

Risk-Aware but Unprepared: Business Students, AI-Assisted Writing, and Lessons from a Professional Verification Failure

Li Jiashuo¹, Nanyang, Normal University, School of Economics and Management,
Alan Robert White^{2*}, Rajamangala University of Technology Krungthep, International
College

* Corresponding author. E-mail address: alan.w@mail.rmutk.ac.th

Received: ; Revised: ; Accepted:

Abstract

The integration of artificial intelligence tools into document preparation has created unprecedented opportunities and risks for business professionals and students. This study examines the educational implications of AI-assisted writing through a mixed-methods design combining a quantitative survey of 50 international business students with semi-structured interviews of 10 students, situated alongside a critical examination of a high-profile professional AI failure. Quantitative results show that students hold strong, internally consistent risk-aware perceptions of AI-generated legal business text (Risk Awareness subscale $M = 4.27$, $\alpha = .73$, $t(49) = 6.11$, $p < .001$) and qualified trust in its utility (Trust/Substitution subscale $M = 3.84$, $\alpha = .78$, $t(49) = 3.14$, $p = .003$), while general awareness of AI's capabilities remains only moderate (General AI Awareness subscale $M = 3.47$, $\alpha = .87$, $t(49) = 0.57$, $p = .571$) and substantive reliance on AI for actual legal drafting is significantly and strongly avoided (Substantive Reliance subscale $M = 2.27$, $\alpha = .83$, $t(49) = -8.59$, $p < .001$, $d = -1.21$). Interview data corroborate and explain this pattern: students report deliberate avoidance of AI for high-stakes drafting due to accuracy concerns, alongside near-total absence of structured instructor feedback on AI use. A parallel analysis of the January 2026 West Midlands Police case, in which AI-generated fabricated evidence informed operational policing decisions and precipitated institutional crisis, illustrates the real-world consequences of the verification failures that students' own risk-aware attitudes already anticipate. Together, these findings indicate that the central educational gap is not attitudinal but structural: students possess the disposition toward caution but lack the formal training and institutional feedback loops needed to convert that disposition into reliable verification practice. The paper proposes a framework for AI literacy education integrating technical competency, critical evaluation skill, ethical

reasoning, and legal understanding, and argues that curricula must close this gap before students enter professional contexts where verification failures carry institutional, not merely academic, consequences.

Keywords: Artificial Intelligence, legal business writing, business education, AI literacy, verification protocols, professional ethics, curriculum development, critical evaluation, legal documents, educational technology

Introduction

Artificial intelligence has fundamentally altered the landscape of professional writing, particularly in legal and business contexts where precision, accuracy, and compliance with formal standards are paramount. Tools based on large language models now assist in drafting contracts, generating policy documents, reviewing agreements, and preparing formal business correspondence. These technological developments promise increased efficiency, reduced costs, and democratized access to legal writing capabilities previously requiring specialized expertise (Alarie, Niblett, & Yoon, 2018; Gouri, Jindal, Gouri, & Gouri, 2025), and recent applied research has begun to operationalize these promises directly into deployable systems, from AI-powered legal documentation platforms that automate clause suggestion and compliance checking (Hitesh & Pallavi, 2025) to chatbot-driven document generators aimed at non-experts and small businesses (Atharva-Mahajan, Bhor, Patil, Rane, & Reshim, 2025; Pramanik, R, & R., 2025).

However, the proliferation of AI writing tools has outpaced educational preparation and professional guidelines for their responsible use. Undergraduate business students, who represent the next generation of professionals working with these technologies, often encounter AI tools without structured instruction on their limitations, appropriate applications, or ethical considerations (Pham & Luong, 2025). Universities face the challenge of integrating rapidly evolving technologies into established curricula while maintaining academic integrity standards and preparing students for workplace expectations.

Recent events have dramatically illustrated the risks of inadequate AI oversight in professional contexts. In January 2026, West Midlands Police in the United Kingdom faced severe institutional consequences after an AI tool generated fabricated evidence that influenced operational decisions regarding security measures for a football match. The case resulted in calls for senior leadership resignation and a damning assessment from the Home Secretary (BBC News, 2026; Gohil, 2026; Stephenson, 2026). This incident provides a stark real-world

illustration of the consequences when organizations and individuals fail to implement appropriate safeguards for AI-assisted work in contexts with legal and public-safety implications — precisely the kind of failure that AI literacy education aims to prevent.

This research examines the intersection of AI technology, legal business writing, and educational practice through a quantitative and qualitative study of international business students and an analysis of this documented professional failure. The study addresses a gap in business education literature by exploring how undergraduate students perceive and engage with AI tools for legal writing purposes, and by connecting these educational findings directly to a real-world case of professional AI failure, asking what students' own attitudes suggest about how such failures might be prevented through curriculum design.

Research Objectives

This study pursues four objectives. First, to quantify undergraduate business students' awareness, perceptions, usage patterns, and curricular expectations regarding the use of artificial intelligence in legal business writing. Second, to determine the internal structure of these attitudes through reliability-validated analysis, distinguishing genuinely unified constructs from composites that obscure meaningful internal variation. Third, to explain the quantitative pattern through qualitative interview data that surfaces the reasoning underlying students' attitudes and behavior. Fourth, to connect these educational findings to a documented real-world case of professional AI failure, in order to evaluate whether student attitudes already anticipate the verification failures that produce institutional harm, and to derive concrete recommendations for AI literacy curriculum design.

Research Questions

RQ1: To what extent are undergraduate business students aware of AI's capabilities and limitations in generating legal business documents?

RQ2: How do students perceive the reliability, utility, and ethical implications of AI-generated legal text, and do these perceptions form a single attitude or multiple distinct ones?

RQ3: To what extent do students actually use AI tools for legal writing tasks, and does usage vary by task stakes (e.g., grammar support versus substantive drafting)?

RQ4: What do students believe should change in university curricula and institutional policy regarding AI use in legal writing instruction, and does this differ from a generic desire for more AI exposure?

RQ5: How do the patterns identified in RQ1–RQ4 relate to documented failures of AI oversight in professional legal and administrative contexts, and what do they suggest about curriculum design needed to prevent such failures?

Significance of the Study

This study contributes to business education and legal writing pedagogy in three ways. First, it provides a reliability-validated empirical account of how business students' attitudes toward AI-assisted legal writing are actually structured, correcting a common methodological shortcut in survey-based AI-literacy research in which multi-item sections are treated as unified constructs without verification. Second, by triangulating quantitative and qualitative data, it identifies the specific mechanism — risk-aware caution unsupported by formal training — that explains low substantive usage, moving the literature beyond a general observation that "usage is low" toward an actionable account of why. Third, by connecting this educational evidence directly to a contemporaneous professional failure with documented institutional consequences, the study demonstrates that the curriculum gap identified among students is not a hypothetical concern but one with a real-world analogue already playing out in professional practice, lending urgency and concrete stakes to the paper's curricular recommendations.

Problem Statement

Despite the rapid integration of generative AI tools into legal and business writing practice, business education has not kept pace with structured instruction on their responsible use. Students are left to develop their own informal risk assessments of AI-generated legal content, largely without institutional guidance, feedback, or formal verification training. Simultaneously, documented cases in professional practice — most notably the West Midlands Police case examined in this paper — demonstrate that even trained professionals operating within institutional structures can fail catastrophically when AI-generated content is not adequately verified before being acted upon. This study investigates the gap between what students intuitively understand about AI's risks and what they are formally taught, and argues that closing this gap is a matter of professional readiness with consequences extending well beyond the classroom.

Conceptual Framework

The study's conceptual framework holds that responsible professional use of AI in legal business writing depends on four interacting components: technical competency

(understanding what AI tools can and cannot generate), critical evaluation skill (the ability to assess AI output for factual and legal accuracy), ethical reasoning (judgment about appropriate versus inappropriate reliance, disclosure, and academic integrity), and legal understanding (knowledge of the jurisdictional, regulatory, and professional standards that AI-generated text must satisfy). The study's four-section questionnaire structure maps onto this framework: the Awareness section operationalizes technical competency, the Perceptions section operationalizes critical evaluation and a portion of ethical reasoning, the Usage section operationalizes behavioral translation of these attitudes into practice, and the Implications section operationalizes students' own assessment of curricular and legal-understanding needs. The central conceptual claim tested by the study is that these four components do not develop in lockstep: a student may possess strong critical evaluation instincts (as reflected in the Risk Awareness subscale) while lacking the technical competency or formal training needed to act on those instincts reliably and consistently — a misalignment this study's results substantiate empirically and the West Midlands Police case illustrates at professional scale.

Theoretical Framework

The study draws on four established models of technology acceptance and behavioral intention, detailed fully in Section 3.1. The Technology Acceptance Model (Davis, 1989) frames perceived usefulness and perceived ease of use as the primary determinants of adoption behavior, underpinning the Perceptions and Implications sections. Rogers' (2003) Diffusion of Innovations Theory frames awareness as the necessary first stage of technology adoption, underpinning the Awareness section. Ajzen's (1991) Theory of Planned Behavior frames actual behavior as a function of attitudes, subjective norms, and perceived behavioral control, underpinning the Usage section. Beetham and Sharpe's (2007) digital literacy framework extends this beyond pure technology acceptance into questions of academic integrity and professional readiness, underpinning the Implications section. Together, these models predict that awareness should precede and enable usage, and that perceptions should predict behavioral intention — a predictive chain this study's results partially confirm and partially complicate, since elevated risk perceptions are shown to suppress rather than channel usage in the absence of formal training, a dynamic the standard TAM/TPB chain does not, on its own, anticipate.

Literature Review

AI Technology in Legal Writing Contexts

Alarie, Niblett, and Yoon (2018) provided foundational analysis of how AI technologies are reshaping traditional legal work, including tasks historically performed by experienced attorneys. Their research identified both short-term benefits such as improved efficiency and transparency, and long-term uncertainties regarding the evolution of legal services, emphasizing that future professionals must develop understanding of both the opportunities and constraints of AI-supported legal environments.

Gouri, Jindal, Gouri, and Gouri (2025) examined how AI-generated legal documents are transforming corporate law practice, highlighting substantial efficiency gains and cost reductions alongside significant risks including data security vulnerabilities and diminished human oversight. This research is particularly relevant for business students learning legal writing, as it reflects practical concerns in documentation processes where AI is increasingly embedded.

Hitesh and Pallavi (2025) described the architecture of an AI-powered Legal Documentation System combining natural language processing, machine learning, and intelligent automation to support contract generation, clause suggestion, document formatting, and risk identification. Their account of the underlying technical capabilities provides a concrete referent for what business students mean, or should mean, when they report awareness of AI's capacity to assist in drafting legal documents, and is directly relevant to interpreting the General Awareness subscale findings reported in this study.

Atharva-Mahajan, Bhor, Patil, Rane, and Reshim (2025) presented a legal chatbot and document generator (ReFlaPsy) designed specifically to assist non-experts and small businesses, combining a recommendation engine for document selection with conversational clarification of legal concepts. Their work demonstrates a deliberate design orientation toward users without formal legal training, which corresponds closely to this study's Trust/Substitution subscale finding that students perceive AI as particularly useful for those lacking legal expertise.

Pramanik, (2025) introduced an AI-powered Legal Documentation Assistant combining natural language processing with optical character recognition to generate and personalize legal documents while maintaining regulatory compliance. Notably, their evaluation explicitly identified periodic misinterpretation of legal terminology, dependence on input quality, and jurisdictional restrictions as persistent limitations even in a purpose-built legal AI system —

empirical confirmation, from the system-design side, of precisely the risks that this study's Risk Awareness subscale shows students intuitively anticipate.

Nigam, Patnaik, Thomas, Shallum, Ghosh, and Bhattacharya (2025) introduced a large-scale anonymized dataset of private legal documents and a model-agnostic generation framework for structured legal drafting, demonstrating substantial improvements in factual accuracy, coherence, and completeness over fine-tuned baselines when generation is decomposed into a planning stage followed by retrieval-grounded section drafting. Their work illustrates that the technical capability gap motivating student caution is an active area of correction in the research community, though such improvements remain at the research stage rather than being uniformly deployed in the consumer-facing tools students actually use.

Bunda, Aida, and Utami (2025) developed an NLP-based system to assist legal practitioners in drafting legal claims, reporting reductions in drafting time of up to seventy percent alongside consistent detection of ambiguous sentences and legal errors. Their findings on practitioner-side efficiency gains and error-detection support the present study's Trust/Substitution finding that students perceive genuine utility in AI assistance, even as Risk Awareness findings show this trust remains conditional on human review.

Davenport (2025) investigated the application of large language models to lengthy legal contracts through a case study of an agreement between a school district and a municipal government, combining hierarchical text segmentation with chain-of-thought prompting to address context-window limitations. The study found that large language models could summarize and analyze long contracts with accuracy and efficiency comparable to human legal analysts, while explicitly flagging confidentiality, bias, and hallucination risk as unresolved ethical concerns requiring human oversight — a finding directly relevant to the verification failures examined later in this paper.

Arfat, Colella, and Marella (2024) compared legal-document summaries generated by two large language models on Italian court decisions, evaluating output using standard natural language processing metrics alongside readability scoring, and found meaningful performance differences between systems on semantic similarity measures. Their comparative methodology underscores that AI legal-text quality is not uniform across tools, a nuance largely absent from student survey responses that treat "AI" as a single undifferentiated category, and a point worth incorporating into AI literacy curricula directly.

Mik (2026) offered a sharply critical counterpoint to the efficiency-focused literature, arguing that large language models do not understand legal meaning, lack genuine contextual reasoning, and instead generate statistically plausible approximations of contractual language.

The analysis warned that AI-generated contracts may appear superficially coherent while containing inconsistent or unsuitable provisions for a given transaction, cautioning against treating fluency as a proxy for legal soundness. This perspective is a necessary counterweight to more optimistic accounts of AI-assisted drafting and underscores why student-reported risk awareness reflects a defensible professional instinct rather than excessive caution.

Giampieri (2024) empirically tested this concern by prompting two leading AI systems to draft supply contracts under two different national legal systems. The study found that both systems produced near-identical contracts regardless of the governing legal system specified, disregarding jurisdiction-specific requirements and conventions, and concluded that human supervision remains necessary to correct the resulting legal fallacies. This finding provides direct empirical support for the human-review expectation that emerged as the single strongest attitude in this study's survey data (Q10, $M = 4.46$).

Lam, Cheng, and Yeong (2023) examined combinations of traditional and generative AI techniques for contract drafting, proposing practical prompt-engineering approaches alongside privacy-preserving, locally hosted alternatives to commercial chatbots for use in professional settings where client confidentiality is a concern. Their attention to confidentiality in professional AI use anticipates an ethical dimension largely unprobed by the present study's instrument and is recommended as an addition to future iterations of the questionnaire.

Magesh, Surani, Dahl, Suzgun, Manning, and Ho (2024) conducted the first preregistered empirical evaluation of commercial AI-driven legal research tools marketed as reducing or eliminating hallucination. Despite vendor claims of hallucination-free output, the study found that leading proprietary tools still produced hallucinated content in roughly one-sixth to one-third of tested queries, substantially better than general-purpose chatbots but far from the marketed guarantee. This finding is significant for legal business education because it demonstrates that even professionally marketed, retrieval-augmented systems require active human verification — the same conclusion students in the present study's interviews reached independently and without exposure to this literature.

Siino, Falco, Croce, and Rosso (2025) conducted a systematic literature review of transformer-based language models in the legal domain spanning 2017 to 2023, identifying legal document analysis, case prediction, and contract review as dominant application categories and documenting a marked increase in methodological sophistication and linguistic diversity over the review period. Dehghani, Dehghani, Ardebili, and Rahnamayan (2025) similarly surveyed large language models across legal document drafting, case analysis, compliance monitoring, and legal education, concluding that current models carry meaningful

risks of bias and overstated reliability in high-stakes legal contexts. Both surveys establish that the technical literature on AI in law has matured considerably in scope and rigor even as the risk profile they document remains substantively unchanged, reinforcing the urgency of the educational intervention this paper proposes.

Basha, Kalyani, and Sandeep (2024) presented "Lexi," a generative AI tool designed to simplify legal language and improve the readability of legal documents for non-specialist audiences, framing improved accessibility as a primary value proposition of AI-assisted legal drafting. Their emphasis on accessibility for non-experts parallels this study's finding that students particularly value AI's capacity to clarify difficult legal vocabulary, a theme echoed directly in interview participant S6's account of using AI to manage legal English as a non-native speaker.

Pedagogical Challenges and Student Perspectives

Regalia (2024) presented a practitioner's perspective on how generative AI is redefining legal writing practice, outlining how contemporary AI tools are improving both the speed and quality of legal writing while reshaping approaches to research and communication. Regalia noted that many students and educators lack sufficient understanding of these tools, creating a critical knowledge gap that mirrors the moderate, rather than elevated, General AI Awareness levels found among the present study's participants.

Pham and Luong (2025) conducted research on undergraduate Legal English students and their difficulties integrating AI tools into writing practice, identifying overreliance, ethical concerns, and diminished creativity as common problems. Their research supports the argument that students require structured guidance to use AI responsibly, and that without systematic instruction, students may not develop the judgment needed to assess AI-generated legal text accurately — a concern the present study's near-absence of instructor feedback on AI use (Q15, $M = 1.94$) suggests remains largely unaddressed.

Mika (2022) examined AI-powered features in legal research platforms, detailing how such tools can enhance citation accuracy and support writing, but warned that overreliance on software output could lead students to accept inaccurate or incomplete results without verification. This observation is particularly relevant in educational contexts where students may not yet possess the judgment needed to critique AI-generated content effectively.

Ethical and Professional Considerations

Smith (2023) examined the relationship between plagiarism, legal writing, and language models, arguing that the use of AI-generated content bears similarities to established legal practices of reusing templates and boilerplate language. The analysis challenges traditional assumptions about originality in legal writing and questions the ethical consistency of current academic integrity approaches.

Cyphert (2023) reviewed legal concerns about plagiarism and copyright infringement when lawyers use generative AI, noting that boilerplate copying is itself standard legal practice but that disciplinary bodies vary considerably in their interpretation of AI-assisted drafting. Cyphert's analysis suggests that students must develop awareness of broader intellectual property frameworks before using AI tools in professional contexts.

Harasta, Novotná, and Savelka (2024) conducted experimental research measuring legal professionals' and law students' perceptions of documents believed to be authored by AI versus humans. The study found a clear preference for documents perceived as human-authored over those believed AI-generated, even when content quality was held constant, illustrating persistent professional skepticism toward AI authorship that students entering the profession will need to navigate. Notably, the same study found that most participants nonetheless expected automated legal document generation to become standard in the near future — a tension between present skepticism and anticipated adoption that closely parallels the qualified-trust pattern identified in the present study's Trust/Substitution subscale.

Casagran, Martí, Guillén, and Casagran (2026) reviewed the use of large language models in legal text generation among EU-based law firms, with particular attention to gender implications of algorithmic bias in generated legal text, and offered recommendations for ensuring fairness and reliability in professional deployment. Their attention to bias as a distinct ethical dimension, separate from factual accuracy, identifies a gap in the present study's instrument: the questionnaire's Risk Awareness items address factual and ethical error broadly but do not separately probe bias or fairness concerns, a limitation addressed in Section 5.4.

Synthesis

Across this literature, a consistent tension emerges between AI's demonstrated efficiency gains in legal document tasks (Gouri et al., 2025; Davenport, 2025; Bunda et al., 2025; Basha et al., 2024) and persistent, well-documented risks of factual error, jurisdictional mismatch, bias, and professional skepticism (Mik, 2026; Giampieri, 2024; Magesh et al., 2024; Harasta et al., 2024; Pramanik et al., 2025; Casagran et al., 2026). Educational research on this topic

(Pham & Luong, 2025; Mika, 2022; Regalia, 2024) consistently identifies a knowledge and judgment gap among students using these tools, but — prior to the present study — has not empirically tested whether students' own attitudes already anticipate the risks the technical literature documents, nor connected those attitudes to a real-world institutional failure case. This is the gap the present study addresses.

Methodology

This research employed a mixed-methods design combining a structured quantitative questionnaire with semi-structured qualitative interviews to assess undergraduate business students' awareness, perceptions, usage patterns, and views on the educational implications of AI-assisted legal business writing. The study was conducted at a university with an international Business degree program, with the quantitative and qualitative components conducted concurrently with overlapping but not identical participant pools.

Quantitative Component

Participants. Fifty undergraduate business students ($N = 50$) from China participated in the survey. All participants were studying in an English-medium international program, ensuring adequate language proficiency for questionnaire completion.

Instrument. A structured questionnaire was developed based on the literature review and organized into four sections: (1) Awareness of AI capabilities in legal business writing, (2) Perceptions of AI benefits and limitations, (3) Usage patterns of AI tools for legal and business writing tasks, and (4) Implications for academic learning and professional preparation. Items used five-point Likert scales ranging from strongly disagree (1) to strongly agree (5), with a neutral midpoint (3). The full 20-item instrument is reproduced below.

AI in Legal Business Writing Questionnaire

Section 1 — Awareness: (1) I am aware that artificial intelligence can assist in drafting legal business documents. (2) I know that AI tools such as ChatGPT can be used for legal contract writing. (3) I understand the types of legal business documents (e.g., contracts, policies) that can be generated using AI. (4) I have read or heard about the use of AI in legal or professional writing. (5) I can identify differences between general writing tools and those designed for legal document generation.

Section 2 — Perceptions: (6) AI-generated legal documents may include factual or legal errors. (7) Using AI tools for legal business writing reduces the need for human legal expertise.

(8) Relying on AI tools for legal writing may create academic or ethical problems. (9) AI tools make legal writing easier for students who lack formal legal training. (10) AI-generated legal documents are generally useful but require human review.

Section 3 — Usage: (11) I have used AI tools such as ChatGPT or Grammarly to write legal or business-related assignments. (12) I use AI tools to help improve grammar, vocabulary, and structure in legal documents. (13) I rely on AI tools to draft contracts, terms, or other business documents for coursework. (14) I have copied content directly from AI-generated outputs into legal assignments. (15) I have received feedback on my use of AI tools for legal writing from lecturers or peers.

Section 4 — Implications: (16) AI tools should be included in legal or business writing instruction at university. (17) Students need training on ethical and responsible use of AI in writing legal documents. (18) AI tools may help me prepare better for legal writing tasks in the workplace. (19) University rules about AI use in legal writing are clear and easy to follow. (20) Using AI in legal writing helps me understand legal language and structure better.

Theoretical framework. The questionnaire is grounded in an integrated framework drawing from established models of technology acceptance and behavioral intention. The Technology Acceptance Model (Davis, 1989) provides the primary foundation: the Perceptions and Implications sections reflect TAM's core constructs of perceived usefulness and perceived ease of use as determinants of adoption behavior. Rogers' (2003) Diffusion of Innovations Theory frames the Awareness section, treating knowledge and awareness as the critical first stage through which individuals engage with new technologies. The behavioral dimensions captured in the Usage section draw on Ajzen's (1991) Theory of Planned Behavior, which holds that actual behavior is shaped by attitudes, subjective norms, and perceived behavioral control. Finally, the Implications section incorporates principles from digital literacy and academic integrity discourse, following Beetham and Sharpe's (2007) digital literacy framework for higher education.

Data analysis. Data were analyzed using descriptive statistics, one-sample *t*-tests, and reliability analysis (Cronbach's alpha). One-sample *t*-tests compared section and subscale means against a threshold of 3.40, selected to represent agreement meaningfully beyond the scale's neutral midpoint of 3.0, consistent with conventions in Likert-based technology-acceptance research that distinguish neutral response from substantive endorsement. Prior to hypothesis testing, internal consistency of each originally specified section was assessed; sections falling below conventional reliability thresholds ($\alpha < .60$) were decomposed into conceptually coherent subscales via inter-item correlation analysis, with items not loading

coherently onto any subscale reported individually. Statistical significance was evaluated at $\alpha = .05$.

Ethical considerations. Informed consent was obtained from all participants. Participation was voluntary with no academic consequences for non-participation. Data were anonymized to protect participant confidentiality. Ethics approval was obtained from the relevant university research ethics committee prior to data collection.

Qualitative Component

Participants. Ten international undergraduate business students from China (pseudonyms S1–S10) participated in semi-structured interviews exploring awareness of AI-assisted legal writing tools, usage patterns, perceived benefits, concerns, and educational needs.

Analysis. Interview transcripts were analyzed thematically, yielding six recurring themes detailed in Section 4.2.

Results and Discussion

Quantitative Survey Results (N = 50)

Internal consistency analysis of the originally specified four sections revealed that two sections — Awareness and Perceptions — did not meet conventional reliability thresholds as five-item composites ($\alpha = -.154$ and $\alpha = .484$, respectively), indicating that, as administered, these sections captured more than one underlying construct. Inter-item correlation analysis identified coherent subscales within each: a three-item General AI Awareness subscale (Q1–Q3) and two three-item Perceptions subscales, Risk Awareness (Q6, Q8) and Trust/Substitution (Q7, Q9, Q10), all with acceptable-to-good reliability (Table 1). Items not loading coherently onto either subscale (Q4, Q5 in Awareness; Q12, Q15 in Usage) are reported individually. The Usage section's reliability improved substantially when restricted to its three substantive-reliance items (Q11, Q13, Q14, $\alpha = .834$); the Implications section retained acceptable reliability in its original five-item form ($\alpha = .715$) and required no decomposition.

Table 1. Reliability-validated subscales: descriptive statistics and one-sample t-tests against threshold = 3.40

Subscale	Items	M	SD	α	t(49)	p	d
General AI Awareness	Q1, Q2, Q3	3.47	0.91	.873	0.57	.571	0.08
Perceptions – Risk Awareness	Q6, Q8	4.27	1.01	.730	6.11	<.001	0.86

Subscale	Items	M	SD	α	t(49)	p	d
Perceptions – Trust/Substitution	Q7, Q9, Q10	3.84	0.99	.775	3.14	.003	0.44
Usage – Substantive Reliance	Q11, Q13, Q14	2.27	0.93	.834	–8.59	<.001	–1.21
Implications	Q16–Q20	3.42	0.69	.715	0.16	.871	0.02

Standalone items not forming reliable composites: Q4 (M = 3.62, SD = 0.92), Q5 (M = 3.54, SD = 1.13), Q12 (M = 3.46, SD = 1.45), Q15 (M = 1.94, SD = 1.08).

Table 2. Full item-level descriptive and inferential statistics (test value = 3.40)

Item	Content (abbreviated)	M	SD	t(49)	p
Q1	Aware AI can assist drafting legal documents	3.38	0.85	–0.17	.869
Q2	Know ChatGPT-type tools usable for contracts	3.36	1.19	–0.24	.813
Q3	Understand types of AI-generable legal documents	3.68	0.98	2.02	.048
Q4	Have read/heard about AI in legal/professional writing	3.62	0.92	1.68	.098
Q5	Can distinguish general vs. legal-specific AI tools	3.54	1.13	0.88	.385
Q6	AI-generated documents may contain factual/legal errors	4.40	0.90	7.83	<.001
Q7	AI reduces need for human legal expertise	3.32	1.50	–0.38	.709
Q8	Reliance on AI creates academic/ethical problems	4.14	1.32	3.95	<.001
Q9	AI eases legal writing for those without formal training	3.74	0.94	2.55	.014
Q10	AI-generated documents useful but need human review	4.46	1.05	7.11	<.001
Q11	Used AI for legal/business assignments	2.92	0.99	–3.44	.001
Q12	Use AI to improve grammar/structure	3.46	1.45	0.29	.770
Q13	Rely on AI to draft contracts/terms for coursework	1.92	0.97	–10.84	<.001
Q14	Copied AI output directly into assignments	1.98	1.24	–8.12	<.001

Q15	Received feedback on AI use from lecturers/peers	1.94	1.08	−9.59	<.001
Q16	AI should be included in legal/business instruction	3.14	0.93	−1.99	.053
Q17	Students need ethics/responsible-use training	3.80	1.03	2.75	.008
Q18	AI may help prepare for workplace legal writing	3.48	1.11	0.51	.613
Q19	University AI policy is clear and easy to follow	3.34	0.89	−0.47	.637
Q20	AI use improves understanding of legal language/structure	3.32	1.08	−0.53	.602

Narrative summary. Four findings define the quantitative results. First, Risk Awareness is the strongest and most reliable finding in the dataset ($d = 0.86$): students hold a well-formed, internally consistent belief that AI-generated legal text may be factually flawed and that reliance on it raises ethical concerns. Second, Trust/Substitution beliefs are also significantly elevated but more moderately ($d = 0.44$): students view AI as useful, particularly for those without formal legal training, but do not believe it substitutes for human legal expertise (Q7 alone does not cross the threshold). Third, General AI Awareness does not differ from the threshold ($d = 0.08$): students know AI exists and broadly what it can produce, but this knowledge is moderate rather than confidently held, and the related but distinct skill of discriminating legal-specific from general-purpose tools (Q5) remains similarly moderate and statistically separate from general awareness. Fourth, Substantive Reliance in Usage is the most extreme finding in either direction ($d = -1.21$): students strongly and consistently avoid using AI for actual contract or document drafting in coursework, a pattern sharper than an undifferentiated Usage section would suggest, because it isolates high-stakes drafting behavior from the low-stakes grammar/structure support captured separately in Q12 (which does not differ from the threshold). The standalone item Q15 (receiving instructor or peer feedback on AI use, $M = 1.94$) is the lowest score in the dataset outside the Substantive Reliance cluster, indicating that regardless of how or how much students use AI, structured pedagogical engagement with that use is almost entirely absent.

Within Implications, the composite does not differ from the threshold, but this masks a meaningful split: Q17 (need for ethics/responsible-use training specifically) is significantly

elevated, while Q16 (AI should be included in instruction, generically) is not. Students are not asking for more AI exposure in the abstract; they are asking specifically to be taught how to evaluate and use it responsibly. Two further Implications items round out this picture. Q19 (university rules on AI use in legal writing are clear and easy to follow) sits just above the scale midpoint but does not differ significantly from the 3.40 threshold ($M = 3.34$, $p = .637$), indicating that existing institutional policy, where it exists, is not yet registering with students as clear guidance — a gap that sits uncomfortably alongside the strong, specific demand for ethics training in Q17. Q20 (using AI in legal writing helps me understand legal language and structure better) similarly does not differ from the threshold ($M = 3.32$, $p = .602$), suggesting students do not yet experience their own AI use as pedagogically productive in itself; this is consistent with the Trust/Substitution subscale finding that students value AI primarily as a practical aid rather than as a vehicle for deeper learning, and reinforces that any educational benefit of AI exposure is more likely to come from structured instruction built around its use than from unsupervised exposure alone.

Qualitative Interview Results (n = 10)

Six themes emerged from thematic analysis of the interview data, each corroborating and adding explanatory depth to the quantitative findings above.

Theme 1: Moderate awareness of AI capabilities. Most participants demonstrated general understanding of generative AI tools but incomplete knowledge of their limitations, consistent with the moderate, non-elevated General AI Awareness subscale finding. S1 noted, "I know AI can help write contracts or explain legal terms, but I am not sure if the information is always correct." S7 added, "AI is useful for generating ideas, but I don't fully understand how it creates answers."

Theme 2: Low actual usage of AI for legal writing. Despite general awareness, participants reported limited direct use of AI for legal writing tasks specifically, mirroring the strongly depressed Substantive Reliance subscale. S2 explained, "I rarely use AI for legal assignments because I worry about making mistakes." Only three of ten participants reported using AI more than occasionally for legal-related coursework.

Theme 3: Perceived benefits of AI-assisted legal writing. Participants identified efficiency and accessibility benefits, particularly for non-native English speakers. S6 noted that AI "helps me understand legal vocabulary that is difficult because English is not my first language," while S4 valued AI's capacity to "save time when creating first drafts." These comments align with the significantly elevated Trust/Substitution subscale, particularly Q9,

and resonate directly with the accessibility-oriented design goals described by Atharva-Mahajan et al. (2025) and Basha et al. (2024) in their respective AI legal-drafting tools.

Theme 4: Concerns about accuracy and reliability. All ten participants expressed concern about the trustworthiness of AI-generated legal information, directly corroborating the Risk Awareness subscale, the strongest quantitative finding. S5 observed, "Sometimes AI sounds confident even when the answer may not be correct," and S10 stated, "I would not submit legal work generated entirely by AI without checking it carefully." Students frequently referenced risks of misinformation, fabricated citations, and outdated legal information — concerns empirically validated by Magesh et al.'s (2024) finding that even commercial legal-AI tools hallucinate in a substantial minority of queries, and by Pramanik et al.'s (2025) documented terminology-misinterpretation limitations in purpose-built legal AI systems.

Theme 5: Ethical and professional concerns. Participants recognized risks of overreliance, accountability gaps, and professional consequences. S8 stated plainly, "If lawyers use AI without reviewing the output, it could affect clients negatively" — a comment that, in retrospect, directly anticipates the institutional dynamics of the West Midlands Police case and echoes the confidentiality and accountability concerns raised by Lam et al. (2023) regarding professional AI deployment.

Theme 6: Need for AI literacy education. All participants supported structured AI training in business education, explaining the significant elevation of Q17 against the non-significant Q16. S2 distinguished these directly: "Universities should teach us how to use AI properly, not just tell us not to use it." S9 added that students should learn "both the advantages and risks before using AI in professional situations."

Summary. The interview data converge on a single explanatory mechanism for the quantitative pattern: students avoid high-stakes AI use not from disinterest or lack of access, but from a reasoned, well-articulated risk calculation that current curricula do nothing to formally support, train, or verify. Students' instincts are sound; their training is absent.

Synthesis of Findings

Three results, read together, define this study's central contribution. Students hold strong, well-formed Risk Awareness ($d = 0.86$) that substantially exceeds their General AI Awareness ($d = 0.08$) — they are more confident that AI can fail than they are confident in their own understanding of how or why it works. This caution translates behaviorally into the strongest finding in the dataset: near-total avoidance of AI for substantive legal drafting ($d = -1.21$), even as low-stakes grammar and structural support (Q12) proceeds at moderate,

undifferentiated rates. And this entire pattern unfolds with almost no institutional scaffolding: Q15, the item asking whether students have received any feedback on their AI use from lecturers or peers, registers the second-lowest mean in the entire 20-item instrument ($M = 1.94$), trailing only the substantive-reliance item it is meant to govern. Students are, in effect, self-regulating a high-stakes technology with no formal training and no feedback loop, and so far, their self-regulation has erred toward caution. The open question this study raises, and that the West Midlands Police case answers in the negative, is how long that caution can be relied upon to hold without formal reinforcement once students enter professional environments with deadline pressure, institutional incentives toward efficiency, and far higher stakes than a graded assignment.

The West Midlands Police Case: A Critical Examination

In January 2026, West Midlands Police faced severe institutional crisis following revelations that an AI tool had generated fabricated evidence used in operational decision-making. The incident centered on decisions regarding security measures for a football match between Maccabi Tel Aviv and Aston Villa. According to BBC News (2026), Chief Constable Craig Guildford faced calls for resignation after an interim report identified multiple failures in decision-making processes.

The AI tool produced incorrect information referencing a match between Maccabi Tel Aviv and West Ham United that never occurred. This fabricated evidence was incorporated into decision-making without adequate verification. Initially, Chief Constable Guildford denied that AI was responsible for the erroneous information, but subsequently admitted before a parliamentary statement that an AI tool had indeed provided the false evidence (BBC News, 2026).

Institutional consequences. Home Secretary Shabana Mahmood informed Parliament that she had lost confidence in Chief Constable Guildford, describing the interim inspector's report as damning. The report identified confirmation bias within decision-making processes and concluded that the force had overestimated threats (BBC News, 2026). Local political leaders called for accountability, with Birmingham Council Leader John Cotton stating that Guildford should resign and West Midlands Mayor Richard Parker describing the position as untenable. West Midlands Police issued public apologies acknowledging errors and promising to rebuild confidence, emphasizing that failures were not committed with intent of deliberate distortion or discrimination (BBC News, 2026). However, this assurance did not mitigate institutional damage or prevent calls for leadership changes.

Critical failure points. The incident reveals four failure points directly relevant to this study's findings. First, a fundamental verification failure occurred when AI-generated information was accepted without adequate fact-checking; basic verification would have revealed that no match between the named teams had occurred. This is precisely the failure mode that students in this study's interviews — none of whom had any institutional training on the subject — already articulated as a risk in the abstract (Theme 4; Risk Awareness subscale). Second, the case illustrates confirmation bias, where AI output aligning with existing assumptions receives reduced critical scrutiny; this is a more sophisticated failure mode than simple inattention, and one not directly captured by the present study's instrument, representing a concrete addition for future curriculum design and questionnaire revision (Section 5.4). Third, accountability problems emerged, with initial denial of AI involvement suggesting organizational reluctance to acknowledge reliance on automated systems — an institutional analogue to the academic-integrity concerns raised by Smith (2023) and Cyphert (2023) regarding disclosure of AI assistance. Fourth, inadequate organizational protocols failed to prevent fabricated evidence from influencing high-stakes decisions affecting public safety and civil liberties, the clearest illustration available of what happens when an organization, unlike the risk-aware students surveyed here, lacks even the informal caution this study finds among undergraduates.

Why Student Attitudes Matter for Professional Practice

The juxtaposition is instructive precisely because it runs in the direction one might not expect. The students in this study, entirely without formal AI literacy training, articulated and behaviorally enacted exactly the verification discipline that a trained, institutionally resourced police force failed to apply. This suggests that the deficit current curricula need to address is not one of attitude formation — students already arrive with workable instincts, plausibly shaped by general digital-literacy exposure and media coverage of AI failures (Theme 1; Regalia, 2024) — but of converting those instincts into reliable, trained, institutionally reinforced practice that survives contact with deadline pressure, authority deference, and the efficiency incentives that the West Midlands case shows can override individual caution within an organization. Mik's (2026) and Giampieri's (2024) technical findings on AI's inability to reliably reproduce jurisdiction-specific or transaction-specific legal reasoning, and Magesh et al.'s (2024) documentation of persistent hallucination even in commercial legal-AI products, together supply the evidentiary basis for treating this caution as warranted rather than

excessive, and for building curricula around specific, named failure modes rather than generic admonitions to "be careful with AI."

Limitations

This study has several limitations. The quantitative sample ($N = 50$) and qualitative sample ($n = 10$) are both drawn from a single institution and a single national cohort, limiting generalizability to business education contexts elsewhere. The questionnaire, while reliability-validated post hoc into coherent subscales, was not validated prior to administration through pilot testing or exploratory factor analysis, and future iterations should incorporate this step, along with additional items addressing algorithmic bias and confidentiality (Casagran et al., 2026; Lam et al., 2023), neither of which the current instrument probes directly. The cross-sectional design captures attitudes at a single point in time and cannot establish whether risk awareness changes with exposure to coursework, internships, or direct experience of AI error. Finally, the connection drawn between student survey data and the West Midlands Police case is illustrative and interpretive rather than causal; the case is presented as a real-world analogue that sharpens the stakes of the educational argument, not as a data point statistically linked to the student sample.

Implications for Curriculum Design

The findings support a four-component AI literacy curriculum corresponding to the conceptual framework outlined in Section 1.5. Technical competency instruction should move beyond declarative awareness ("AI can draft contracts") toward functional understanding of how specific tools generate text, including the architectural basis for errors described by Hitesh and Pallavi (2025) and Nigam et al. (2025), so that General Awareness scores reflect operational understanding rather than passive familiarity. Critical evaluation training should build directly on, rather than presume the absence of, students' existing Risk Awareness, providing structured verification protocols, named examples of documented failures such as the hallucination rates reported by Magesh et al. (2024) and the jurisdictional fallacies documented by Giampieri (2024), and the West Midlands Police case itself as a teaching case demonstrating consequences at institutional scale. Ethical reasoning instruction should respond directly to the specific demand expressed in Q17 and Theme 6, addressing disclosure norms, academic integrity boundaries consistent with evolving interpretations discussed by Smith (2023) and Cyphert (2023), and bias considerations raised by Casagran et al. (2026). Finally, legal understanding instruction should close the jurisdictional and regulatory gap that AI tools

themselves cannot reliably close, per Giampieri's (2024) and Mik's (2026) findings, ensuring students understand that fluent AI output is not evidence of legal soundness. This conclusion is reinforced by two further findings that have not yet been addressed. The non-significant result for Q19 indicates that institutional AI-use policy, to the extent it currently exists, is not perceived by students as clear, suggesting that policy clarity itself should be treated as a curriculum deliverable rather than an administrative afterthought; a policy document is only as useful as students' ability to locate, understand, and act on it. Similarly, the non-significant result for Q20 indicates that unsupervised AI exposure alone does not appear to deepen students' understanding of legal language or structure, which argues against curricular approaches that merely permit or encourage AI use without building structured reflection or critique into the assignment itself; the pedagogical value of AI exposure in this domain appears to be conditional on deliberate instructional design, not incidental to use. Across all four components, the near-zero score on Q15 indicates the most urgent and lowest-cost intervention available to institutions: simply establishing a feedback mechanism, of any kind, between instructors and students regarding AI use in legal writing assignments, since currently none meaningfully exists.

Conclusion and Suggestions

This study set out to examine how undergraduate business students understand, evaluate, and use artificial intelligence in legal business writing, and to situate those findings against a real-world case of professional AI failure. The results show a more precise and, in some respects, more encouraging picture than the literature's general framing of student AI use as uncritical or unregulated would suggest. Students in this sample exhibit strong, internally consistent risk awareness and behaviorally enact that awareness by avoiding AI for substantive legal drafting, reserving its use for lower-stakes language support. What they conspicuously lack is not caution but structure: formal technical training, explicit verification protocols, and any meaningful feedback from instructors on how they are currently using these tools. The West Midlands Police case demonstrates, at institutional scale, precisely what happens when this structure is absent even among trained professionals operating within established organizational protocols — fabricated AI output, accepted without verification, that produced confirmation bias, institutional denial, and ultimately the loss of public confidence in senior leadership. The lesson for business education is not that students need to be persuaded AI is risky; the survey data show they already believe this. The lesson is that this belief must be

matched with structured, institutionally supported training before students carry an informal, unreinforced version of this caution into professional environments where the consequences of a single unverified AI output can extend far beyond a single grade.

Several avenues for future research emerge from the findings of this study. First, the present research was limited to undergraduate business students within a single institutional context. Future studies could examine whether similar patterns of AI awareness and usage exist across different universities, disciplines, and cultural settings. Comparative research involving law, accounting, engineering, or information systems students may reveal whether risk perceptions and verification practices vary according to disciplinary training and professional expectations.

Second, while this study identified a strong awareness of AI-related risks, it relied primarily on self-reported behaviors. Future research could incorporate observational methods, experimental designs, or analysis of student-produced assignments to determine whether reported verification practices align with actual behavior when students use AI tools for academic work.

Third, longitudinal studies would be valuable in examining how students' AI competencies evolve over time. Tracking students from university into professional employment could provide insight into whether informal caution develops into systematic verification practices or whether workplace pressures increase reliance on unverified AI outputs.

Fourth, future research should evaluate the effectiveness of specific educational interventions. Experimental studies comparing different forms of AI literacy training, verification frameworks, and instructor feedback mechanisms could identify the most effective approaches for promoting responsible AI use in legal and business writing contexts.

Finally, the inclusion of the West Midlands Police case highlights the importance of examining AI governance beyond educational settings. Future studies could investigate how professionals in legal, business, and public-sector organizations verify AI-generated information, identify barriers to effective oversight, and assess how institutional policies influence trust, accountability, and decision-making. Such research would help bridge the gap between student AI use and the broader organizational challenges associated with integrating generative AI into professional practice.

References

- Ajzen, I. (1991). The theory of planned behavior. *Organizational Behavior and Human Decision Processes*, 50(2), 179–211.
<https://www.sciencedirect.com/science/article/pii/074959789190020T>
- Alarie, B., Niblett, A., & Yoon, A. H. (2018). How artificial intelligence will affect the practice of law. *University of Toronto Law Journal*, 68(supplement 1), 106–124.
<https://utppublishing.com/doi/abs/10.3138/utlj.2017-0052>
- Arfat, Y., Colella, M., & Mareello, E. (2024). Legal text analysis using large language models. In *Proceedings* (pp. 258–268). https://doi.org/10.1007/978-3-031-70242-6_25
- Atharva-Mahajan, A., Bhor, S., Patil, R., Rane, S., & Reshim, P. (2025). Legal chatbot and document generator using ReFlaPsy. *World Journal of Advanced Research and Reviews*.
<https://doi.org/10.30574/wjarr.2025.26.1.1370>
- Basha, T. Y., Kalyani, B., & Sandeep, Y. (2024). Generative artificial intelligence in legal drafting. In *2024 International Conference on Computational Intelligence for Green and Sustainable Technologies (ICCIGST)* (pp. 1–6).
<https://doi.org/10.1109/iccigst60741.2024.10717541>
- BBC News. (2026, January 15). West Midlands Police chief faces calls to quit over Maccabi Tel Aviv ban. <https://www.bbc.com/news/live/c394zlr8e12t>
- Beetham, H., & Sharpe, R. (2007). *Rethinking pedagogy for a digital age: Designing and delivering e-learning*. Routledge.
https://www.academia.edu/download/110680527/Learning_and_e-learning.pdf
- Bunda, R. H., Aida, W., & Utami, U. (2025). Digital reform in legal complaint drafting: Construction of an AI-based planning system to reduce the risk of obscure interim decisions. *Awang Long Law Review*. <https://doi.org/10.56301/awl.v8i1.1766>
- Casagran, C. B., Martí, L. B., Guillén, S. R., & Casagran, E. B. (2026). Navigating the AI legal landscape: Gender implications of large language models in legal text generation. *IDP. Revista de Internet, Derecho y Política*. <https://doi.org/10.7238/idp.v0i44.432721>

Cyphert, A. B. (2023). Generative AI, plagiarism, and copyright infringement in legal documents. *Minnesota Journal of Law, Science & Technology*, 25, 49. <https://scholarship.law.umn.edu/cgi/viewcontent.cgi?article=1564&context=mjlst>

Davenport, M. J. (2025). Enhancing legal document analysis with large language models: A structured approach to accuracy, context preservation, and risk mitigation. *Open Journal of Modern Linguistics*, 15(2), 232–280. <https://doi.org/10.4236/ojml.2025.152016>

Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly*, 13(3), 319–340. <https://misq.umn.edu/misq/article-abstract/13/3/319/191>

Dehghani, F., Dehghani, R., Ardebili, Y. N., & Rahnamayan, S. (2025). Large language models in legal systems: A survey. *Humanities and Social Sciences Communications*, 12. <https://doi.org/10.1057/s41599-025-05924-3>

Giampieri, P. (2024). AI-powered contracts: A critical analysis. *International Journal for the Semiotics of Law*, 38, 403–420. <https://doi.org/10.1007/s11196-024-10137-z>

Gohil, N. (2026). West Midlands police chief apologises after AI error used to justify Maccabi Tel Aviv ban. *The Guardian*. <https://www.theguardian.com/uk-news/2026/jan/14/west-midlands-police-chief-apologises-ai-error-maccabi-tel-aviv-ban>

Gouri, H., Jindal, P., Gouri, R., & Gouri, V. (2025). AI-generated legal documents: Challenges and opportunities for corporate law firms. *Exploration of AI in Contemporary Legal Systems*, 73–84. <https://www.igi-global.com/chapter/ai-generated-legal-documents/365944>

Harasta, J., Novotná, T., & Savelka, J. (2024). It cannot be right if it was written by AI: On lawyers' preferences of documents perceived as authored by an LLM vs a human. *Artificial Intelligence and Law*, 1–38. <https://www.muni.cz/en/research/publications/2417778>

Hitesh, D., & Pallavi, D. (2025). AI-powered legal documentation system. *International Journal of Scientific Research in Engineering and Management*. <https://doi.org/10.55041/ijssrem47956>

- Lam, K.-Y., Cheng, V. C. W., & Yeong, Z. K. (2023). Applying large language models for enhancing contract drafting. In *Proceedings* (pp. 70–80). <https://ceur-ws.org/Vol-3423/paper7.pdf>
- Magesh, V., Surani, F., Dahl, M., Suzgun, M., Manning, C. D., & Ho, D. E. (2024). Hallucination-free? Assessing the reliability of leading AI legal research tools. *arXiv*. <https://doi.org/10.48550/arxiv.2405.20362>
- Mik, E. (2026). Contractual deepfakes: Can large language models generate contracts? *arXiv*. <https://doi.org/10.48550/arxiv.2602.09384>
- Mika, K. (2022). Friend or foe? Lexis artificial intelligence (AI) in legal writing. *Proceedings: Online Journal of Legal Writing Presentations*, 3(1), 24. https://engagedscholarship.csuohio.edu/fac_articles/1261/
- Nigam, S., Patnaik, B. D., Thomas, A. V., Shallum, N., Ghosh, K., & Bhattacharya, A. (2025). Structured legal document generation in India: A model-agnostic wrapper approach with VidhikDastaavej. *arXiv*. <https://doi.org/10.48550/arxiv.2504.03486>
- Pham, T. T. D., & Luong, M. H. (2025). Students' challenges in employing AI tools in legal writing. *International Journal of TESOL & Education*, 5(2), 75–86. <https://i-jte.phamho.com/index.php/journal/article/view/526>
- Pramanik, S., Sahana Reddy R, Deepika R. (2025). Intelligent legal documentation: Leveraging AI for precision and productivity. In *2025 2nd International Conference on New Frontiers in Communication, Automation, Management and Security (ICCAMS)* (pp. 1–6). <https://doi.org/10.1109/iccams65118.2025.11233958>
- Regalia, J. (2024). From briefs to bytes: How generative AI is transforming legal writing and practice. *Tulsa Law Review*, 59, 193. <https://digitalcommons.law.utulsa.edu/cgi/viewcontent.cgi?article=3322&context=tlr>
- Rogers, E. M. (2003). *Diffusion of innovations* (5th ed.). Free Press. <https://www.scirp.org/reference/referencespapers?referenceid=3275637>

Siino, M., Falco, M., Croce, D., & Rosso, P. (2025). Exploring LLMs applications in law: A literature review on current legal NLP approaches. *IEEE Access*, 13, 18253–18276. <https://doi.org/10.1109/access.2025.3533217>

Smith, M. L. (2023). Language models, plagiarism, and legal writing. *University of New Hampshire Law Review*, 22, 361. https://scholars.unh.edu/unh_lr/vol22/iss2/9/

Stephenson, M. (2026). How AI destroys policing institutions. *Policing Insights*. <https://policinginsight.com/feature/analysis/how-ai-destroys-policing-institutions/>