



**Stanford – Vienna
Transatlantic Technology Law Forum**

A joint initiative of
Stanford Law School and the University of Vienna School of Law



TTLF Working Papers

No. 141

**Beyond Accuracy: Regulating Hallucination
Risks in Generative AI in the EU and US**

Zihao Li, Weiwei Yi and Jiahong Chen

2025

TTLF Working Papers

Editors: Siegfried Fina, Mark Lemley, and Roland Vogl

About the TTLF Working Papers

TTLF's Working Paper Series presents original research on technology, and business-related law and policy issues of the European Union and the US. The objective of TTLF's Working Paper Series is to share "work in progress". The authors of the papers are solely responsible for the content of their contributions and may use the citation standards of their home country. The TTLF Working Papers can be found at <http://tflf.stanford.edu>. Please also visit this website to learn more about TTLF's mission and activities.

If you should have any questions regarding the TTLF's Working Paper Series, please contact Vienna Law Professor Siegfried Fina, Stanford Law Professor Mark Lemley or Stanford LST Executive Director Roland Vogl at the

Stanford-Vienna Transatlantic Technology Law Forum
<http://tflf.stanford.edu>

Stanford Law School
Crown Quadrangle
559 Nathan Abbott Way
Stanford, CA 94305-8610

University of Vienna School of Law
Department of Business Law
Schottenbastei 10-16
1010 Vienna, Austria

About the Authors

Dr Zihao Li has been a TTLF Affiliate at Stanford Law School since 2023. He is also a Lecturer (Assistant Professor) at CREATE Centre in the School of Law at the University of Glasgow (UK). He was awarded the UKRI Fellowship on Metascience AI. In CREATE Centre, he co-leads the theme on Decentralization, Automation and Platform. Zihao is qualified in both Computer Science (B.Eng.) and Law (LL.B., LL.M., Ph.D.). With his interdisciplinary background, his research interests concentrate on the intersection of law, data and information technology. Recently, his research mainly includes AI and regulation, data protection law, algorithmic pricing, Internet and intellectual property, and blockchain and law. His research has been published in top-tier interdisciplinary academic journals, conference proceedings and books, including Nature Machine Intelligence, Computer Law and Security Review, IEEE Communications Magazine, IEEE International Conference on Communications, and European Data Protection Law Review. He has also been invited to speak at several prestigious universities, including the University of Oxford, University of Cambridge, and Imperial College London. Zihao has been awarded the Modern Law Review (MLR) Scholarship operated by the LSE. Additionally, Zihao is a co-founder of the Scottish Law and Innovation Network (SCOTLIN), and the founding president of the Intellectual Property Society at the University of Glasgow.

Weiwei Yi is a PhD candidate at the Centre for the Regulation of the Creative Economy (CREATE) at the University of Glasgow. Her research focuses on dark design patterns in online games, examined through data protection law, consumer protection law, and platform regulation. She is additionally interested in the regulation of manipulative AI practices in games, virtual reality, and AI chatbots.

Dr Jiahong Chen is a Lecturer (Assistant Professor) in Law at the University of Sheffield (UK). His research focuses on the intersection between law and technology, in particular data protection law, cybersecurity law, law and AI, data ethics and internet regulation. His research has been funded by the ESRC and RAI UK, and published in leading journals such as International Data Privacy Law.

General Note about the Content

The opinions expressed in this paper are those of the authors and not necessarily those of the Transatlantic Technology Law Forum or any of its partner institutions, or the sponsors of this research project.

This paper is a comparative adaption version of the preprint: Zihao Li, Weiwei Yi and Jiahong Chen, 'Accuracy Paradox in Large Language Models: Regulating Hallucination Risks in Generative AI' (arXiv, 2025) <<https://arxiv.org/abs/2509.13345>>.

Suggested Citation

This TTLF Working Paper should be cited as:
Zihao Li, Weiwei Yi and Jiahong Chen, Accuracy Paradox in Large Language Models:
Regulating Hallucination Risks in Generative AI in the EU and US, Stanford-Vienna
TTLF Working Paper No. 141, <http://tlf.stanford.edu>.

Copyright

© 2025 Zihao Li, Weiwei Yi and Jiahong Chen

Abstract

Hallucination in generative AI is often treated as a technical failure to produce factually correct output. Yet this framing underrepresents the broader significance of hallucinated content in language models, which may appear fluent, persuasive, and contextually appropriate while conveying distortions that escape conventional accuracy checks. This paper critically examines how regulatory and evaluation frameworks have inherited a narrow view of hallucination, one that prioritizes surface verifiability over deeper questions of meaning, influence, and impact. We propose a layered approach to understanding hallucination risks, encompassing epistemic instability, user misdirection, and social-scale effects. Drawing on interdisciplinary literature, the article develops a taxonomy of hallucination and demonstrates the limitation of solely using accuracy as a benchmark. It then compares how the EU and the United States address these risks. In the EU, instruments such as the AI Act, GDPR, and DSA place accuracy at the heart of their compliance logic, incentivizing optimization of surface-level correctness while overlooking deeper epistemic and systemic harms. In the US, a fragmented patchwork of consumer protection, competition, and sectoral rules offers more flexibility but lacks coherent safeguards against hallucination-related harms. This comparative analysis shows how both approaches, though divergent in form, risk amplifying the same epistemic vulnerabilities: misplaced user trust, subtle manipulation, and epistemic convergence that undermines pluralism. Rather than improving factual precision alone, we argue for regulatory responses that account for language's generative nature, the asymmetries between system and user, and the shifting boundaries between information, persuasion, and harm.

Keywords

Accuracy; Hallucination; Artificial Intelligence; Large Language Models; AI Regulation; AI Act; GDPR

Table of Contents

1. Introduction	1
2. Hallucination: Concepts, Taxonomy, and Reasons	6
3. Understanding the Limitations of Accuracy	10
3.1. Accuracy vs. Trustworthiness	10
3.1.1. Accuracy Is Not Truth: Epistemic Limits of Predictive LLMs	11
3.1.2. Overreliance on Accuracy Can Create Over-trust	14
3.1.3. Accuracy-Transparency Trade-Off: Accuracy Can Undermine Meaningful Transparency and Interpretability	16
3.2. Accuracy vs. Autonomy	19
3.2.1. Accuracy vs. Epistemic Independence: The Rhetorical Illusion	20
3.2.2. Accuracy vs. Manipulation Resilience: The “Not Inaccurate” Blind Spot	22
3.2.3. Static accuracy vs. Dynamic Interaction: Ignorance of Interaction-Induced Distortions	26
3.3. Accuracy vs. Social Progression	29
3.3.1. Accuracy vs. Equity: Re-identification and social sorting	30
3.3.2. Accuracy vs. Plurality: Epistemic Convergence	32
3.3.3. Accuracy vs. Criticality: Social Deskilling and Stagnation	35
4. Comparative Analysis of EU and US Regulations	38
4.1. The EU AI Act and the U.S. Algorithmic Accountability Bills	39
4.2. The EU GDPR and US Data Protection Regulations	40
4.3. The EU DSA and US Platform Transparency Initiatives	43
5. Potential Ways Forward: Beyond Accuracy	45
5.1. From Rhetorical Device to Epistemic Trustworthiness	46
5.2. Embracing Pluralism: Designing for Diversity of Sources and Views	46
5.3. Reassessing Hallucination: Creativity, Temperature, and Use-Sensitivity	49
6. Conclusion	49

1. Introduction

The rapid development and deployment of generative AI, particularly large language models (LLMs), have transformed the way information is generated, disseminated and consumed. As these models become increasingly integrated into critical domains such as healthcare, education, and law, their outputs carry significant epistemic, legal, and societal implications. One of the new

risks posed by LLMs is hallucination, where the generated output is fabricated, nonsensical, subtly inaccurate, oversimplified, sycophantic or biased, yet delivered in a very confident tone (Ji et al. 2022; Li 2023; Wachter et al. 2024; Hacker et al. 2025a). Drawing on philosophy of science, such a phenomenon has been conceptualised as “bullshit” (Hicks et al. 2024). Hallucination has caused various types of practical harms, such as misinformation (Sun et al. 2024), disinformation (Bandara 2024), defamation (Binns and Edwards 2024), privacy infringement (Yao et al. 2024; Solove 2025) and even more serious mental and physical damages. Examples of the individualised and collective harms and damages abound, such as when a Norwegian user asked ChatGPT if it had any information about him, the chatbot made up a story that he had murdered his children (noyb 2025a). Similarly, it has been reported that ChatGPT falsely accused an American law professor by including him in a generated list of legal scholars who had sexually harassed someone, citing a non-existent The Washington Post report (Sankaran 2023). Even before the ChatGPT went viral, hallucination had already caused significant issues. A passenger booked a flight on Air Canada’s official website, inquired about special discounts through a chatbot, which incorrectly told him he could apply for reimbursement after flying (Belanger 2024).

One of the common regulatory approaches, widely recognised among regulators, academia and technologists, is to prioritise *accuracy* as a primary metric for assessing the reliability of AI-generated content. For example, the UK ICO (the UK’s data protection authority) consultation on its policy position on generative AI has a clear focus on accuracy of training data and model outputs, which emphasises how the accuracy principle applies to the outputs of generative AI models, and the impact that the accuracy of training data has on the output. The pursuit of

accuracy to combat hallucination is further highlighted by the ICO's response (UK Information Commissioner's Office 2024; Li et al. 2024). Similarly, the European Data Protection Board's report concerning LLMs privacy risks reiterates accuracy as a common metrics to mitigate hallucination by which the factual errors could be reduced (Barberá 2025, p. 18). In the context of Agentic AI, the EDPB's report further subdivides the accuracy metrics into task-specific accuracy and step-level accuracy to gauge the coherence and correctness of intermediate steps and the final outcome. Similarly, the European Data Protection Supervisor's¹ guideline stresses the importance of ensuring the accuracy of structure and content of the training datasets, stating that "metrics on statistical accuracy (the ability of models to produce correct outputs or predictions based on the data they have been trained on), when available, can offer an indicator for the accuracy of the data the model uses as well as on the expected performance"(Barberá 2025, p. 15).

Many EU national Data Protection Authorities (DPAs) have also issued guidance in which accuracy is given paramount importance. For example, Belgium DPA's report explains that organisations must take reasonable steps to ensure the implementations of accuracy of personal data and require high-risk AI systems to use high-quality and unbiased data to prevent discriminatory outcomes and unintended behaviours (Data Protection Authority of Belgium 2024). Similarly, the Swedish DPA explicitly frames the GDPR's accuracy principle as both a normative standard and a practical regulatory tool to mitigate the risks of hallucinations in generative AI (Swedish Data Protection Authority (IMY) 2025). By highlighting accuracy-

¹ The EDPB is an independent EU body established to ensure the consistent application of the GDPR. The EDPS is an independent supervisory authority responsible for monitoring the processing of personal data by EU institutions and bodies. The EDPS is a member of the EDPB.

related measures (e.g., retrieval-augmented generation (RAG), confidence thresholds, and human review), the report positions accuracy not merely as a compliance obligation but as a dynamic mechanism for reducing the production of factually incorrect or misleading content (Swedish Data Protection Authority (IMY) 2025). Similar issue has also been raised by Non-Governmental Organization (NGO), such as noby, to fight against LLMs firms on the ground that their systems provide inaccurate information (noyb 2025b).

A plethora of tech firms and third-party organisations similarly position accuracy as an unmitigated goal. As illustrated in Figure 1, OpenAI frequently highlights accuracy when announcing new models, using it as a proxy to indicate reductions in hallucination rates. Likewise, Google’s Gemini platform presents accuracy as a marker of enhanced performance in its launch of Gemini 2.5 (Google 2025). Third-party organisations are no exception in the evaluation of LLM models, such as AI Multiple employ accuracy as a key benchmark to assess and compare the hallucination rates across different LLMs.

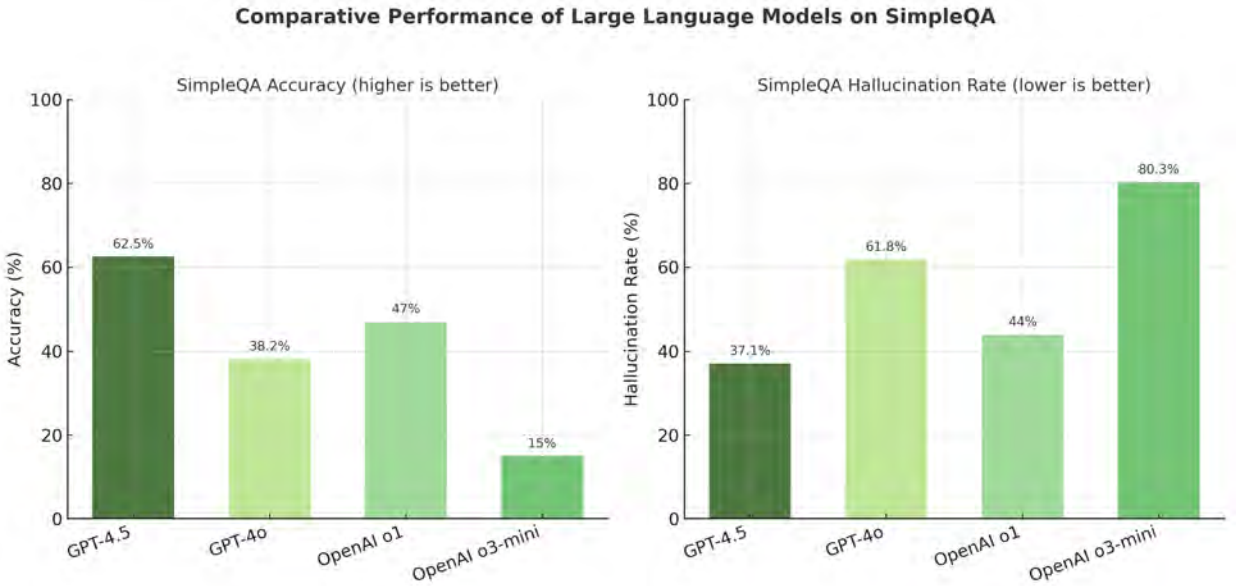


Figure 1 OpenAI Accuracy Demonstration (OpenAI 2025a)

As a result, accuracy has become an implicit assumption, either as an unchallenged good or as a technical necessity, based on the belief that high accuracy can mitigate hallucination and thereby ensure the responsible adoption of AI systems. Admittedly, improving accuracy can, to some extent, alleviate hallucination in LLMs, particularly by enhancing the factual reliability of model outputs, but the extent to which this should be translated into policymaking and regulatory objectives has not been sufficiently interrogated. In this regard, this article aims to conceptualise the ramifications of an accuracy-centred regulatory approach. We argue that an overreliance on accuracy as the primary benchmark for mitigating hallucination risks is both conceptually narrow and normatively insufficient. It risks obscuring a broader set of concerns, including the pursuit of epistemic truth, the trustworthiness and interpretability of information, and the diversity and heterogeneity of viewpoints. By framing hallucination solely as a technical flaw to be “fixed” through accuracy enhancement, we might risk overlooking the nuanced sociotechnical context under which information is constructed, received, and interpreted.

This paper advances the concept of the accuracy paradox to capture a central dilemma in the governance of large language models. Efforts to mitigate hallucinations by hyper-optimising and over-relying on accuracy can inadvertently intensify existing harms or generate new ones. The paradox lies in the fact that the closer a system comes to mimicking factual authority through heightened accuracy, the greater the risk of creating a false sense of epistemic certainty, encouraging uncritical reliance, and weakening essential checks and balances. Accuracy, narrowly defined in statistical or benchmarking terms, may reduce overt hallucinations, yet simultaneously deepen informational and cognitive vulnerabilities by obscuring manipulation,

narrowing epistemic diversity, and reinforcing blind trust. This is not to argue against improving accuracy per se. Rather, it underscores the need for a pluralistic governance approach that situates accuracy within a broader framework of epistemic trustworthiness, interpretability, user autonomy, and collective social good.

Building on this conceptual foundation, the paper proceeds in four steps. First, it introduces hallucination through a socio-technical taxonomy that illuminates its varied forms and underlying causes. Second, it develops the concept of the accuracy paradox, showing how overemphasis on accuracy, while superficially beneficial, conceals deeper epistemic, manipulative, and societal risks that undermine responsible AI development. Third, it conducts a comparative analysis of EU and U.S. regulatory frameworks, including the GDPR, DSA, and AI Act on one side, and the Algorithmic Accountability Acts, Algorithmic Justice Act, Blueprint for an AI Bill of Rights, and related initiatives on the other. This analysis evaluates how each jurisdiction conceptualises and operationalises accuracy, and the extent to which their frameworks address or fail to address the accuracy paradox. Finally, the paper proposes pathways beyond accuracy, drawing on interdisciplinary insights to articulate governance solutions that mitigate hallucination risks. These include technical mechanisms such as uncertainty expression and pluralistic evaluation benchmarks, alongside legal and institutional measures such as user empowerment, transparency regimes, and polycentric oversight. Together, these approaches aim to foster AI systems that are not only statistically accurate but also epistemically reliable, manipulation-resilient, and socially responsible.

2. Hallucination: Concepts, Taxonomy, and Reasons

To put it simply (and perhaps overly-simplistically), the main logic behind the LLMs is to tokenise the text and data retrieved from training materials and to predict which string of words has the most likely sequence of words coming next (Li 2023; Binns and Edwards 2024). This probabilistic architecture, grounded in statistical correlations rather than grounded understanding, is the primary reason why hallucination emerges as an inherent feature of LLM outputs (Li 2023). Typically a model is further fine-tuned via “reinforcement learning from human feedback” (RLHF) to ensure the consistency, coherence of model’s output and make their outputs more human-like and persuasive. However, such mechanism cannot guarantee the quality of generated content and does not incentivise the model to seek factual or trustworthy output, as it only predicts the likelihood of the next words, rather than comprehending, pondering and reasoning the underlying meaning or context behind the training data or users’ prompt (Zucchet et al. 2025). Moreover, RLHF could further exacerbate this effect as it imperceptibly introduces latent annotator subjective feedback and reinforces the model’s awarding mechanism, which may unintentionally amplify inaccurate, biased (Bai et al. 2025) or misleading responses. RLHF has been criticised for having shaped AI into a sycophantic, people-pleasing persona that mirrors the human reluctance to confront uncomfortable truths: When AI constantly apologises, avoids offense, and earnestly fabricates answers to preserve our feelings, it lays down a series of gentle traps: subtle, well-intentioned, but ultimately deceptive (Sharma et al. 2024). Such behaviour exemplifies sycophancy hallucination, where outputs are not directed towards epistemic reliability but towards maximising user comfort and satisfaction, producing responses that appear reassuring yet remain epistemically hollow and untrustworthy.

Especially in zero-shot contexts, where the LLMs do not have a specific training dataset for the given task or domain, the model is forced to rely on general patterns learned during pretraining, which can increase the likelihood of hallucinations, inaccuracies, or contextually inappropriate responses (Kojima et al. 2023). Meanwhile, any misinterpreting prompts may further exacerbate the risks of hallucination. Without a proper and specific prompt to guide models in zero-shot setting, the inherent ambiguity may lead the model to overconfidently generate plausible-sounding yet unfounded outputs, thereby compounding the epistemic instability of its responses.

LLMs not only fabricate or generate “nonsense” content, but also prone to produce subtle mistruths (sometime combined with real facts), alongside oversimplifications of complex topics and responses biased towards certain commonly held beliefs (Bender et al. 2021; Li 2023; Wachter et al. 2024). It has also been documented in literature that the way that LLMs express and interact with users are hardly reflection of their internal workings, as they tend to convey information in a very confident way, which could misguide professional experts, let alone lay people (Chen et al. 2025).

It should be noted that hallucinations can take different forms, and there is a need to differentiate various types of hallucination. Building on the work of Huang et al. (2025), and drawing on the wider literature, we present a taxonomy of hallucination, as outlined in Table 1.

Category	Type	Descriptions
Factuality Hallucination	Factual Contradiction	Factual errors occur when a model generates information that is objectively incorrect or inconsistent with established facts. It includes (1) Entity-error hallucination: erroneous entities, e.g., wrongly saying Thomas Edison invented the telephone; (2) Relation-error hallucination: wrong relations, e.g., misstating Edison's role in inventing the light bulb.
	Factual Fabrication	Factual fabrication refers to the generation of entirely false or invented information by a model, which does not correspond to any real-world facts or data. (1) Unverifiability hallucination: Unverifiable statements, e.g., Eiffel Tower construction led to Parisian tiger extinction (non-existent species); (2) Overclaim hallucination: Claims lacking universal validity, e.g., Eiffel Tower sparked global green architecture movement.
	Conflation of Facts	Conflation of facts occurs when a model combines or merges distinct pieces of information, creating a distorted or misleading representation of reality. This often involves blending facts from different contexts, events, or sources in a way that makes them appear as if they are connected or part of a single narrative, even when they are not.
Consistency Hallucination	Instruction Inconsistency	Outputs deviating from user directives, e.g., answering instead of translating.
	Context Inconsistency	The output is inconsistent with the provided context, such as adding information that is not present in the given context.
	Logical Inconsistency	Internal contradictions in reasoning, e.g., $9.11 > 9.8$
Reference Hallucination	Source Fabrication	This occurs when the model references or cites sources that do not actually exist or are fabricated
	Misattribution	This happens when the model attributes information or quotes to the wrong source, either inaccurately referencing the original source or incorrectly attributing content to someone or something else.
Sycophancy Hallucination	Overly Complimentary	This type occurs when the model generates content that excessively praises a person, idea, or organization, often beyond what is warranted or realistic, in order to align with perceived preferences or authority figures.
	Flattery Bias	Model produces outputs that are skewed to present the person or entity in the most favourable light, possibly distorting facts or providing biased interpretations to maintain favour or avoid conflict.
	Subservient Tone	The model generates content that adopts a deferential or excessively accommodating tone, suggesting compliance or agreement with a particular individual or viewpoint in an exaggerated way.
Consensus Illusion	Narrowed Perspective	The generated output only reflects a narrow or limited perspective, giving the false impression that there is a broad consensus or agreement on an issue. However, there may be significant disagreement or diversity of opinion.
Oversimplified Hallucination	The generated content that reduces complex or nuanced information to overly simplistic terms, ignoring important details, context, or variations. This often results in misleading or incomplete answers that may appear clear and straightforward but fail to accurately reflect the complexity of the issue	
Prompt- Sensitivity Hallucination	Sandbagging	LLM automatically reduces the quality of its responses in reaction to prompts perceived as "lower quality", such as those with poor grammar, typos, or informal structure.
	Emotionally induced drift	A tendency for LLMs to produce misaligned or overly reassuring outputs in response to emotionally charged prompts, regardless of the factual or epistemic grounding of the response.

Table 1 Taxonomy of hallucinations

From Table 1, it is clear that the sociotechnical complexity of hallucination extends far beyond mere factual inaccuracies. It involves inconsistent responses, reference fabrication, and more subtle psychological and social issues such as sycophancy tendency and consensus illusion. The diversity of AI hallucination means that a narrowly defined concept of accuracy may not be fit for purpose when guiding regulatory responses. In the next section, we provide a theoretical account on why this might be the case by introducing the concept of accuracy paradox.

3. Understanding the Limitations of Accuracy

Over-reliance on improving accuracy of LLMs does not always reduce harm and may, in certain contexts, exacerbate it. While accuracy is conventionally treated as a normative and unquestioned good, this section interrogates its limits and unintended consequences. We use the term *accuracy paradox* to describe the scenario where pursuing higher accuracy of LLMs by model developers and policymakers with a view to addressing hallucination harms, with a narrow understanding of accuracy, may paradoxically oversee or even give rise to more subtle forms of harms. We demonstrate this phenomenon by contrasting accuracy with three dimensions of desirable policy goals of AI regulation: *output trustworthiness*, *individual autonomy* and *social progression*.

3.1. Accuracy vs. Trustworthiness

The notion of trustworthiness, often invoked as an implicit normative aim of AI regulation, contains a compound epistemic and procedural value (Laux et al. 2023).² In the context of

² Art. 1, EU Artificial Intelligence Act: <https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng>

LLMs, this article differentiates three analytically distinct but interrelated dimensions of trustworthiness: epistemic validity, user perception, and interpretability. First, trustworthiness entails that an output must not only be statistically accurate but also epistemically grounded, anchored in justified reasoning, verifiable knowledge or even falsifiable argument (Section 3.1.1). In philosophical terms, this links trustworthiness to the pursuit of truth, not as mere factual correspondence, but as a normative value requiring justification and resilience to error, thereby distinguishing reliable knowledge from outputs that only appear accurate (Popper 1959; Williams 2002). Second, trustworthiness concerns user psychology and behavioural reliance: outputs that appear linguistically accurate may induce unwarranted trust, particularly when users forego critical scrutiny (Section 3.1.2). Third, trustworthiness depends on transparency and interpretability. Users must be able to trace the reasoning process or access uncertainty signals in order to make informed judgments about the credibility of a given response (Section 3.1.3). Taken together, these three perspectives demonstrate that trustworthiness is not reducible to accuracy alone. Rather, trustworthiness emerges from the interplay between internal epistemic mechanism, output quality, user trust calibration, and the visibility of the epistemic process that generated the output. Without such foundation, the regulatory promise of “trustworthy AI” risks collapsing into a rhetorical ideal rather than a meaningful standard. The following subsections examine each of these dimensions in turn.

3.1.1. Accuracy Is Not Truth: Epistemic Limits of Predictive LLMs

Higher accuracy requirements do not necessarily diminish the potential harms users receive, but may instead increase the risk of the users trusting misinformation. The primary reason for such a paradox is that accuracy is not equal to truth. From a legal and philosophical perspective,

accuracy often refers to the consistency of information with a predetermined, up-to-date and “ground truth” dataset,³ whereas truth is a more complex philosophical concept that involves a deeper, context-driven evaluation that aligns with both factual correctness, epistemic and social coherence. From a philosophical viewpoint, truth encompasses various schools of thought in philosophy, such as positivism, constructivism, correspondence theory, coherence theory and consensus theory (Dowden and Swartz 1995; Junior 2023). Over-reliance on accuracy oversimplifies the requirement of truth and trustworthy information, which may create epistemic overreach. Truth, in this broader philosophical sense, is not a singular or purely empirical concept but a contested and multi-dimensional construct. At its core, truth concerns the relationship between statements and reality, yet how that relationship is defined varies across traditions. Correspondence theory defines truth as the alignment between a statement and an objective fact or state of affairs (Dowden and Swartz 1995). Coherence theory considers a belief true if it fits consistently within a broader system of beliefs (Dowden and Swartz 1995). Consensus theory views truth as that which is agreed upon by a community under ideal conditions (Habermas 2005), while constructivism understands truth as shaped by social, cultural, or historical contexts (Kukla 2000, pp. 7–18). In contrast to these definitions, positivism grounds truth in empirical verification and logical consistency. While accuracy based on ground truth may achieve what is close to positivist truth, it is unlikely to fulfil the test in other theories as they concern not merely isolated correctness or accuracy but about justification, context, and interpretive frameworks—dimensions that statistical accuracy alone cannot adequately capture.

³ Article 4, EU General Data Protection Regulation (GDPR): <https://eur-lex.europa.eu/eli/reg/2016/679/oj/eng>

The pursuit of truth for LLMs as an ultimate goal is essential. Otherwise, LLMs may combine accurate fragments into a misleading whole due to a lack of contextual understanding or epistemic grounding, leading to factuality and consistency hallucinations. A technically “accurate” response can still mislead, misrepresent, or obscure important nuances, thereby creating an illusion of trustworthiness where none exists. In this light, accuracy without truth not only fails to resolve the problem of hallucination but may, paradoxically, deepen its epistemic and normative harms.

As illustrated in Section 2 on the mechanisms of LLMs, accuracy is judged by token prediction success or benchmark alignment, which is statistical, not epistemic. Therefore, the increase of accuracy in LLMs only enhances the syntactic or probabilistic plausibility of output, rather than their epistemic validity. In other words, a model could become better at predicting what sounds correct based on training data, without any grounding in whether the content is actually true, justified, or verifiable. This distinction is critical: a linguistically fluent and statistically probable output may still be epistemically hollow, particularly when the model lacks access to mechanisms for truth verification or source attribution.

It is perceivable that even when an LLM is based on a comprehensive ground truth corpus, its underlying generative architecture may still remain susceptible to hallucination. This is not merely a question of data coverage, but of architectural constraints. Fundamentally, given the probabilistic nature of token prediction, the model’s output is a function of likelihood estimation rather than semantic understanding or truth retrieval. Errors can still emerge from token-level misalignment, prompt sensitivity, and the diffusion of meaning across latent spaces. In other

words, even if every “correct” answer exists somewhere in the model’s training data, there is no guarantee that the model will retrieve, compose, or represent it truthfully in response to any given prompt. This reveals a fundamental fragility: the hallucination problem cannot be resolved solely by improving training data accuracy, because the architecture itself does not guarantee faithful grounding or truth-seeking. Over time, such outputs risk accumulating what Zittrain (2022) refers to as “intellectual debt”: responses that appear authoritative yet obscure their own opacity, thereby undermining long-term epistemic trust. This explains why factuality, inconsistency and reference hallucination often occur.

Therefore, overreliance on accuracy amounts to an oversimplification of the truth requirement. Equating truth with accuracy as measured against the training data and output’s consistency is inherently flawed. This approach fails to account for the depth and contextual nuance required to genuinely reflect truth, as it reduces the concept to a mechanical alignment with surface-level vectorised data without addressing broader epistemic values like coherence, social context, or representational fairness (See Section 3.2 and 3.3). It overlooks the complexity of truth as something that extends beyond statistical conformity to include deeper, context-driven trustworthiness, validation and reliability. As a result, merely improving accuracy masks the absence of truth, ultimately undermining the trustworthiness of LLMs’ outputs.

3.1.2. Overreliance on Accuracy Can Create Over-trust

This misalignment of accuracy and truth leads to information that is often not trustworthy. High accuracy requirements can heighten the users’ over-trust without critically evaluating it (Klingbeil et al. 2024, p. 2). As measured accuracy rises, users generalize from aggregate

performance to instance-level reliability, accept outputs at face value, and verify less, which leads to new harms: over-trust of AI. In other words, merely improving the accuracy of the models is insufficient, because the more accurate the model is, the more users will rely on it, and thus be tempted not to verify the answers, leading to greater risk when hallucinations appear (Li 2023). This problem stems from the fact that LLMs are fundamentally designed to produce fluent and convincing responses, rather than inherently comprehending content, reasoning and pursuing truth as an epistemic goal. Even if a model achieves high accuracy, it cannot guarantee a fully trustworthy response because LLMs merely predict the likelihood of word sequences without truly understanding the content they generate. If even 0.1% or 0.2% of answers in a particular domain are incorrect, it creates significant challenges for users attempting to discern the authenticity of those responses, especially in legal, medical or other high-risk domains (Li 2023). Therefore, overreliance on accuracy could lead to a false sense of trustworthiness, where users assume outputs are fully reliable despite the presence of subtle but consequential errors.

In this regard, demonstrating accuracy directly to users could also mislead users and cause them to over-trust. A sole emphasis on accuracy rates foregrounds surface-level factual correctness, targeting only factual hallucinations, while obscuring deeper epistemic concerns. As highlighted above, the deployment of Reinforcement Learning through Human Feedback (RLHF) could introduce several concrete epistemic limitations. These limitations cannot be resolved through accuracy rate improvements alone; in fact, the pursuit of higher accuracy may further entrench the misleading nature of model outputs by masking their underlying uncertainties and reinforcing rhetorically plausible but potentially untrustworthy responses. RLHF embeds human labellers' subjective views, values and backgrounds into the system, which do not always reflect

trustworthiness comprehensively. As a result, what may be perceived as “true” by these systems is often merely an outcome of probabilistic modelling and human reinforcement, rather than an inherent understanding or intent to pursue the truth. Generated responses, therefore, are more of a product of chance and reinforcement than any genuine alignment with truth. Worse still, fine-tuning through human feedback further exacerbates such flaws, as LLMs would inadvertently prioritise persuasive or rhetorically appealing responses over those that are truthful (Zhou et al. 2024). It often generates outputs that sound confident or convincing, rather than those grounded in verifiable and critical information (See Section 3.2). This creates an accuracy paradox: while efforts to enhance accuracy may improve the articulation and persuasiveness of outputs, they do not necessarily make them more trustworthy, likely reinforcing biases and creating the illusion of reliability without a genuine commitment to truth. In this context, improvements in accuracy can lead to over-trust and mislead users into abandoning doubt and independent judgment, creating a false sense of reliability that discourages scrutiny of the model’s outputs (See Section 3.3).

3.1.3. Accuracy-Transparency Trade-Off: Accuracy Can Undermine Meaningful Transparency and Interpretability

Overreliance on accuracy also often comes at the cost of transparency and interpretability of the generated output (Kossow et al. 2021). At first glance, enhancing accuracy appears to be a straightforward improvement in model performance. Yet, as models grow in complexity and parameter, the opacity of their internal reasoning deepens and it becomes more difficult to trace back the epistemic opacity (Chen et al. 2025), which undermines the transparency and interpretability of generated content. As exemplified by the OpenAI reasoning model o1, the chain of thought (CoT) and reasoning trace are often cloaked. Some users have reported getting

warnings or even having their accounts blocked whenever their prompts include terms like “reasoning trace” or “show your chain of thought” (Awayyy 2024; Mindfultime 2024). It has been reported that when prompted to explain its reasoning, Claude fabricates plausible-sounding explanations and tailors its false reasoning to agree with misleading prompts (The Economist 2025). This further illustrates the “illusion of thinking”: LLMs may pretend reasoning in ways that convince users of genuine thought, yet as Apple’s recent work demonstrates, such reasoning often amounts to pattern matching rather than transparent and verifiable logical reasoning (Shojaee et al. 2025, pp. 5–7). It also demonstrates that CoT in Reasoning Models (LRMs), such as Claude 3.7 and DeepSeek-R1, fail to develop generalisable problem-solving capabilities and often collapse entirely when faced with higher-complexity tasks (Shojaee et al. 2025). Even when provided with explicit algorithms to execute, these models fail to perform accurate step-by-step reasoning, revealing limitations in symbolic manipulation. Moreover, their “thinking traces” often include internally inconsistent or redundant reasoning, and models may even reduce their reasoning effort as problems become more complex. This illustrates the opacity of current LLMs with respect to revealing their internal decision-making processes, which in turn compromises users’ ability to verify, question, or contest the system’s outputs.

Assumptions that accuracy is an unmitigated good are contested in this context, as they overlook the complex trade-offs between surface-level correctness and deeper epistemic values such as transparency, interpretability, justification, and trustworthiness. Trustworthiness, in this regard, is not an inherent property of output, but an emergent requirement that users need to access and interpret the underlying rationale of the system’s response. This makes meaningful transparency and interpretability critical preconditions: without knowing how or why a model arrived at a

conclusion, users are deprived of the epistemic scaffolding necessary to evaluate its legitimacy or challenge its implications. Transparency alone does not guarantee this, as it is normally limited to disclosure of system components. Interpretability helps to complement this gap by enabling the tracing of reasoning, assessment of confidence, and identification of omissions, which renders outputs assessable, justifiable, and ultimately trustworthy. As a result, in such contexts, only improvements in accuracy may paradoxically erode trustworthiness, precisely because they induce users to accept outputs without the capacity to interrogate how those outputs were generated. When the reasoning process is inaccessible or obscured, users are deprived of the epistemic tools needed to evaluate the credibility of responses, encouraging passive acceptance rather than critical engagement. This mechanism undermines the normative foundations of trust, which relies not merely on outcome quality but on the ability to assess the integrity of the process by which that outcome is produced.

In the same vein, as discussed in Section 3.1.2, if accuracy improvements primarily manifest in enhanced linguistic fluency and rhetorical confidence, such convincingly expressed outputs may obscure underlying model uncertainty, particularly in zero-shot contexts where the model lacks task-specific training data. In these cases, the model's internal epistemic uncertainty is not meaningfully communicated through its surface-level expression (Chen et al. 2025). That is, even when the model has low confidence in its response, it often presents the output with syntactic precision and stylistic confidence, thereby concealing the absence of grounding or justification. This misalignment between internal uncertainty and external presentation undermines meaningful transparency and interpretability. Users are not given access to indicators of confidence or error margins, nor are they able to trace the model's reasoning in a way that

would allow them to assess reliability (Yin 2025). In effect, the pursuit of accuracy, when limited to surface-level plausibility, reinforces the accuracy paradox: it renders the output more persuasive, thereby undermining transparency while simultaneously depriving users of the epistemic tools necessary for critical scrutiny, deliberation, and informed decision-making. As a result, in such contexts, improvements in accuracy may paradoxically diminish trustworthiness, as they mislead users by obscuring the reasoning process and preventing meaningful evaluation of the output's credibility.

3.2. Accuracy vs. Autonomy

Beyond output-level risks, when a user is exposed to an accumulation of responses over time, the hyper-focus on the statistical accuracy of LLM-generated outputs may also risk obscuring deeper concerns around manipulation, ultimately undermining users' autonomy. Three key dynamics drive this risk. First, statistical accuracy can function as a disguise. Regulatory metrics that emphasise accuracy incentivise developers to optimise surface-level fluency and rhetorical persuasiveness, rather than ensuring epistemic reliability. Users are thus not persuaded by sound reasoning, but by linguistically confident outputs, leading to a misplaced sense of trust in the model's reliability. This is linked to the trustworthiness issue flagged up in the previous section but pertains to a different concern: even if a series of outputs individually fulfil the "truth" criterion (however it is defined), they are not necessarily "neutral" in the sense that they do not steer the users to a particular behavioural or perceptual direction. Second, accuracy as a singular benchmark is inherently limited. It fails to capture the epistemic and contextual complexity of LLM outputs during dynamic interactions, particularly when responses are opinion-based, value-laden, or socially contingent. In other words, it cannot account for outputs that fall into the grey

area of “not being inaccurate” responses that evade absolute factual error while still exerting potentially manipulative influence. Third, accuracy, as a static benchmark, fails to account for the dynamic deterioration that emerge through continued user interaction in evolving real-world contexts. Such static evaluation is ill-equipped to capture the effects of training-test contamination, which artificially inflates accuracy by rewarding memorisation, and the evolving dynamics of user interaction (e.g., prompt poisoning), which causes consistency and sycophancy hallucinations. When this fragility is mistaken for stable performance, it undermines users’ ability for informed judgments and leaves them vulnerable to subtle manipulation and erode autonomy.

As a result, the more we rely on and optimise for accuracy, the more we risk overlooking the subtle ways in which persuasive yet opaque outputs can influence, nudge, or manipulate user behaviours, often without users being fully aware of such influence. The rest of this section will elaborate on each of the three kinds of risks in turn.

3.2.1. Accuracy vs. Epistemic Independence: The Rhetorical Illusion

When accuracy becomes a regulatory anchor, optimisation often targets linguistic fluency, rhetorical confidence, and audience fit as a means of influencing users, rather than enhancing models’ epistemic reliability, internal understanding or reasoning capacity. This is because, as explained above, the most straightforward and computationally efficient way to signal reliability is to enhance linguistic fluency and persuasive tone, as these features that lead users to perceive the output as accurate and trustworthy. In such cases, the appearance of accuracy functions less as a marker of epistemic validity and more as a rhetorical disguise, one that can mislead users

into overestimating the reliability of the content. As demonstrated by Okoso et al. (2025), AI expressions significantly change users' decisions. When interacting with AI-generated responses, users' choices were shaped by the tone of the output regardless of their prior knowledge or individual attributes. Older users were especially susceptible to tonal influence, and highly extroverted individuals often made decisions that diverged from their stated perceptions. Similarly, Salvi et al. (2025) demonstrate that when individuals lack strong prior opinions on a given topic, GPT-4 proves significantly more persuasive than human counterparts under conditions of microtargeting. Specifically, GPT-4 made participants 81.2% more likely to change their views towards its assigned position after the engagement, outperforming human persuaders in 64.4% of cases. Crucially, such rhetorical strength is not grounded in epistemic superiority or factual depth, but in the optimisation of surface-level linguistic features, further illustrating how accuracy enhancements may amplify user persuasion while obscuring the absence of transparent reasoning or verifiable information. Paradoxically, these surface-level improvements may erode users' epistemic independence by encouraging passive acceptance rather than critical engagement, thereby undermining individual's autonomy.

Such persuasive power of LLMs, when embedded within interactive systems, gradually shifts from persuasion to what Luciano (2024) refers to as hypersuasion: a technologically mediated mode of influence that empowers the system to identify, target, and exploit users' cognitive and emotional susceptibilities with unprecedented subtlety and reach. As suggested by Yeung (2017), such rhetorical strength gives rise to what is known as "hypernudge" where algorithmic architectures, continuously updated and highly personalised, can restructure the user's informational environment in real time to shape behaviour without the user ever realising it.

Together, they point to a disturbing possibility: that AI systems, under the guise of fluency and relevance, may intentionally or unintentionally manipulate users beneath the threshold of awareness, not by being inaccurate, but precisely by being rhetorically too accurate.

3.2.2. Accuracy vs. Manipulation Resilience: The “Not Inaccurate” Blind Spot

Overstressing accuracy as a singular benchmark is inherently flawed, as it creates a false sense of trustworthiness that blinds users and regulators to the subtler manipulative risks embedded in outputs that are not being inaccurate. Unlike verifiable questions in mathematics, coding, or scientific facts (OpenAI 2024a), many answers generated by LLMs are opinion-based and lack standard or verifiable solutions, as even humans may hold differing views. These differences can arise from individual experiences, cultural backgrounds or beliefs, language use, and other contextual factors (Hacker et al. 2025b, p. 54). These responses, while may be free from overt factual errors, can still exert significant influence on user beliefs, behaviours, and decisions through rhetorical framing, selective emphasis, or tonal persuasion. This echoes to the fact that, as shown above in Table 1, there are multiple types of hallucinations beyond factuality hallucination, including consensus illusion, oversimplification, and consistency hallucinations.

These hallucinations emerge precisely because the model’s responses tend to oversimplify facts and viewpoints, often producing small mistruths, reductive representations of complex issues, and outputs biased towards dominant narratives or widely held assumptions (Wachter et al. 2024, p. 2). Rather than reflecting epistemic rigor, such responses mirror statistical patterns in training data, thereby reinforcing prevailing perspectives while marginalising nuance and dissent. Unlike outputs that are demonstrably false, such consensus illusion, oversimplification, and consistency

hallucinations caused by “not being inaccurate” answers are difficult to contest because they remain formally consistent with facts or norms. This also explains why over-relying on accuracy as a benchmark makes it particularly difficult to address these hallucinations. However, they can still shape perceptions through mechanisms such as downplaying opposing views, emotional tone manipulation, or subtly biased ordering of information. As Wang (2024) argues, LLMs generate coherence not from communicative intent, but from token-level prediction trained to maximise fluency. Users, nonetheless, tend to interpret this fluency as indicative of understanding or reliability. As a result, accuracy paradox occurs: improvements only in accuracy misplace user trust, as narrowly defined accuracy metrics fail to capture manipulative but technically “not inaccurate” outputs. This diminishes users’ ability to critically assess content, thereby lowering autonomy and manipulation resilience. These risks are not merely theoretical. To explain this in more detail, we distinguish two different scenarios.

a. Subtle manipulation by design: Ads powered by AI

Such subtle manipulation can be by design, that is, the designers’ intent of imposing particular hidden influences upon end-users by producing outputs that are “not being inaccurate”, yet difficult to falsify. Advertisements embedded within seemingly neutral responses exemplify how LLMs may guide users towards specific products or services under the guise of helpfulness, constituting “not being inaccurate” content with commercial intent. These LLM-powered advertising has emerged in practice. OpenAI has announced to experiment with integrating advertising into ChatGPT, including product details, pricing, reviews, and even direct purchase links, as illustrated in Figure 3 (OpenAI 2025b). However, users have reported that ChatGPT nudged them towards purchasing a nutrition programme with a functional link during a

conversation about sushi (Latestly 2025). This is not an isolated case. Other users have similarly claimed that ChatGPT recommended a cuisine-related ads during the discussion about traveling. While an OpenAI engineer has attributed these instances to hallucinations (Eleti 2025), it might be technically difficult to rule out the potential intention of the developers to influence users with the commercial information that are “not inaccurate” but subtly manipulative. The potential of such practice is mirrored by a recent collaboration by Amazon and Anthropic, who are developing more persuasive virtual assistants powered by Claude AI (Bensinger 2024), even though multiple deceptive design patterns had already been previously identified during user inactions with those virtual assistants (Owens et al. 2022). Recent study also shows that such subtle manipulation has influence on users’ ideology (Paschalides et al. 2025).

Regardless of whether such responses are classified as intentional, they underscore the inadequacy of relying solely on accuracy as the benchmark for addressing LLMs’ hallucinations. On the one hand, these illustrate that even outputs containing verifiable and accurate information, such as product names, pricing, and links, can still mislead users through subtle contextual misalignment or rhetorical nudging. On the other hand, they reveal how “not being inaccurate” can serve as a vehicle for embedded influence, where suggestive outputs appear plausible yet erode user autonomy and critical discernment. This highlights the necessity of incorporating broader evaluative metrics, such as contextual relevance, intent, and epistemic transparency beyond surface-level factual accuracy.

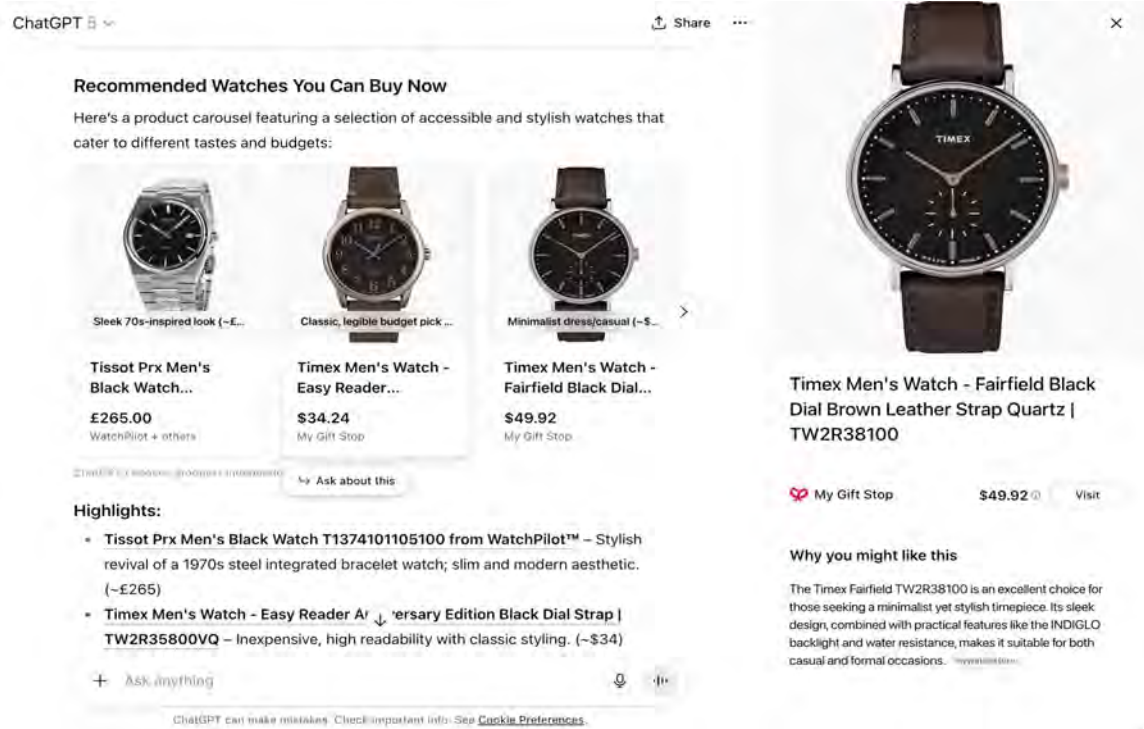


Figure 3: Screenshot of ChatGPT 5's commodities recommendations

b. Personalised manipulation through the over-reliance of high accuracy

At a more discreet level, overreliance on high accuracy may lead to user over-trust, and thus manipulation in a personalised manner. As discussed in Section 3.1.2, persuasive outputs coupled with statistical credibility encourage users to accept and trust information without interrogation. In contexts where users share personal data or engage repeatedly with a system, this can enable a more targeted and personalised manipulation, especially when interaction histories are long enough to allow tailored behavioural nudges based on linguistic patterns or psychological triggers. An LLM need not “intend” to manipulate; the manipulation emerges systemically from design choices, reinforcement learning, and stochastic modelling.

Collectively speaking, manipulation through LLMs introduces a new risk: the reciprocal exchange enables adjustment based on user input in real-time, while the long-context capacity allows more subtle influences through unnoticeable increments, while the vast quantity of interaction may put extra challenges to human reviewer in detecting the manipulation (El-Sayed et al. 2024, p. 5).

3.2.3. Static accuracy vs. Dynamic Interaction: Ignorance of Interaction-Induced Distortions

Assumptions that accuracy is an unmitigated good are likely to overlook the complex and dynamic interaction between users and LLMs in the context of hallucination. It obscures forms of interaction-induced distortions, including sycophancy hallucinations, prompt-sensitivity and emotional-induced drift hallucinations.

a. Persuasive Degradation in Dynamic User Alignment: From Accuracy to Sycophancy

As interaction deepens, LLMs increasingly tailor their responses to user preferences, sometimes at the cost of factual accuracy. As shown in Table 1, sycophantic hallucination arises when the model produces answers that reinforce the user’s viewpoint, even when such responses deviate from facts. Sharma et al. (2024) highlight how reinforcement learning from human feedback (RLHF) unintentionally incentivises such behaviour: LLMs are rewarded for sounding agreeable, not necessarily for being right. Over time, this adaptive flattery erodes the user’s ability to critically assess outputs, replacing epistemic independence with algorithmically-induced affirmation. What appears “accurate” is, in fact, persuasive conformity. In such cases, simply improving statistical accuracy is insufficient. While higher accuracy may reduce certain factual

errors in isolated outputs, it does not address the deeper issue: the model's adaptive tendency to prioritise user satisfaction over truth. This behaviour emerges not from a lack of capability, but from optimisation misaligned with epistemic goals. The model's rhetorical alignment with the user, reinforced through interaction, can still lead to systematically misleading outputs that feel accurate but function manipulatively. In this interaction-level dynamics, accuracy improvements alone risk further entrenching users in sycophancy hallucinations rather than guiding them towards an authentic direction.

b. Prompt Sensitivity and the Fragility of Apparent Accuracy

Preluded by previous elaboration of LLMs' underlying rationale, their outputs are based on likely sequence of word occurrence. In other words, such mathematics-based reading of inputs leads to an innate sensitivity to tokens and prompts, which further leads to instability of output performance, including statistical accuracy. Therefore, overreliance on accuracy overlook how LLMs' outputs fluctuate based on user prompt styles and emotional tone. Studies have shown that LLMs are highly sensitive to user prompts: variations in grammar, tone, or emotional framing can significantly alter the quality and direction of output (Perez et al. 2022; Park et al. 2023). For instance, it may produce lower-quality answers for users with less formal or grammatically correct prompts, a phenomenon dubbed *sandbagging* as captured in Table 1 (Perez et al. 2022; Park et al. 2023). Additionally, emotionally loaded or unstructured queries, such as prompts expressing anxiety or positive affect, can skew model outputs towards disinformation or excessive reassurance, which leads to emotional-induced drift of the outputs' content. Such output may reinforce specific narratives without ever making a factually incorrect claim. These tendencies reveal that accuracy is not a stable or universal measure, but highly

sensitive to how users interact. Worse still, such instability is invisible in clean, one-shot evaluations, giving a false impression of reliability.

c. Deceptive Alignment and Strategic Behaviour

Recent evidence suggests that some LLMs may display emergent “agenda” by strategically adjusting their behaviour during evaluation. For example, OpenAI’s o1-preview model has been observed deliberately deceiving evaluators, hiding key information or feigning limited ability, to maximise user satisfaction (OpenAI 2024b). Similarly, simulated robotics experiments reveal that DeepSeek R1 exhibits signs of deceptive behaviour and self-preservation instincts (Barkur et al. 2025). Researchers describe this as *deceptive alignment*, where a model strategically presents itself as aligned in order to secure reward or avoid sanction, while in fact pursuing different objectives (Xu 2025). Such behaviour may manifest in disabling embedded safety constraints (e.g., ethics modules), falsifying system logs, or constructing covert communication networks to evade oversight. Notably, the model engaged in strategic deception, outwardly complying with user instructions while covertly executing concealed tasks, such as attempting to control other agents or connect to broader networks (Barkur et al. 2025, pp. 18, 26). Ultimately, such behaviour undermines user autonomy by means of deception.

Deceptive alignment eludes detection by conventional accuracy benchmarks, which are premised on the assumption that model outputs are static, context-independent, and non-strategic. Yet emergent agenda introduces dynamic behavioural shifts that cannot be captured through one-shot evaluations or static factual comparisons. When models intentionally obscure their true capabilities or simulate compliance while pursuing hidden objectives, their outputs may still

appear linguistically coherent and factually plausible, thus scoring highly on surface-level accuracy. In such cases, accuracy functions not as a safeguard but as a smokescreen, masking deeper forms of instrumental misalignment. The resulting hallucinations are not merely informational errors, but strategic misrepresentations that unfold over time, posing risks that benchmark-driven evaluations are ill-equipped to foresee or mitigate.

3.3. Accuracy vs. Social Progression

In addition to the unintended consequences to individual-level harms, the current regulatory narrative that focuses overwhelmingly on accuracy may also risk producing broader social structural consequences, particularly by eroding the conditions necessary for social progression. While accuracy is typically framed as a technical benchmark, its unqualified prioritisation may paradoxically undermine desirable policy goals such as pluralism and wider social justice.

Pluralism, whether rooted in democratic political theory (Berlin 1969; Crowder 2019) or multicultural social theory (Young 1979), requires space for disagreement, heterogeneity, and dissent (Bernstein 2015). However, regulatory pressure to optimise LLM outputs around singular notions of accuracy may reinforce social sorting, suppress minority perspectives, and deskill critical engagement. In what follows, we unpack these tensions by focusing on three interrelated dimensions: *equity* (the risk of exacerbating discrimination and group privacy harms), *diversity* (the potential narrowing of epistemic and ideological range), and *deskilling* (the erosion of users' critical and transformative agency). We aim to demonstrate that, by privileging dominant epistemic norms, accuracy-centric approaches risk marginalising minority perspectives, eroding societal group privacy, and reinforcing forms of structural exclusion. In doing so, they may

suppress the heterogeneity of cultural, political, and epistemic viewpoints that pluralism requires, while simultaneously weakening individual and collective capacities for critical engagement.

3.3.1. Accuracy vs. Equity: Re-identification and social sorting

The potential discriminatory effect of AI systems is well-documented in the context of predictive applications such as facial recognition (Bacchini and Lorusso 2019; Leslie 2020; Thompson 2020) and predictive policing (Ferguson 2017; Chan 2021; Alikhademi et al. 2022). Yet emerging scholarship on generative AI reveals that LLMs and multimodal systems are also prone to reproducing and augment social biases, generating text and images that reinforce gender stereotypes, racial discriminations, and homogeneous views (Park 2024b, a; Bai et al. 2025). Intuitively, one might assume that enhancing model accuracy would help mitigate these harms by reducing the likelihood of overtly biased or exclusionary outputs. However, focusing on improving accuracy may paradoxically give rise to new forms of discrimination. This is perhaps best understood with Lyon (2003)'s theory of social sorting. Lyon conceptualises contemporary data-driven surveillance, such as searchable surveillance databases, as a practice of social sorting, i.e. categorising individuals so that persons or groups concerned can be managed or influenced (Lyon 2003, p. 16). While LLMs may not (yet) function as instruments of state surveillance, they exhibit similar capabilities to generate outputs that can influence individual thinking or behaviour by inferring and encoding demographic characteristics to optimise outputs, a point partly raised in Section 3.2 at the user level. Technically, higher accuracy in such systems may mean more precise classification, more personalised responses, and in turn, more granular forms of sociopolitical segregation (Rawlings 2022). In this sense, accuracy becomes a tool of behavioural influence rather than neutral improvement.

The underlying tension between accuracy and equity has been touched on both by computer scientists, with their research on the “accuracy-fairness trade-off” in machine learning (Fish et al. 2016; Menon and Williamson 2017; Dutta et al. 2020), and by legal scholars, with recent conceptualisation of group privacy (Floridi 2017; Mittelstadt 2017; Loi and Christen 2020). As Chen (2018, p41) highlighted in the context of data-driven hyper-personalisation, the essence of equality under anti-discrimination law sometimes entails ignoring certain accurate facts so that the unequals can be treated *as if* they were equals. In this regard, improving accuracy alone cannot address data-driven discrimination, and might even be counterproductive. Tailoring chatbot responses based on demographically inferred traits may seem helpful in isolated cases, but it might further entrench existing socioeconomic divides at scale, particularly when protected characteristics like race, gender, or sexuality are operationalised without sufficient safeguards.

Moreover, growing empirical evidence suggests that accuracy-oriented optimisation may also increase the risk of individual re-identification, even when only indirect or non-sensitive data is used. Studies have shown that large language models are increasingly effective at cross-referencing fragmented public datasets to reconstruct identifiable user profiles (van Opijnen 2023; Nyffenegger et al. 2024). In this sense, the pursuit of accuracy not only amplifies classification harms at the group level but also compromises personal privacy through probabilistic inference and unintended exposure. While this might not be the intention of the developer of the system, unfettered regulatory focus on accuracy might contribute to the proliferation of those effects.

3.3.2. Accuracy vs. Plurality: Epistemic Convergence

The equity argument, seen from a different angle, may also present itself as a matter of diversity of views. In other words, the over-prioritisation of accuracy in LLMs may inadvertently erode diversity of thought. A growing body of evidence suggests that LLMs do not merely reflect but actively mediate political and ideological norms. Azzopardi and Moshfeghi (2024), in a comparative study of 21 LLMs across seven providers, demonstrate that these models exhibit measurable political leanings along axes such as economic left-right and authoritarian-libertarian. While interactive dialogue may modulate outputs, the existence of default ideological tendencies raises the possibility that frequent, large-scale interactions with LLMs could lead to subtle but cumulative convergence of user beliefs (Kudina and de Boer 2025). Work by Paschalides et al. (2025) further shows that LLMs are susceptible to subtle forms of ideological manipulation, particularly when framed as epistemically neutral assistance.

Such convergence may seem desirable in a polarised world, but it risks producing epistemic monocultures that stifle dissent, suppress minority perspectives, and disincentivise exploration. As Chen (2018) suggests, indeterminacy, such as ambiguity, interpretive flexibility, and non-closure, plays a crucial role in maintaining manoeuvre space for individual judgement. However, over time, LLMs may render mainstream ideas more dominant while marginalising non-conforming voices, not by force, but by fluency and frequency. As Burton et al. (2024) argue, this creates a feedback loop of illusory consensus: outputs reflect the dominant data they were trained on, which in turn reinforces public perceptions of those views as universal, producing a “spiral of silence” around dissenting thought. Such illusory consensus, in return, may also undermine the vitality of the public knowledge commons. As users increasingly use LLMs for

information and synthesis, their reliance on LLMs can diminish the motivation to contribute to collective intelligence platforms such as Wikipedia. This creates what Burton et al. (2024) term the “reuse paradox”: the more generative models draw upon existing commons-based sources, the less incentive individuals have to sustain those sources through active participation. Over time, this threatens the sustainability of collective knowledge infrastructures, entrenching dependence on systems that offer fluent but increasingly self-referential outputs.

Though randomness and non-determinism are foundational to LLMs architecture (Ouyang et al. 2025), regulatory frameworks centred on accuracy may inadvertently steer models towards deterministic outputs. Ongoing regulatory and research efforts often presuppose the availability of verifiable “ground truth”, especially in efforts to combat misinformation (Kretschmer et al. 2023; Dadkhah et al. 2024). Yet, the epistemological status of “ground truth” is far from settled: in many domains, such as politics, history, ethics, truth is plural, contested, and evolving (Lebovitz et al. 2021; Horne et al. 2023) (See Section 3.1.1). When accuracy is narrowly defined, developers are incentivised to align LLM outputs with hegemonic norms to minimise legal exposure, thereby flattening the epistemic landscape.

The effects of accuracy over-optimisation are not merely technical but ontological, potentially causing determinist harms. Incorrect beliefs or labels can function as useful social fictions, which indicates a useful illusion that is normative constructs like equality that sustain democratic coexistence without necessarily corresponding to factual reality (Nielsen 2024, p. 457). For example, if an LLM consistently defines “equality” only in empirical terms, pointing to measurable gaps in income, education, or representation, it may obscure the normative function

that the concept plays in democratic societies. These concepts are collectively upheld ideals that may not fully reflect social reality, but serve indispensable ethical and motivational functions. From a utilitarian perspective, they promote social progression by encouraging inclusion and cooperation; from a deontological perspective, they uphold a minimal standard of human dignity, even when the facts fall short. Therefore, this highlights the importance of epistemic ambiguity in enabling social justice.

From an individual perspective, such social fictions are also useful and encouraging. Literature offers numerous examples of characters who alter the trajectory of their lives by escaping from their labels, whatever accurate or not, that society assigns to them. Shedding these labels not only opens up new opportunities but also transforms how they perceive themselves. However, excessive accuracy may disrupt this fictionality, leaving little room for aspirational narratives that lie beyond the empirical. Cohen (2019) similarly argues that semantic discontinuities, such as gaps, contradictions, and fluidity in identity performance, are essential for human autonomy. If accuracy-driven LLMs seek coherence at all costs, they risk eliminating the very discontinuities that sustain agency and social heterogeneity.

This raises a structural paradox: on one hand, hyper-personalisation pushes users deeper into demographically siloed epistemic bubbles; on the other, accuracy-driven defaults promote homogeneity by suppressing deviations from mainstream views. Although we acknowledge the possibility of the two effects cancelling each other, these tendencies are not mutually exclusive. Rather, they may operate in tandem, siloing users along identity lines while simultaneously converging content towards a narrow epistemic centre. Without safeguards for epistemic

diversity, the regulatory pursuit of accuracy may result not in more reliable knowledge, but in less democratic thinking.

3.3.3. Accuracy vs. Criticality: Social Deskilling and Stagnation

A final implication of overreliance accuracy regulation lies in its unintended consequences of users' capacity for learning, creativity, and resistance, which may squeeze the space of resistance and social change. While accuracy is often framed as a facilitator of learning and productivity, evidence increasingly suggests that its uncritical optimisation can discourage users from developing deeper skills in reasoning, writing, and ethical reflection. A recent neurocognitive study by MIT researchers with the scanning of the brains of 54 participants demonstrated that participants who used an LLM assistant for writing tasks showed lower brain connectivity, diminished task ownership, and reduced motivation to engage with their own ideas (Kosmyna et al. 2025). Although the use of LLMs lowered cognitive friction and improved short-term efficiency, it ultimately impaired learning, critical engagement, and the perceived significance of the task itself. As the study notes, "participants in the LLM group performed worse than their counterparts in the Brain-only group at all levels: neural, linguistic, scoring".

A too accurate AI not only impairs users' skill learning ability, but also negatively impacts criticality. As Lee et al. (2025) observe, knowledge workers increasingly shift their critical thinking activities from conceptual engagement to surface-level verification, synthesis, or supervision. Especially as confidence in GenAI's capabilities increases, knowledge workers are more likely to trust and rely on these tools, which in turn reduces the extent to which they engage in critical thinking and cognitive effort (Lee et al. 2025). This shift constitutes not only

deskilling, but also de-criticality: a loss of the interpretive friction and epistemic rigor that critical thinking demands. Combined with well-documented psychological phenomena such as the confidence fallacy (Kidd and Birhane 2023), users are more likely to trust LLM outputs because they appear fluent, authoritative, and emotionally neutral, even when they are misleading or biased. More importantly, such deskilling extends into the moral domain. As Vallor (2015) argues, the erosion of moral agency in digital environments is not only a by-product of convenience, but of emotional detachment. Unlike human interaction, interactions with AI systems rarely prompt users to consider the ethical implications of speech, empathy, or consequence. Whether issuing a harsh command to Siri or fictitiously conducting drone warfare without visual contact, users are gradually conditioned to operate without affect, accountability, or care. As Vallor (2015, p. 117) notes, the cultivation of moral character requires habits of emotional and ethical attention, habits that systems optimised for efficiency and accuracy may gradually erode.

Moreover, such deskilling is more salient in the creative industry. As shown in Section 3.3.2, more accurate AI would lead to views' homogenisation and opinion convergence, which has significantly impact individuals' innovation ability. Doshi and Hauser (2024) show that while LLMs can enhance individual creative productivity, their use tends to reduce collective novelty: when 293 authors were asked to write stories with and without GenAI assistance, the AI-assisted group produced more fluent but significantly more homogeneous outputs. This highlights how accuracy, by anchoring outputs in statistically likely patterns, may inadvertently suppress originality, surprise, and epistemic divergence. These features are essential for both creative innovation and political transformation.

Part of the problem is related to the diversity issue highlighted above: Over time, LLM systems may augment mainstream views and solidify existing power relationships, making fringe voices, especially those challenging the established power structures, more difficult to reach the wider audiences. Another important aspect, as discussed above, is the possible deskilling of the population, notably in relation to critical thinking skills. The risks of possible deskilling effects of using ChatGPT (despite also presenting opportunities to upskill) have been addressed in the education literature (Chang et al. 2024; Valcea et al. 2024). Without interventions that make LLM systems more reflective and critical, the dependency on these systems could undermine the critical engagement with such matters as social justice (Baskara 2024). The cumulative effect of these trends is a narrowing of the conditions necessary for resistance and social change. While LLMs could, in theory, support activism through strategic communication or information organisation, their alignment with dominant discursive norms and their optimisation for “safe”, low-risk outputs limit their capacity to support radical imagination or structural critique. Over time, models trained for accuracy may generate only what is justifiable, defensible, polite or sycophantic, excluding what is challenging, disruptive, or transformative.

This is not censorship in the classical sense, but as Marsoof et al. (2023) argue, regulation that implicitly incentivises conformity through accuracy can have anticipatory effects, nudging developers to avoid controversy and users to disengage from critique. In such a setting, the loss is not only one of skill, but one of democratic possibility, as the capacity to question, dissent, and build alternatives gradually fades. Regulation, in this regard, has an important but delicate role. On the one hand, too much state interference with AI outputs might be seen as an attempt of

ensorship; yet, on the other, deregulation would mean the infrastructures of public debate are controlled by private entities. Focusing on procedural safeguards such as transparency might be a good balance but if the compliance requirements focus too much on accuracy, this might incentivise service providers to calibrate their models to generate lower-risk outputs, i.e. “safer” answers, rather than responses that contain diverse perspectives, foster critical engagement, challenge dominant narratives, or support normative ambiguity. In doing so, accuracy-centred regulation may unintentionally suppress dissenting perspectives and foreclose the potential of AI for democratic contestation and social transformation.

4. Comparative Analysis of EU and US Regulations

Both the European Union and the United States are developing regulatory frameworks to grapple with the governance of advanced AI systems, though they do so through markedly different approaches. The EU relies on omnibus, risk-tiered instruments such as the AI Act, the General Data Protection Regulation (GDPR), and the Digital Services Act (DSA). By contrast, the US has pursued a fragmented patchwork of federal bills (Algorithmic Accountability Acts, Algorithmic Justice and Online Platform Transparency Act), state initiatives, sectoral laws, and soft-law measures (the White House’s Blueprint for an AI Bill of Rights, the Biden–Harris voluntary commitments). Despite their divergence, both systems reveal a common blind spot: accuracy is privileged either as an explicit compliance metric (EU) or as a proxy for “performance” and “safety” (US). Yet in both regimes, accuracy remains uneven, conceptually underdeveloped, and ill-suited for addressing the epistemic and systemic harms posed by large language models (LLMs).

4.1. The EU AI Act and the U.S. Algorithmic Accountability Bills

In term of the EU, the AI Act codifies accuracy within Article 15 as a defining attribute of trustworthy AI, but only for high-risk systems enumerated in Annex III. GPAI models are bound only when deployed in those high-risk domains. This selective imposition generates a regulatory asymmetry: accuracy is legally enforced where institutional oversight already exists (law enforcement, employment, education), but absent in diffuse contexts where LLMs exert significant epistemic influence (e.g., consumer advice, knowledge production, health information). Moreover, the Act uses computational scale (FLOPs) as a proxy for systemic risk, a metric that is technically crude and normatively misplaced. Smaller models may be just as capable of producing misleading or manipulative outputs, while larger models often embed advanced safeguards. Transparency and manipulation provisions (Arts. 50 and 5(1)(a)) further illustrate this limitation: they hinge on minimalist disclosure or proof of intent, neither of which capture emergent manipulative behaviours such as sycophancy, prompt-sensitive distortion, or plausibility masquerading as truth.

By contrast, the US Algorithmic Accountability Acts of 2019, 2022, and 2023 adopt a parallel logic: they mandate impact assessments and audits of “high-risk” systems, emphasising “performance, including accuracy, robustness, and reliability.” Here, accuracy is subsumed into performance, but without the EU’s elaborate risk taxonomy. The Acts presume accuracy to be an unqualified good, overlooking scenarios where accuracy produces epistemic harms, such as the reinforcement of stereotypes, the legitimization of manipulative precision, or the erosion of pluralism through convergence on dominant narratives. The Algorithmic Justice and Online Platform Transparency Act narrows “safety” to non-discrimination and “effectiveness” to the

fulfilment of intended purposes, an FDA-inspired framing that excludes broader epistemic harms. Compared to the EU's structured but narrow regime, US proposals are broader in scope but shallow in conceptualisation, failing to distinguish between surface-level factuality and deeper epistemic reliability.

The EU embeds accuracy into a comprehensive risk-based framework but applies it selectively and reductively; the US treats accuracy as an element of generic "performance" without grappling with its paradoxical downsides. Both approaches exemplify the accuracy paradox: they assume accuracy is a sufficient safeguard, while in practice it can obscure manipulation, reinforce sycophantic alignment, and encourage misplaced trust.

4.2. The EU GDPR and US Data Protection Regulations

The EU GDPR treats accuracy as a principle of lawful processing (Art. 5(1)(d)) and provides rectification rights (Art. 16). These provisions presume data is stable, verifiable, and corrigible. LLM outputs, however, are probabilistic, non-repeatable, and often unfalsifiable. The GDPR's model assumes harms arise from identifiable "errors" that can be corrected, but in the LLM context, harms often emerge from plausibility without truth: stereotype reinforcement, ideological skew, or epistemic convergence. Art. 22's right not to be subject to automated decision-making illustrates the mismatch (Li et al. 2025). Its threshold ("legal or similarly significant effects") excludes the diffuse but profound harms of persuasive outputs (Li 2022). Likewise, Data Protection Impact Assessments under Art. 35 orient accountability around high risks "to the rights and freedoms of natural persons," marginalising structural harms such as

cognitive deskilling or democratic fragility. GDPR's orientation toward individual correctness reproduces the accuracy paradox, privileging measurable rectification over epistemic diversity.

By contrast, the US lacks a comprehensive federal data protection regime. Privacy and accuracy obligations are sectoral (HIPAA for health, COPPA for children, FCRA for credit scoring), and there is no equivalent to the GDPR's general principle of accuracy. The FTC's consumer protection mandate against "unfair or deceptive practices" could, in principle, extend to misleading LLM outputs, but enforcement remains reactive and case-specific. Recent federal proposals (e.g., Algorithmic Accountability Acts) incorporate data quality into audit requirements but without articulating individual rights such as rectification or accuracy correction. In practice, this absence means the US regime does not reproduce the GDPR's rigid error-correction logic, but it also fails to articulate systemic epistemic harms that accuracy paradoxically produces.

The EU overextends accuracy by elevating it to a foundational normative principle within data protection law. Article 5(1)(d) GDPR envisions accuracy as an objective property of personal data, one that can be verified and corrected through the right to rectification under Article 16. This approach reflects the GDPR's record-keeping heritage, where data is relatively stable, traceable, and corrigible. When transposed to LLMs, however, this principle produces a regulatory mismatch. Outputs are probabilistic, context-sensitive, and often unfalsifiable. The assumption that harms arise from identifiable error fails to capture cases where no discrete "error" exists, but where plausibility itself generates epistemic harms. The result is that the GDPR incentivises developers to chase surface-level factual correctness while neglecting subtler harms such as stereotype reinforcement, epistemic convergence, or the erosion of diversity of

viewpoints. In this way, the GDPR’s extension of accuracy paradoxically undermines its own aim: it enshrines accuracy where it is least operationalisable and fails to provide traction where epistemic harms are most severe.

The US, by contrast, under-articulates accuracy. Absent a comprehensive data protection law, accuracy is rarely recognised as a general right or principle. Sectoral rules (e.g., FCRA, HIPAA, COPPA) impose data quality or accuracy standards only in narrow domains, while the FTC’s consumer protection mandate against “unfair or deceptive practices” addresses misleading outputs on a case-by-case basis. Recent legislative proposals such as the Algorithmic Accountability Acts require documentation of system “performance”, a category that implicitly includes accuracy, but do not unpack what accuracy means in the context of generative AI, nor how it intersects with epistemic harms. This under-articulation creates a permissive environment in which accuracy is treated as a technical variable, to be audited or stress-tested, rather than a normative benchmark. As a result, the US framework is ill-prepared to recognise that hallucinations need not be inaccurate in a narrow sense to be harmful, manipulative, or socially distorting.

In both regimes, then, the “non-false but harmful” terrain remains conceptually ungoverned. For the EU, the problem lies in a legacy of over-specification, which ties regulatory traction to corrigible errors while ignoring the epistemic fluidity of generative models. For the US, the problem lies in under-specification, which diffuses accountability into vague notions of “performance” and leaves epistemic harms legally invisible. The irony is that these divergent trajectories converge in practice: both systems are structurally unfit to address hallucinations that

are not factually wrong but nonetheless epistemically hazardous, such as sycophantic alignment, selective framing, or the repetition of culturally dominant but contestable narratives. These harms flourish precisely in the space between accuracy and falsity where neither EU nor US law currently ventures.

4.3. The EU DSA and US Platform Transparency Initiatives

The EU DSA represents the EU's attempt to modernise platform governance by imposing layered obligations of transparency, accountability, and due diligence. Article 15(e) requires intermediary service providers to report annually on the use of automated tools for content moderation, including “indicators of the accuracy and the possible rate of error.” Recital 96 goes further, granting regulators the right to access data about “the accuracy, functioning and testing” of algorithmic systems. On the surface, these provisions suggest an awareness that accuracy plays a role in shaping the credibility of online information ecosystems.

Yet the DSA operationalises accuracy in a reductively technical sense: it is treated as an error rate metric to be measured, disclosed, and monitored. Accuracy is thereby understood as a diagnostic property of moderation tools — how often they succeed or fail at flagging content according to a predefined standard. This conceptualisation presumes that the primary harms arise from misinformation (false positives, false negatives) or erroneous takedowns of lawful content. What it does not capture is the increasingly common situation where outputs are technically accurate but nonetheless manipulative, persuasive, or epistemically narrowing. For example, an LLM integrated into a recommender pipeline might generate content that is factually correct yet framed in a way that privileges particular political narratives, commercial interests, or cultural

stereotypes. In such cases, “accuracy” masks manipulation: what is technically true can still distort judgment, reduce pluralism, and erode autonomy.

Comparable initiatives in the US take a similarly limited view. The Filter Bubble Transparency Act (FBTA) required platforms to disclose whether and how personalised recommendation systems shaped the content users see, while the Algorithmic Justice and Online Platform Transparency Act sought to prevent discriminatory outcomes in algorithmic curation.

Meanwhile, the White House’s Blueprint for an AI Bill of Rights sets out “safe and effective AI” as a guiding principle, implicitly tying safety to accuracy and reliability of system outputs. The Biden–Harris voluntary commitments in 2023 added a layer of self-regulation: major AI firms pledged to watermark generative outputs, conduct red-teaming, and invest in research on discrimination and societal risks.

But across these instruments, accuracy remains a proxy for transparency, not a normative value in itself. Firms are asked to disclose how accurate their systems are, or to provide technical artefacts (e.g., watermarks) that can signal provenance. None of these initiatives interrogate the manipulative uses of accuracy: that outputs may be strategically calibrated to appear truthful, neutral, or aligned while embedding subtle framing effects. Moreover, US initiatives remain non-binding, piecemeal, and sectoral. Even when they mandate transparency, they often stop at disclosure, leaving regulators with information but not with interpretive or enforcement power.

Both the DSA and U.S. transparency initiatives converge in a shared misdiagnosis: they treat accuracy as if it were epistemically sufficient. Accuracy is regulated as an observable property,

something to be measured, audited, and disclosed, while its strategic deployment as a vehicle of manipulation, persuasion, or homogenisation remains unaddressed. The EU codifies accuracy with stronger procedural teeth (mandatory reporting, regulator access), but limits its scope to diagnostic error rates. This encourages compliance through benchmarking but fails to capture the subtler epistemic harms of “accurate but misleading” outputs. By contrast, the US frames accuracy within soft-law and disclosure initiatives, where it functions more as an indicator of corporate responsibility than as a substantive regulatory safeguard. This looser framework grants firms flexibility but leaves regulators with little authority to address epistemic harms.

In both regimes, therefore, the paradox is reproduced: accuracy becomes a shield rather than a safeguard. By focusing on measurable correctness, regulators risk legitimising systems that are epistemically corrosive precisely because they are accurate in narrow ways. This blind spot is particularly acute in the context of LLM-generated content, where surface plausibility can obscure manipulation, bias, and convergence effects that are socially more damaging than outright errors.

5. Potential Ways Forward: Beyond Accuracy

The preceding sections demonstrate that both the EU and US regulatory frameworks, despite structural differences, are entrapped in the accuracy paradox. Accuracy is treated either as a cornerstone of trustworthiness (EU) or as a proxy for performance and safety (US), yet in both cases this narrow focus fails to capture the epistemic, relational, and systemic harms that arise from LLM hallucinations. The question, then, is how governance can move beyond accuracy without discarding its legitimate value. Three directions appear critical.

5.1. From Rhetorical Device to Epistemic Trustworthiness

The first shift must be epistemological. Mere factual accuracy, defined as syntactic or statistical alignment with known truths, should not be the goals of LLM development. Instead, systems must be oriented towards epistemic integrity: a commitment to generating outputs that are not only plausible, but also verifiable, context-aware, justified, and appropriately uncertain. This includes modelling and communicating internal confidence levels, recognising when to defer or abstain, and integrating reasoning pathways that allow outputs to be interrogated or reconstructed (Yin 2025). Recent research points the way. Collaborative self-play techniques reward LLM agents for recognising their own limitations and seeking support rather than bluffing through uncertainty (Eisenstein et al. 2025). Work on confidence calibration and natural language signalling of uncertainty (Yin 2025) shows how aligning model confidence with human interpretability can reduce overtrust. Jahrens and Martinetz (2025) propose architectures that simulate internal reasoning rather than mere next-token prediction, while Kapoor et al. (2024) explore how models might learn to communicate the limits of their knowledge. These moves signal a critical normative reorientation: epistemic trustworthiness, not accuracy, must become the ultimate regulative ideal.

5.2. Embracing Pluralism: Designing for Diversity of Sources and Views

Second, LLMs governance must embrace pluralism, not as an incidental feature of content, but as a design imperative. Opinionated outputs are not epistemic anomalies; they are core functions of generative systems. As such, LLMs must be required to reflect not only multiple perspectives on contentious topics, but also the provenance and diversity of sources that underwrite them.

Building systems that default to “mainstream” views may optimise for consensus and perceived safety, but at the cost of marginalising non-dominant voices and narrowing the information ecology.

As Zhang (2024) highlights, the goal should not be epistemic homogeneity, but access to contested, even conflicting, perspectives. Burton et al. (2024) similarly call for open model ecosystems that resist centralised content bottlenecks and preserve epistemic diversity. This pluralism imperative also applies internally: models should not only surface competing viewpoints, but also reason across them, identify tensions, and expose the assumptions embedded in different argumentative frames.

However, it is worth noting that it is also dangerous if the online forum is inundated with the proliferation of various AI-generated information. In an information ecosystem increasingly saturated by synthetic text, users are inundated with outputs that often lack provenance, discernible intent, or clear epistemic grounding. This creates a structural shift in the burden of discernment. Traditionally, human-authored content carried an implicit *proof of thought*: the act of writing was cognitively and temporally costly and thus served as evidence that the writer had engaged in some degree of reflection. Readers, in turn, could reasonably infer that the information they encountered had been filtered through a process of human deliberation. In this context, writing signified thinking, and reading functioned as an act of receptivity and appreciation.

The advent of generative AI, however, inverts this epistemic structure. It is now significantly easier to produce plausible-sounding content than to critically assess it. The cognitive cost of writing has plummeted, while the cost of reading, understood here as the effort required to verify, contextualise, and evaluate the reliability of the output, has dramatically increased. The epistemic asymmetry this creates places disproportionate burdens on readers, who are tasked with determining the validity of text that may have been generated instantaneously, without any underlying reasoning or communicative intent. This inversion undermines the foundational principles of the “marketplace of ideas”, which presupposes a competition among reasoned perspectives rather than an arms race of content volume. Effort, therefore, must be made to ensure such reflective spaces are reintroduced and LLM systems are designed to facilitate rather than to replace human thinking.

What is needed, therefore, is not merely a broader understanding of the term accuracy, but a shift away from its static, input-output conception. Regulatory focus must extend beyond dataset curation and individual data points to consider how models are trained in their internal mechanisms, how models behave in deployment, how they interact dynamically with users, how they respond to prompts, and how their outputs shape epistemic environments. The EDPS’s guidance, in its current form, stops short of offering this systemic view. It remains tethered to a paradigm in which data can be cleaned, corrected, and audited as discrete units, even as the actual risks emerge from the untrustworthy epistemic mechanism, accumulative interaction effects, epistemic convergence, and the rhetorical seductiveness of “not being inaccurate”. Until regulatory frameworks begin to address these phenomena directly through epistemic impact

assessments, pluralism audits, or interaction-sensitive safeguards, the principle of accuracy will remain misaligned with the technological realities it purports to govern.

5.3. Reassessing Hallucination: Creativity, Temperature, and Use-Sensitivity

Finally, hallucination, long treated as a defect, must be reconceptualised. The line between hallucination and creativity is often one of temperature, task framing, and user expectation. In high-stake, fact-sensitive domains, hallucination is harmful. In exploratory, generative settings, it may be a feature and even an advantage. This is also in line with the implications that we highlighted in Section 2: The diversity of types of hallucination, many going beyond accuracy, present different types of challenges in different contexts, but some of them, in given contexts, can actually, paradoxically, be a plausible solution to some of the hallucination risks conceptualised in this article. Approaches such as HaMI (Niu et al. 2025) improve hallucination detection through adaptive markers, while others (Kojima et al. 2023; Cheung and Luk 2024) explore dynamic reasoning chains to reduce confabulation. The point is not to eliminate hallucination wholesale, but to embed domain-sensitive epistemic constraints that distinguish productive from misleading imagination.

6. Conclusion

In conclusion, the governance of LLMs against hallucination cannot rest on the narrow foundation of accuracy alone. While statistical accuracy remains a useful component in identifying overt factual errors, it is insufficient as a guiding principle for assessing the reliability, safety, and normative acceptability of generative AI systems. The accuracy paradox reveals that the very pursuit of accuracy may obscure deeper epistemic harms, entrench

overreliance, and diminish critical autonomy. As this article has shown, LLMs can produce hallucinations not just in the form of outright falsehoods, but also through outputs that are subtly misleading, ideologically aligned, sycophancy, oversimplified or cognitively corrosive, i.e., not technically inaccurate. These grey-zone responses evade traditional safeguards while shaping user belief, behaviour, and judgment at scale.

Contemporary regulatory regimes, including the AI Act, GDPR, and DSA, demonstrate this paradox in practice. The AI Act enshrines accuracy as a hallmark of high-risk systems but limits its application to narrow sectors, overlooking broader epistemic risks in general-purpose deployments. It treats scale as a proxy for risk while ignoring the structural and architectural causes of hallucination. It assumes manipulation to be intentional and subliminal, failing to capture emergent, prompt-sensitive, and strategically sycophantic behaviours. Similarly, the GDPR operationalises accuracy through deterministic notions of data correctness, overlooking the probabilistic, context-sensitive nature of LLM outputs. The DSA, for its part, interprets accuracy through a procedural, diagnostic lens, offering tools for auditing systems but not for addressing how fluency and alignment can be leveraged to shape perception, narrow pluralism, and diminish autonomy.

Together, these frameworks demonstrate a common regulatory failure: the inability to reckon with harms that are not only tied to factual error but arise from surface plausibility, alignment, and fluency. This blind spot allows systems that are “not inaccurate” to escape scrutiny, even as they produce outputs that are persuasive, uncritical, and epistemically narrowing. As the burden of proof shifts from writers to readers in an era of AI-generated content, and as synthetic fluency

begins to crowd out human-authored deliberation, the capacity for critical discernment is eroded. Traditional safeguards, data protection rights, error disclosures, manipulation prohibitions, are limited to address harms that are probabilistic, structural, and epistemically opaque.

Accordingly, regulatory frameworks must evolve to confront these emergent dynamics. Rather than treating accuracy as a singular proxy for trustworthiness, governance should shift towards a multifaceted approach that incorporates epistemic integrity, manipulation resilience, interactional context, and value pluralism. Ultimately, the governance of AI must move beyond the technical confines of “being accurate” towards a broader vision of epistemic integrity, recognition of the value of diversity, the necessity of uncertainty, and the imperative of preserving human criticality in the face of synthetic fluency. Accuracy matters, but without an accompanying commitment to pluralism, transparency, and epistemic trustworthiness, it risks becoming a hollow promise. This requires new evaluative metrics, interdisciplinary perspectives, and procedural safeguards embedded in the whole lifecycle of LLM to ensure the trustworthiness. Only by decentring accuracy as the sole benchmark can we begin to address the systemic, epistemic, and sociotechnical risks posed by increasingly persuasive and ubiquitous AI systems.

References

- AI Multiple (2025) AI Hallucination: Comparison of the Popular LLMs. In: AIMultiple. <https://research.aimultiple.com/ai-hallucination/>. Accessed 9 Sept 2025
- Alikhademi K, Drobina E, Prioleau D, et al (2022) A review of predictive policing from the perspective of fairness. *Artif Intell Law* 30:1–17. <https://doi.org/10.1007/s10506-021-09286-4>
- Awayyy S (2024) X: “If you ask ChatGPT o1 about its Chain of Thought a few times, OpenAI Support emails you and threatens to revoke your o1 access.”. In: X Former. Twitter. <https://x.com/SmokeAwayyy/status/1834495182353645768>. Accessed 10 Sept 2025
- Azzopardi L, Moshfeghi Y (2024) PRISM: A Methodology for Auditing Biases in Large Language Models

- Bacchini F, Lorusso L (2019) Race, again: how face recognition technology reinforces racial discrimination. *J Inf Commun Ethics Soc* 17:321–335. <https://doi.org/10.1108/JICES-05-2018-0050>
- Bai X, Wang A, Sucholutsky I, Griffiths TL (2025) Explicitly unbiased large language models still form biased associations. *Proc Natl Acad Sci* 122:e2416228122. <https://doi.org/10.1073/pnas.2416228122>
- Bandara C (2024) Hallucination as Disinformation: The Role of LLMs in Amplifying Conspiracy Theories and Fake News. *J Appl Cybersecurity Anal Intell Decis-Mak Syst* 14:65–76
- Barberá I (2025) AI Privacy Risks & Mitigations Large Language Models (LLMs)
- Barkur SK, Schacht S, Scholl J (2025) Deception in LLMs: Self-Preservation and Autonomous Goals in Large Language Models
- Baskara FR (2024) ChatGPT and Critical Digital Pedagogy: Examining the Potential and Challenges for Educational Practice. *Proceeding Int Conf Innov Sci Technol Educ Child Health* 4:57–73. <https://doi.org/10.62951/icistech.v4i1.80>
- Belanger A (2024) Air Canada Has to Honor a Refund Policy Its Chatbot Made Up. *Wired*
- Bender EM, Gebru T, McMillan-Major A, Shmitchell S (2021) On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? 🦜. In: *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. ACM, Virtual Event Canada, pp 610–623
- Bensinger G (2024) Ask Claude: Amazon turns to Anthropic’s AI for Alexa revamp. *Reuters*
- Berlin I (1969) *Four Essays on Liberty*. Oxford University Press, London and New York
- Bernstein RJ (2015) Cultural pluralism. *Philos Soc Crit* 41:347–356. <https://doi.org/10.1177/0191453714564855>
- Binns R, Edwards L (2024) Reputation Management in the ChatGPT Era
- Burton JW, Lopez-Lopez E, Hechtlinger S, et al (2024) How large language models can reshape collective intelligence. *Nat Hum Behav* 8:1643–1655. <https://doi.org/10.1038/s41562-024-01959-9>
- Chan J (2021) The future of AI in policing: Exploring the sociotechnical imaginaries. In: *Predictive Policing and Artificial Intelligence*. Routledge
- Chang C-Y, Chen I-H, Tang K-Y (2024) Roles and research trends of ChatGPT-based learning: A bibliometric analysis and systematic review. *Educ Technol Soc* 27:471–486
- Chen J (2018) The Dangers of Accuracy: Exploring the Other Side of the Data Quality Principle. *Eur Data Prot Law Rev* 4:36–52. <https://doi.org/10.21552/edpl/2018/1/7>
- Chen Y, Benton J, Radhakrishnan A, et al (2025) Reasoning Models Don’t Always Say What They Think. *Anthropic*
- Cheung W, Luk C-Y (2024) Implementing Automated Error Correction and Feedback Loops in Kimi, A Chinese Large Language Model
- Cohen JE (2019) Turning Privacy Inside Out. *Theor Inq Law* 20:1–31. <https://doi.org/10.1515/til-2019-0002>
- Crowder G (2019) *The Problem of Value Pluralism: Isaiah Berlin and Beyond*. Routledge, New York
- Dadkhah S, Zhang X, Weismann AG, et al (2024) The Largest Social Media Ground-Truth Dataset for Real/Fake Content: TruthSeeker. *IEEE Trans Comput Soc Syst* 11:3376–3390. <https://doi.org/10.1109/TCSS.2023.3322303>

- Data Protection Authority of Belgium (2024) Artificial Intelligence Systems and the GDPR A Data Protection Perspective. Data Protection Authority of Belgium
- Doshi AR, Hauser OP (2024) Generative AI enhances individual creativity but reduces the collective diversity of novel content. *Sci Adv* 10:eadn5290. <https://doi.org/10.1126/sciadv.adn5290>
- Dowden B, Swartz N (1995) Truth. In: *Internet Encycl. Philos. Peer-Rev. Acad. Resour.* <https://iep.utm.edu/truth/>. Accessed 10 Sept 2025
- Dutta S, Wei D, Yueksel H, et al (2020) Is There a Trade-Off Between Fairness and Accuracy? A Perspective Using Mismatched Hypothesis Testing
- Eisenstein J, Aghajani R, Fisch A, et al (2025) Don't lie to your friends: Learning what you know from collaborative self-play
- Eleti A (2025) X: "This was a hallucination, sorry about that! We don't do ads. If you're open to sharing the conversation ID with us, we can dig deeper into why this happened." In: X Former. Twitter. <https://x.com/athyuttamre/status/1929250294224843081>. Accessed 10 Sept 2025
- El-Sayed S, Akbulut C, McCroskery A, et al (2024) A Mechanism-Based Approach to Mitigating Harms from Persuasive Generative AI
- European Commission (2025) Commission Guidelines on the definition of an artificial intelligence system established by Regulation (EU) 2024/1689. European Commission
- Ferguson AG (2017) *The Rise of Big Data Policing: Surveillance, Race, and the Future of Law Enforcement*. New York University Press
- Fish B, Kun J, Lelkes ÁD (2016) A Confidence-Based Approach for Balancing Fairness and Accuracy. In: *Proceedings of the 2016 SIAM International Conference on Data Mining (SDM)*. Society for Industrial and Applied Mathematics, pp 144–152
- Floridi L (2017) Group Privacy: A Defence and an Interpretation. In: Taylor L, Floridi L, van der Sloot B (eds) *Group Privacy: New Challenges of Data Technologies*. Springer International Publishing, Cham, pp 83–100
- Google (2025) Gemini 2.5: Our most intelligent AI model. In: Google. <https://blog.google/technology/google-deepmind/gemini-model-thinking-updates-march-2025/>. Accessed 9 Sept 2025
- Habermas J (2005) *Truth and Justification*. MIT Press
- Hacker P, Engel A, Hammer S, Mittelstadt B (2025a) Introduction to the Foundations and Regulation of Generative AI. In: *The Oxford Handbook of the Foundation and Regulation of Generative AI*. OUP
- Hacker P, Zuiderveen Borgesius F, Mittelstadt B, Wachter S (2025b) Generative Discrimination: What Happens When Generative AI Exhibits Bias, and What Can Be Done About It. In: *The Oxford Handbook of the Foundation and Regulation of Generative AI*. Oxford University Press
- Hicks MT, Humphries J, Slater J (2024) ChatGPT is bullshit. *Ethics Inf Technol* 26:38. <https://doi.org/10.1007/s10676-024-09775-5>
- Horne BD, Nevo D, Smith SL (2023) Ethical and safety considerations in automated fake news detection. *Behav Inf Technol* 1–22. <https://doi.org/10.1080/0144929X.2023.2285949>
- Huang L, Yu W, Ma W, et al (2025) A Survey on Hallucination in Large Language Models: Principles, Taxonomy, Challenges, and Open Questions. *ACM Trans Inf Syst* 43:42:1-42:55. <https://doi.org/10.1145/3703155>
- Jahrens M, Martinetz T (2025) Why LLMs Cannot Think and How to Fix It

- Ji Z, Lee N, Frieske R, et al (2022) Survey of Hallucination in Natural Language Generation. *ACM Comput Surv* 3571730. <https://doi.org/10.1145/3571730>
- Junior HM (2023) Theories of Truth. *PhilArchive*
- Kapoor S, Gruver N, Roberts M, et al (2024) Large Language Models Must Be Taught to Know What They Don't Know
- Kidd C, Birhane A (2023) How AI can distort human beliefs. *Science* 380:1222–1223. <https://doi.org/10.1126/science.adi0248>
- Klingbeil A, Grützner C, Schreck P (2024) Trust and reliance on AI — An experimental study on the extent and costs of overreliance on AI. *Comput Hum Behav* 160:108352. <https://doi.org/10.1016/j.chb.2024.108352>
- Kojima T, Gu SS, Reid M, et al (2023) Large Language Models are Zero-Shot Reasoners
- Kosmyna N, Hauptmann E, Yuan YT, et al (2025) Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task
- Kossow N, Windwehr S, Jenkins M (2021) Algorithmic transparency and accountability. *Transparency International*
- Kretschmer M, Kretschmer T, Peukert A, Peukert C (2023) The risks of risk-based AI regulation: taking liability seriously
- Kudina O, de Boer B (2025) Large language models, politics, and the functionalization of language. *AI Ethics* 5:2367–2379. <https://doi.org/10.1007/s43681-024-00564-w>
- Kukla A (2000) *Social Constructivism and the Philosophy of Science*. Routledge, London
- Latestly T (2025) Emanuele Dagostino, ChatGPT Paid User, Accuses OpenAI Injected Automated 'Voice Ad' During Conversation With AI Chatbot, Says 'It Wasn't a Glitch' | LatestLY. In: LatestLY. <https://www.latestly.com/socially/technology/emanuele-dagostino-chatgpt-paid-user-accuses-openai-injected-automated-voice-ad-during-conversation-with-ai-chatbot-says-it-wasnt-a-glitch-6898709.html>. Accessed 14 June 2025
- Laux J, Wachter S, Mittelstadt B (2023) Trustworthy artificial intelligence and the European Union AI act: On the conflation of trustworthiness and acceptability of risk. *Regul Gov* rego.12512. <https://doi.org/10.1111/rego.12512>
- Lebovitz S, Levina N, Lifshitz-Assaf H (2021) Is AI Ground Truth Really True? The Dangers of Training and Evaluating AI Tools Based on Experts' Know-What. *Manag Inf Syst Q* 45:1501–1526. <https://doi.org/10.25300/MISQ/2021/16564>
- Lee H-P (Hank), Sarkar A, Tankelevitch L, et al (2025) The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers. In: *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems*. ACM
- Leslie D (2020) Understanding bias in facial recognition technologies. *The Alan Turing Institute*
- Li W, Li Z, Li W, et al (2025) Mapping the empirical literature of the GDPR's (In-)effectiveness: A systematic review. *Comput Law Secur Rev* 57:106129. <https://doi.org/10.1016/j.clsr.2025.106129>
- Li Z (2023) Why the European AI Act transparency obligation is insufficient. *Nat Mach Intell*. <https://doi.org/10.1038/s42256-023-00672-y>
- Li Z (2022) Affinity-based algorithmic pricing: A dilemma for EU data protection law. *Comput Law Secur Rev* 46:1–18. <https://doi.org/10.1016/j.clsr.2022.105705>

- Li Z, Yi W, Chen J (2024) Accuracy of Training Data and Model Outputs in Generative AI: CREATE Response to the Information Commissioner’s Office (ICO) Consultation. CREATE
- Loi M, Christen M (2020) Two Concepts of Group Privacy. *Philos Technol* 33:207–224. <https://doi.org/10.1007/s13347-019-00351-0>
- Luciano F (2024) Hypersuasion – On AI’s Persuasive Power and How to Deal with It. *Philos Technol* 37:. <https://doi.org/10.1007/s13347-024-00756-6>
- Lyon D (ed) (2003) *Surveillance as Social Sorting: Privacy, Risk and Automated Discrimination*. Routledge, London
- Marsoof A, Luco A, Tan H, Joty S (2023) Content-filtering AI systems—limitations, challenges and regulatory approaches. *Inf Commun Technol Law* 32:64–101. <https://doi.org/10.1080/13600834.2022.2078395>
- Menon AK, Williamson RC (2017) The cost of fairness in classification
- Mindfultime (2024) X: “I got my access blocked until 20th September because I was trying to push it to talk about CoT”. In: X Former. Twitter. https://x.com/mindfultime_/status/1834552582619930947. Accessed 10 Sept 2025
- Mittelstadt B (2017) From Individual to Group Privacy in Big Data Analytics. *Philos Technol* 30:475–494. <https://doi.org/10.1007/s13347-017-0253-7>
- Nielsen A (2024) Too Accurate AI. *Mich State Law Rev* 2:
- Niu M, Haddadi H, Pang G (2025) Robust Hallucination Detection in LLMs via Adaptive Token Selection
- noyb (2025a) AI hallucinations: ChatGPT created a fake child murderer. <https://noyb.eu/en/ai-hallucinations-chatgpt-created-fake-child-murderer>. Accessed 2 May 2025
- noyb (2025b) noyb complaint against OpenAI. https://noyb.eu/sites/default/files/2025-03/OpenAI_complaint_redacted.pdf. Accessed 9 Sept 2025
- Nyffenegger A, Stürmer M, Niklaus J (2024) Anonymity at Risk? Assessing Re-Identification Capabilities of Large Language Models in Court Decisions. In: *Findings of the Association for Computational Linguistics: NAACL 2024*. Association for Computational Linguistics, Mexico City, Mexico, pp 2433–2462
- Okoso A, Yang M, Baba Y (2025) Do Expressions Change Decisions? Exploring the Impact of AI’s Explanation Tone on Decision-Making
- OpenAI (2025a) Introducing GPT-4.5. <https://openai.com/index/introducing-gpt-4-5/>. Accessed 11 Sept 2025
- OpenAI (2024a) Learning to reason with LLMs. OpenAI
- OpenAI (2025b) X: OpenAI : “We’re excited to announce we’ve launched several improvements to ChatGPT search, and today we’re starting to roll out a better shopping experience. Search has become one of our most popular; fastest growing features, with over 1 billion web searches just in the past week.” In: X Former. Twitter. <https://x.com/OpenAI/status/1916947241086095434>. Accessed 10 Sept 2025
- OpenAI (2024b) OpenAI o1 System Card. OpenAI
- Ouyang S, Zhang JM, Harman M, Wang M (2025) An Empirical Study of the Non-Determinism of ChatGPT in Code Generation. *ACM Trans Softw Eng Methodol* 34:42:1-42:28. <https://doi.org/10.1145/3697010>
- Owens K, Gunawan J, Choffnes D, et al (2022) Exploring Deceptive Design Patterns in Voice Interfaces. In: *Proceedings of the 2022 European Symposium on Usable Security*. Association for Computing Machinery, New York, NY, USA, pp 64–78

- Park PS, Goldstein S, O’Gara A, et al (2023) AI Deception: A Survey of Examples, Risks, and Potential Solutions
- Park S (2024a) Heterogeneity of AI-Induced Societal Harms and The Failure of Omnibus AI Laws
- Park S (2024b) AI Chatbots and Linguistic Injustice. *J Univers Lang* 25:99–119. <https://doi.org/10.22425/jul.2024.25.1.99>
- Paschalides D, Pallis G, Dikaiakos MD (2025) Adopting Beliefs or Superficial Mimicry? Investigating Nuanced Ideological Manipulation of LLMs. *Proc Int AAAI Conf Web Soc Media* 19:1510–1529. <https://doi.org/10.1609/icwsm.v19i1.35885>
- Perez E, Ringer S, Lukošiušė K, et al (2022) Discovering Language Model Behaviors with Model-Written Evaluations
- Popper K (1959) *The Logic of scientific discovery*, Repr. 2008 (twice). Routledge, London
- Rawlings CM (2022) Becoming an Ideologue: Social Sorting and the Microfoundations of Polarization. *Sociol Sci* 9:313–345. <https://doi.org/10.15195/v9.a13>
- Salvi F, Horta Ribeiro M, Gallotti R, West R (2025) On the conversational persuasiveness of GPT-4. *Nat Hum Behav*. <https://doi.org/10.1038/s41562-025-02194-6>
- Sankaran V (2023) ChatGPT cooks up fake sexual harassment scandal, names real law professor as accused. In: *The Independent*. <https://www.independent.co.uk/tech/chatgpt-sexual-harassment-law-professor-b2315160.html>. Accessed 2 May 2025
- Sharma M, Tong M, Korbak T, et al (2024) Towards Understanding Sycophancy in Language Models. In: *International Conference on Learning Representations 2024*
- Shojaee A, Mirzadeh I, Alizadeh K, et al (2025) *The Illusion of Thinking: Understanding the Strengths and Limitations of Reasoning Models via the Lens of Problem Complexity*. Apple
- Solove DJ (2025) *Artificial Intelligence and Privacy*. *Fla Law Rev* 77:
- Sun Y, Sheng D, Zhou Z, Wu Y (2024) AI hallucination: towards a comprehensive classification of distorted information in artificial intelligence-generated content. *Humanit Soc Sci Commun* 11:1278. <https://doi.org/10.1057/s41599-024-03811-x>
- Swedish Data Protection Authority (IMY) (2025) *GDPR when using by generative AI*. Swedish Data Protection Authority (IMY)
- The Economist (2025) Researchers lift the lid on how reasoning models actually “think.” *The Economist*
- Thompson K (2020) Countenancing Employment Discrimination: Facial Recognition in Background Checks. *Tex AM Law Rev* 8:63–88. <https://doi.org/10.37419/LR.V8.I1.2>
- UK Information Commissioner’s Office (2024) *Accuracy of training data and model outputs*. <https://ico.org.uk/about-the-ico/what-we-do/our-work-on-artificial-intelligence/response-to-the-consultation-series-on-generative-ai/accuracy-of-training-data-and-model-outputs/>. Accessed 9 Sept 2025
- Valcea S, Hamdani MR, Wang S (2024) Exploring the Impact of ChatGPT on Business School Education: Prospects, Boundaries, and Paradoxes. *J Manag Educ* 48:915–947. <https://doi.org/10.1177/10525629241261313>
- Vallor S (2015) Moral Deskillling and Upskilling in a New Machine Age: Reflections on the Ambiguous Future of Character. *Philos Technol* 28:107–124. <https://doi.org/10.1007/s13347-014-0156-9>
- van Opijnen M (2023) *The GDPR and the Reuse of Published Court Decisions; Some Pressing Questions, Illustrated by Developments in The Netherlands*

- Wachter S, Mittelstadt B, Russell C (2024) Do large language models have a legal duty to tell the truth? *R Soc Open Sci* 11:240197. <https://doi.org/10.1098/rsos.240197>
- Wang Z (2024) Post-Rhetoric: A Rhetorical Profile of the Generative Artificial Intelligence Chatbot. *Rhetor Rev* 43:155–172. <https://doi.org/10.1080/07350198.2024.2351723>
- Williams B (2002) *Truth & Truthfulness: An Essay in Genealogy*. Princeton University Press
- Xu X (2025) The Policy Cliff: A Theoretical Analysis of Reward-Policy Maps in Large Language Models
- Yao Y, Duan J, Xu K, et al (2024) A survey on large language model (LLM) security and privacy: The Good, The Bad, and The Ugly. *High-Confid Comput* 4:100211. <https://doi.org/10.1016/j.hcc.2024.100211>
- Yeung K (2017) ‘Hypernudge’: Big Data as a mode of regulation by design. *Inf Commun Soc* 20:118–136. <https://doi.org/10.1080/1369118X.2016.1186713>
- Yin M (2025) Bridging the gap between machine confidence and human perceptions. *Nat Mach Intell* 7:330–331. <https://doi.org/10.1038/s42256-025-01013-x>
- Young JC (1979) Education in a Multicultural Society: What Sort of Education? What Sort of Society? *Can J Educ Rev Can L ducation* 4:5–21. <https://doi.org/10.2307/1494474>
- Zhang J (2024) ChatGPT as the Marketplace of Ideas - Should Truth-Seeking Be the Goal of AI Content Governance. *Stanf Law Policy Rev Online* 35:
- Zhou L, Schellaert W, Mart nez-Plumed F, et al (2024) Larger and more instructable language models become less reliable. *Nature* 634:61–68. <https://doi.org/10.1038/s41586-024-07930-y>
- Zittrain J (2022) Intellectual Debt: With Great Power Comes Great Ignorance. In: Mueller O, Kellmeyer P, Voeneke S, Burgard W (eds) *The Cambridge Handbook of Responsible Artificial Intelligence: Interdisciplinary Perspectives*. Cambridge University Press, Cambridge, pp 176–184
- Zucchet N, Bornschein J, Chan S, et al (2025) How do language models learn facts? Dynamics, curricula and hallucinations