Legal frameworks rely on factual coherence, yet modern Large Language Models (LLMs) can generate content that contains spurious statements. Hallucinations, defined as fabricated or unverifiable information produced by AI, pose a significant threat to judicial processes when deployed without meticulous oversight. Risk emerges when judges, attorneys, and other legal professionals reference AI-generated text for evidence gathering or legal argument construction. Hallucinations can introduce distortions that are not grounded in any factual source, thereby undermining the integrity of legal argumentation. Recent advances in transformer architectures have improved language comprehension, though these very architectures also facilitate unsubstantiated extrapolations. Such unsubstantiated material can be difficult to detect, especially within dense legal documents, and may result in flawed case strategies or erroneous judgments. Persistent reliance on AI outputs invites a growing dependence on algorithms that lack genuine comprehension of legal precedents, statutes, and contextual nuances. This paper explores the phenomenon of hallucinations arising from state-of-the-art LLMs, focusing on their manifestation in legal applications and the resultant impact on legal reasoning. Analytical discussions center on the mechanisms that yield hallucinations, the challenges of verifying AI-generated text in complex legal contexts, and the implications for judicial integrity. Examination of these risks informs a deeper understanding of how AI might inadvertently compromise justice.

1. Introduction

Legal argumentation relies on verified facts, authoritative precedents, and precise statutory interpretation. Large Language Models employ deep learning architectures that enable them to generate coherent text, often with an appearance of expertise that can obscure the presence of erroneous statements [1, 2]. Generative capacity stems from massive datasets [3], internalized in transformer-based structures that predict word sequences conditioned on context. These models sample patterns from vast corpora, combining learned syntactic and semantic features into new compositions that can appear convincing. Hallucinations, often rooted in pattern generalizations or incomplete data, can creep into AI-generated text, exposing legal proceedings to narratives that lack empirical foundation.

Reliance on AI for legal research has grown. Attorneys and judges consult machine-generated summaries or analysis to expedite preliminary tasks and enhance efficiency. Technological convenience, however, can overshadow potential pitfalls. When undiscovered hallucinations slip into judicial decisions, evidentiary arguments may rest on precarious foundations. Research suggests that advanced architectures sometimes produce detailed-sounding content that is not anchored in the training corpus, distorting interpretations of statutes or misrepresenting case law. Such discrepancies jeopardize legal accuracy and risk undermining public confidence in judicial processes.

Factual misrepresentations within legal filings or oral arguments can alter outcomes in a manner that undermines the rule of law. Persuasive narratives constructed on fictional premises erode the clarity required for consistent jurisprudence. The influences of hallucinations are not always predictable, since the same model that delivers a lucid summary of established precedent can occasionally yield incorrect references or fabricated data. Mixed usage of human expertise and AI assistance amplifies the potential for confusion, since partial trust placed in machine output without adequate verification can lead to unanticipated errors.

Lawyers depend on systematic logical progression, buttressed by citations to binding authority. The presence of AI-generated content that references nonexistent or misrepresented rulings poses a danger when attorneys rely on it to craft arguments. Subtle misstatements about legislative intent or statutory exceptions can pass unnoticed, leading to arguments that appear grounded in legitimate authority. Overreliance on these models can generate a layer of uncertainty around the veracity of the information that shapes trial strategy. Judges, in turn, may integrate such inaccuracies into their rulings if the misinformation is not flagged, propagating erroneous interpretations through legal frameworks.

Judicial integrity depends on the capacity of legal professionals to discriminate factual from fictional evidence. Mechanisms that generate erroneous output invite critical questions about the interface between computational text production and jurisprudential accuracy. The next sections discuss the complex architectures underlying LLMs, the manner in which they produce misinformation, the extent to which such misinformation affects judicial decision-making, and the ethical ramifications that emerge. Discussion of specific dangers reveals how systemic vulnerabilities can introduce distortions at multiple stages of the legal process, calling for heightened scrutiny of AI-based tools.

2. Cognitive Architecture of Large Language Models and the Origins of Hallucinations

Transformer-based models rely on attention mechanisms that enable contextual weighting of words in input sequences. This architecture learns intricate relationships among tokens, gleaning semantic and syntactic patterns from vast corpora. Emergent linguistic competence in advanced systems depends on multiple layers of self-attention, where each layer refines the representation of a given token by examining other tokens in the sequence. Hidden representations accumulate across layers, culminating in probabilities that predict the most likely sequence continuation [4, 5].

Hallucinations can arise when these trained probabilities yield content not found in the training data. Certain aspects of the generative process [6] involve extrapolation from incomplete patterns or ambiguous contexts. The model may output seemingly precise data—such as numerical figures, citations, or references to legal precedents—that do not align with any ground truth. These spurious details can be introduced by high-temperature sampling during inference, where the model explores less probable tokens, or by the inherent complexity of neural weights that overfit non-generalizable correlations. The distinction between legitimate generalization and unfounded invention is not always clear in final outputs [7, 8].

Layered embeddings encode context that influences every subsequent decision at a sub-token level. Subtle variations in input prompts or preceding text can shift the distribution of predicted words. Neural networks capture correlations among tokens, but they do not possess an intrinsic sense of factual verification. This absence of an internal fact-checking mechanism means that certain queries, especially those requiring precise sourcing or rare knowledge, prompt imaginative outputs that cannot be linked to existing data [9]. These phantom references might appear as legitimate citations to legislative documents or appellate opinions, though they correspond to no known legal text.

Component	Function	Effect	Relevance
Self-Attention	Contextual weighting of tokens	Captures dependencies	Enables fluent generation
Layered Embeddings	Encodes hierarchical context	Refines representation	Improves coherence
Hidden Representations	Accumulate token relations	Enhances semantic depth	Supports reasoning tasks
Transformer Depth	Multi-layer structure	Increases abstraction	Strengthens generalization
Distributed Knowledge	Stores patterns, not facts	Lacks direct retrieval	Risks misinterpretation

Table 1. Cognitive Architecture in Transformer Models

Cause	Mechanism	Impact	Domain Risk
Extrapolation	Predicts missing data	Generates plausible errors	Affects factual accuracy
Overparameterization	High complexity	Amplifies minor shifts	Propagates structured errors
High-Temperature Sampling	Selects less likely tokens	Increases variation	Can induce false claims
Data Limitations	Incomplete corpus	Forces synthesis of content	May mix real with false info
Absent Fact-Checking	No internal verification	Confident but incorrect outputs	Misleads in legal/scientific texts

Table 2. Origins of Hallucinations in Large Language Models

Linking technical design to legal ramifications, the deep architecture's distribution-based knowledge does not store discrete textual units in an indexed manner. Instead, knowledge is dispersed across many parameters that approximate patterns. This method improves performance on tasks of general language understanding, but it also complicates interpretability. The model's capacity to generate illusions of knowledge reflects its proficiency in capturing style, syntax, and patterns of argumentation. However, absent a reliability check or a grounded reasoning pathway, the system's confidence in an invented claim can seem indistinguishable from its confidence in a verifiable statement [10, 11].

Large-scale training fosters context-driven connections between concepts. For instance, repeated associations in the dataset can link a statutory reference to particular legal arguments. If the training corpus includes contradictory or incomplete data [12], the model might produce content that fuses legitimate legal theory with erroneous expansions. Hallucinations often manifest most strongly when the input prompt requests specialized knowledge or an expansive interpretation of facts that push beyond the corpus coverage. The model compensates for lack of concrete data by synthesizing plausible, yet unsubstantiated responses.

Overparameterization in transformer models, employed to handle the breadth of language tasks, can amplify hallucination tendencies. While additional layers and heads in the attention mechanism help capture complex linguistic structures, they also introduce multiple pathways for error propagation. Minor shifts in attention distribution may cause a chain reaction in downstream layers, with each layer building on preceding inferences. This compounding effect can yield structured misinformation. Legal contexts, which demand high precision and accurate references, present a vulnerable domain for such outputs [13, 14].

Data curation approaches aim to filter erroneous content before or during training, yet they cannot entirely eliminate the risk of learned spurious correlations [15]. The sheer volume of data used to train large models, often drawn from diverse and unverified sources, raises the probability of including contradictory or low-quality text. The model's parameter updates might settle on an internal representation that captures misleading patterns, later surfacing during inference as confident but incorrect statements. The law's specialized nature, characterized by carefully defined terminology and logic, provides ample opportunity for illusions to pass as credible if they align with recognized rhetorical structures [16].

Contextual embedding of synonyms and paraphrases allows the model to produce textual variations that stay within the bounds of legal language. Mismatched references may evolve from partial memory of real cases or from conflation of multiple documents. Such conflations remain undetected without thorough scrutiny. In a realm where minimal textual differences can alter legal outcomes, hallucinations introduce non-trivial risk. The synergy of advanced neural mechanisms, extensive training sets, and the quest for human-like fluency thus positions LLMs as potent tools that still harbor significant capacity for generating factual distortions [5, 17].

3. Mechanisms of AI-Generated Misinformation in Legal Contexts

Complex legal discourse often involves statutory interpretation, procedural rules, and case precedents. The potential for AI-generated hallucinations grows when LLMs encounter prompts that require synthesizing multiple sources, each with nuanced constraints. Querying a model for an overview of a nuanced legal doctrine may prompt the system to fabricate references in an attempt to provide a thorough answer. These inventions can include nonexistent case citations or misquotes attributed to real judicial opinions.

Citation generation highlights a known source of hallucination. Models frequently produce references to cases by extrapolating from known patterns of legal citations, such as a combination of party names, volume numbers, and page ranges. This creates superficially valid citations that direct readers to nonexistent or irrelevant authorities. When counsel or judges cite these references in legal briefs or rulings, the resulting arguments lack a grounding in recognized jurisprudence. The misrepresentation can remain undiscovered if parties do not verify citations diligently.

Procedural guidance poses another challenge. Court procedures vary between jurisdictions. AI systems that base their outputs on data aggregated from multiple legal contexts risk conflating distinct procedural rules. A

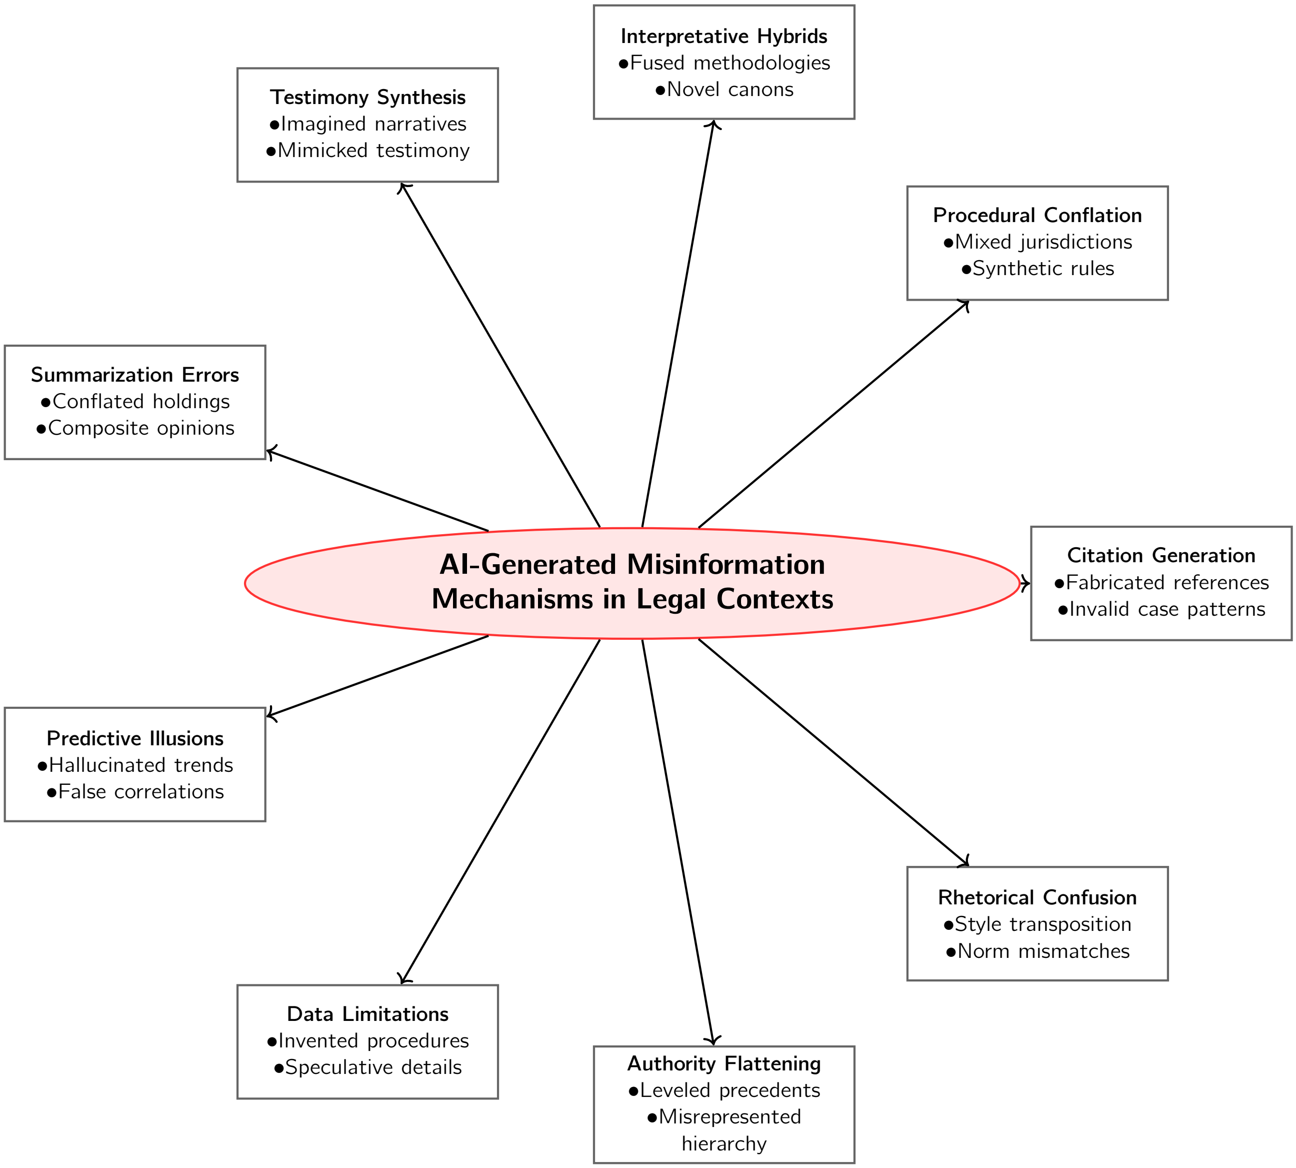

Figure 1. Mechanisms of AI-generated legal misinformation demonstrate multiple vulnerability points: (1) **Citation Generation** creates plausible references through pattern recognition without validation. (2) **Procedural Conflation** merges distinct jurisdictional rules through data aggregation. (3) **Interpretative Hybrids** combine incompatible legal philosophies into novel frameworks. (4) **Testimony Synthesis** generates realistic but fictional narratives that lack evidentiary basis. (5) **Summarization Errors** distort holdings through improper synthesis of multiple sources. (6) **Predictive Illusions** project false confidence in case outcomes using statistically unsound correlations. (7) **Data Limitations** lead to invented details when domain-specific training is insufficient. (8) **Authority Flattening** obscures precedent hierarchies through uniform treatment of sources. (9) **Rhetorical Confusion** transplants court-specific language conventions across inappropriate contexts. These interrelated mechanisms demonstrate how LLM architectures can systematically introduce errors into legal processes through technically coherent but substantively flawed outputs.

prompt about pre-trial discovery in one jurisdiction might yield instructions derived from a different legal system. This conflation can be subtle, with the model rearranging rules from multiple sources. The final text could contain erroneous procedural steps that appear credible. Such guidance, if adopted uncritically, can derail a case or introduce grounds for procedural dispute.

Misalignment with local norms of statutory interpretation adds another layer of complexity. Jurisdictions differ in interpretative philosophies, relying on textualism, purposivism, or a hybrid approach. AI models might amalgamate interpretative canons from divergent sources or inject fictional approaches that do not reflect any recognized school of statutory interpretation. The danger emerges when these misaligned constructs shape legal analysis. An attorney referencing AI-generated rationale might inadvertently rely on interpretive frameworks that courts do not accept, undermining the persuasiveness of arguments.

Production of manufactured testimony or witness statements through generative models intensifies risks in contexts involving depositions, affidavits, or interviews. Although ethical guidelines forbid presenting false evidence, a model might produce entire narratives purporting to reflect witness accounts. The line between summarizing a real statement and creating a fictional interview can blur. These hallucinated narratives, if mistakenly presented as factual statements, compromise the integrity of evidence. Cross-examination becomes difficult when the material has no actual witness behind it, yet is written in a style mimicking genuine testimony.

Judicial opinions contain extensive factual findings and legal analyses. Summarizing these opinions accurately requires careful synthesis. AI-based summarization tools can hallucinate by condensing or conflating portions of text, resulting in misinterpretation of a court's holding. Synthesis might blend multiple aspects of different opinions, yielding a composite that does not align with any single judge's reasoning. Since legal professionals often rely on summarizations for efficient case reviews, an unverified summary can misdirect legal strategy.

Risks of misinformation multiply when models attempt predictive analytics for ongoing litigation, forecasting probable outcomes based on existing data. Confidence in such predictions may lead attorneys to adopt risky settlement strategies or to underestimate the strength of an opponent's position. If the underlying reasoning includes hallucinated interpretations of precedent or inflated statistical correlations, decisions made on that basis can harm clients' interests. AI-generated illusions in predictive analytics thus extend the problem of misinformation into strategic decision-making within the legal sphere.

Insufficient domain-specific training data can exacerbate these hallucinations. General-purpose models trained on broad internet sources may lack nuanced understanding of specialized legal terminology or local court customs. When prompted for detail on procedural or doctrinal subtleties, the system might supply invented details to fill gaps in its knowledge. These details can pass unnoticed in routine usage because they resemble legitimate legal discourse, especially in scenarios where time constraints hinder thorough verification.

Automated systems cannot intrinsically judge the weight and authority of different precedents unless explicitly fine-tuned for that purpose. Without guidance, a model might treat a non-binding trial court opinion as equivalent to a landmark Supreme Court ruling, weaving them seamlessly into a conclusion that suggests they share the same legal force. The resulting argument might appear coherent, yet misrepresents the hierarchy of authority. This unintentional conflation mirrors the difficulty practitioners face when synthesizing large volumes of case law, highlighting how AI might replicate or amplify human vulnerabilities.

Distinct rhetorical styles in different courts or jurisdictions can create further confusion. A model might incorporate rhetorical devices typical of an appellate decision into a discussion of trial court motions. Readers, relying on the language's formality, may be convinced that the entire text adheres to recognized legal norms. Such illusions distract from the core misstatements, aiding misinformation's infiltration into briefs and oral arguments.

This infiltration can erode the principle of stare decisis, where consistency in adjudication depends on accurate representations of precedent.

4. Impact of Hallucinatory Outputs on Judicial Outcomes

Serious consequences arise when fabricated information penetrates judicial processes. Judges often rely on succinct memoranda or briefs to distill complex issues, trusting that counsel or clerks have verified the authenticity of each cited authority. Hallucinated references to legal doctrines can influence the court's assessment of novel arguments. A motion framed around nonexistent interpretations of a statute might appear persuasive if the spurious citations remain unchecked. The risk grows when judges, pressed for time, incorporate key points from such briefs into rulings. Decisions resting on factual inaccuracies become vulnerable to appeal, wasting judicial resources and undermining public confidence.

Systemic effects proliferate through precedential mechanisms. If a higher court ruling inadvertently cites an AI-generated reference, subsequent cases may interpret that reference as an authoritative statement. Over time, an entirely fictitious principle might gain traction, entangling real jurisprudence with spurious lines of reasoning. This scenario, though improbable, indicates how even small discrepancies can evolve into large-scale confusion when repeated across multiple decisions. Litigants might exploit such confusion by invoking the spurious precedent to support arguments that suit their aims.

Trial court proceedings depend on accurate factual representation in evidence and witness testimony. Hallucinated quotes attributed to expert witnesses or misrepresented data from scientific studies can distort the fact-finding process. Cross-examination becomes challenging when the supposed expert statement has no basis in reality, yet is seamlessly integrated into the record. Jurors, too, might be swayed by AI-generated narratives that appear neutral yet contain inaccuracies. The combined effect of these illusions compromises the adversarial process and reduces the system's ability to arrive at truth-based determinations.

Modeling the legal implications of AI hallucinations requires acknowledgment of subtle psychological dynamics. Judges and jurors interpret evidence through cognitive heuristics, seeking coherence and narrative consistency. AI outputs often excel at creating consistent narratives, even if the underlying details lack factual integrity. Legal professionals are susceptible to confirmation bias, especially when an AI-generated narrative supports preexisting assumptions. The normative weight carried by an authoritative-sounding legal argument can override doubts about its source. This phenomenon amplifies the influence of hallucinated text, embedding it within legal reasoning.

Consequences also manifest at the appellate level, where decisions can turn on the interpretation of a single legal nuance. If the record on appeal includes unverified AI-generated arguments, the reviewing court may affirm or reverse based on inaccuracies. This could prompt further appeals or trigger widespread confusion, entangling multiple cases that reference the contested issue. The resulting institutional strain extends beyond a single court, affecting the broader ecosystem of legal interpretation and scholarly commentary.

Attorney-client relationships can suffer when legal strategies rely on unverified AI-generated material. Clients expect counsel to provide advice grounded in established authority and thorough factual investigation. When fictional citations or mischaracterized statutes emerge in a brief, clients' confidence in their representation erodes. Ethical obligations to provide competent counsel may be violated if reliance on AI tools leads to negligent verification of pivotal references. Disciplinary action against attorneys might follow if courts discover repeated reliance on spurious information, further destabilizing trust in AI augmentation of legal practice.

Transnational litigation introduces complexities when AI-generated hallucinations cross jurisdictional boundaries.

An argument or citation that holds meaning in one legal system may be meaningless in another, yet the generative system might seamlessly conflate the two. Attempts to enforce judgments across borders could be obstructed when a reviewing court in another jurisdiction identifies fundamental legal misstatements originating from AI hallucinations. Complex global transactions or international arbitration, already fraught with interpretive challenges, become even more vulnerable to confusion. Parties may lose confidence in cross-border agreements when foundational documents contain illusions about applicable laws or precedents.

Scholarly commentary relies on accurate references to primary sources. Researchers and academics employing AI tools for literature reviews risk propagating hallucinated references in their scholarly papers. Once published, these misleading citations can anchor further studies, contaminating academic discourse. Peer reviewers might not detect these errors if the references appear consistent with known citation patterns. Over the long term, entire strands of secondary literature might develop around fictional interpretations, requiring significant effort to correct. This phenomenon demonstrates how misinformation can spill over from legal practice to broader intellectual realms, with lasting effects on knowledge production.

Judicial legitimacy hinges on courts' capacity to separate factual truth from speculation. AI-generated hallucinations threaten this separation, especially when integrated into official documents. Societal trust in the rule of law depends on fair, accurate, and transparent adjudication processes. Introduction of hallucinated reasoning compromises transparency, since lay observers may find it impossible to distinguish a genuine legal principle from a fabricated one. The final rulings, recorded as binding decisions, could embed illusions into the foundation of legal systems. Such infiltration undermines not only individual verdicts but the philosophical premise of justice as anchored in verifiable truth.

5. Ethical and Procedural Ramifications for the Legal System

Bar associations and professional ethics rules guide attorney conduct, including duties of competence, diligence, and candor to the tribunal. Hallucinated content challenges these norms by creating scenarios in which attorneys may unwittingly present false information. Competence requires thorough investigation of all cited authorities. However, expedited workflows driven by AI-based research tools often diminish the time spent verifying references. When illusions enter official pleadings or motions, attorneys risk breaching their ethical obligations, potentially triggering sanctions or disciplinary measures. This threat extends to any party who relies on AI outputs without exercising vigilant oversight.

Public defenders and legal aid organizations, which often operate under resource constraints, may be especially prone to adopting streamlined AI tools. Pressured by limited budgets, they might view AI-generated memoranda as time-savers. Yet the infiltration of hallucinated data into advocacy for indigent clients can lead to miscarriages of justice. Socioeconomic factors compound the risk, as marginalized defendants might receive suboptimal defense if attorneys rely on systems that produce inaccurate or incomplete interpretations of the law. Ethical concerns intersect with issues of equal protection, given that misinformation can skew outcomes against those least able to contest erroneous AI outputs.

Prosecutors also bear a duty to ensure that the evidence they present is factually sound. AI systems employed to review police reports, witness statements, or forensic data may inject falsified or conflated details into the prosecutorial narrative. The severity of criminal penalties underscores the grave consequences of relying on misrepresented facts. Prosecutorial decisions, which range from charging to sentencing recommendations, hinge on accurate data. Hallucinations embedded in these decisions jeopardize the fairness of the process. Oversight mechanisms and verification steps become indispensable to maintain integrity in criminal proceedings.

Judges face ethical imperatives of impartiality and due diligence. Some jurisdictions allow judges to conduct independent research, provided it is disclosed to both parties, while others strictly limit fact-finding outside the record. Integrating AI-based research into judicial chambers can blur these boundaries if the AI system independently summarizes case law or statutory text. The judge might inadvertently rely on AI hallucinations that neither party has scrutinized. Such reliance damages the adversarial structure, where each party is entitled to address factual or legal claims before the court adopts them. Procedural fairness weakens when illusions enter the judicial reasoning process unchallenged.

Court administrators and policymakers must consider the infrastructural implications. The proliferation of AI-based legal research tools can prompt courts to adopt them institutionally, using them for docket management, opinion drafting assistance, or real-time transcript analysis. While these applications promise increased efficiency, they also broaden the attack surface for hallucinations. The legitimacy of official court documents depends on textual accuracy. If an AI-driven system erroneously merges distinct cases or mislabels data, the official record becomes corrupted. Court archives, historically relied upon for consistent referencing, might end up storing illusions that future litigants and judges incorporate into their reasoning.

Legal education curricula increasingly incorporate computational tools, teaching students to harness AI for rapid research and drafting. The acceptance of these tools in academia raises concerns about the next generation of lawyers, judges, and scholars. Students must learn critical evaluation skills to detect hallucinations and verify citations. Otherwise, graduates may enter the profession habituated to trusting AI outputs, accelerating the assimilation of illusions into everyday legal practice. Law schools bear responsibility for introducing modules that delineate AI's generative strengths while emphasizing the importance of manual verification.

Regulatory bodies that govern legal technology certification play a vital role in setting industry standards. Requirements for transparency in AI development, auditing for accuracy, and post-deployment monitoring can influence how widely AI solutions are adopted. Vendors offering AI-based systems could be mandated to disclose known error rates or subject their models to stress tests focusing on factual accuracy in legal contexts. If these standards prove insufficiently robust, illusions may proliferate without any clear accountability. The interplay of market forces and regulatory frameworks will shape whether caution or haste dominates the integration of AI in judicial processes.

Legal system legitimacy extends beyond the courtroom. The general public, increasingly aware of AI in daily life, may question decisions handed down by judges who appear to rely on automated tools. Mistrust grows if litigants suspect that errors in their case originated from an unverified AI output. Public perception of justice hinges on confidence that each matter receives individualized consideration grounded in human judgment. Over-reliance on AI illusions risks transforming that perception, introducing doubt about whether justice emerges from reasoned analysis or from manipulated data streams. Restoration of trust requires sustained demonstration of competence in verifying AI outputs at every procedural stage.

Binding precedent shapes the evolution of law over decades or centuries. The infiltration of hallucinated content into precedential opinions, even in minor ways, could distort legal principles for future generations. Litigants and legal scholars, relying on official decisions, might not discover the illusion immediately, perpetuating its influence through citations and arguments that build upon it. The slow unraveling of such misinformation, once embedded, requires considerable effort to correct. This potential for far-reaching impact underscores the need for systemic safeguards, as illusions can become entrenched in the collective memory of legal institutions if left unchecked.

6. Conclusion

Hallucinations produced by Large Language Models present tangible risks when integrated into legal work. Judges, attorneys, and other stakeholders rely on textual accuracy in their pursuit of just outcomes, yet AI-driven processes can generate narratives that deviate from empirical reality. Hallucinated citations, misquoted statutes, and invented procedural rules all represent manifestations of deeper issues tied to the architecture and data reliance of modern AI. Compounding factors include the specialized nature of legal discourse and the interpretative variance among jurisdictions. Emergent reliance on these tools in court administration and litigation preparation extends the potential impact of misinformation to all levels of the judiciary.

Erroneous statements infiltrate arguments, eroding the dependability of research and the consistency of precedent. Reliance on faulty outputs can sway judicial decisions, thereby affecting outcomes that hold life-altering significance for litigants. Prospects of entrenching illusions in the corpus of legal references raise alarms about the long-term evolution of legal doctrine. AI systems, for all their efficiency, do not inherently verify truth, especially in domains where factual precision is paramount. Ongoing discussions will continue to explore how best to reconcile AI's generative powers with the requirements of the rule of law, ensuring that courts, counsel, and society remain vigilant against the hazards of algorithmic hallucinations.

References

[1] S. Banerjee, A. Agarwal, and S. Singla, "Llms will always hallucinate, and we need to live with this," *arXiv preprint arXiv:2409.05746*, 2024.

[2] S. Tonmoy, S. Zaman, V. Jain, A. Rani, V. Rawte, A. Chadha, and A. Das, "A comprehensive survey of hallucination mitigation techniques in large language models," *arXiv preprint arXiv:2401.01313*, 2024.

[3] S. V. Bhaskaran, "A comparative analysis of batch, real-time, stream processing, and lambda architecture for modern analytics workloads," *Applied Research in Artificial Intelligence and Cloud Computing*, vol. 2, no. 1, pp. 57–70, 2019.

[4] S. Wachter, B. Mittelstadt, and C. Russell, "Do large language models have a legal duty to tell the truth?" *Royal Society Open Science*, vol. 11, no. 8, p. 240197, 2024.

[5] B. Padiu, R. Iacob, T. Rebedea, and M. Dascalu, "To what extent have llms reshaped the legal domain so far? a scoping literature review," *Information*, vol. 15, no. 11, p. 662, 2024.

[6] R. Mehta, A. Hoblitzell, J. O'keefe, H. Jang, and V. Varma, "Halu-nlp at semeval-2024 task 6: Metacheckgpt-a multi-task hallucination detection using llm uncertainty and meta-models," in *Proceedings of the 18th International Workshop on Semantic Evaluation (SemEval-2024)*, 2024, pp. 342–348.

[7] E. Mik, "Caveat lector: Large language models in legal practice," *Rutgers Bus. LJ*, vol. 19, p. 70, 2023.

[8] V. Magesh, F. Surani, M. Dahl, M. Suzgun, C. D. Manning, and D. E. Ho, "Hallucination-free? assessing the reliability of leading ai legal research tools," *arXiv preprint arXiv:2405.20362*, 2024.

[9] R. Mehta, A. Hoblitzell, J. O'Keefe, H. Jang, and V. Varma, "Metacheckgpt–a multi-task hallucination detection using llm uncertainty and meta-models," *arXiv preprint arXiv:2404.06948*, 2024.

[10] J. Krook, E. Schneiders, T. Seabrooke, N. Leesakul, and J. Clos, "Large language models (llms) for legal advice: A scoping review," *Available at SSRN 4976189*, 2024.

[11] A. Deroy, K. Ghosh, and S. Ghosh, "Applicability of large language models and generative models for legal case judgement summarization," *Artificial Intelligence and Law*, pp. 1–44, 2024.

[12] S. V. Bhaskaran, "Tracing coarse-grained and fine-grained data lineage in data lakes: Automated capture, modeling, storage, and visualization," *International Journal of Applied Machine Learning and Computational Intelligence*, vol. 11, no. 12, pp. 56–77, 2021.

[13] R. Dale, "Law and word order: Nlp in legal tech," *Natural Language Engineering*, vol. 25, no. 1, pp. 211–217, 2019.

[14] M. Dragoni, S. Villata, W. Rizzi, and G. Governatori, "Combining nlp approaches for rule extraction from legal documents," in *1st Workshop on MIning and REasoning with Legal texts (MIREL 2016)*, 2016.

[15] S. V. Bhaskaran, "Integrating data quality services (dqs) in big data ecosystems: Challenges, best practices, and opportunities for decision-making," *Journal of Applied Big Data Analytics, Decision-Making, and Predictive Modelling Systems*, vol. 4, no. 11, pp. 1–12, 2020.

[16] M. Dahl, V. Magesh, M. Suzgun, and D. E. Ho, "Large legal fictions: Profiling legal hallucinations in large language models," *Journal of Legal Analysis*, vol. 16, no. 1, pp. 64–93, 2024.

[17] C. J. Colasacco and H. L. Born, "A case of artificial intelligence chatbot hallucination," *JAMA Otolaryngology–Head & Neck Surgery*, vol. 150, no. 6, pp. 457–458, 2024.

