

An initial foray into AI judicial drafting: a comparative experiment in a constitutional privacy case

Kudakwashe Mawowa ¹ {0009-0002-5880-838X}, Donrich Willem Thaldar¹ {0000-0002-7346-3490}

¹ School of Law, University of KwaZulu-Natal, Durban, South Africa

Abstract. This empirical study investigates how well can Artificial Intelligence (AI) draft a South African High Court judgment using OpenAI’s ChatGPT-4o large language model and case data from *De Jager v Netcare Limited and Others* 2025 JDR 0793 (GP)—an informational privacy and data protection dispute. Drawing from local and international experiments, the study develops a value-driven prompting framework and distills a rubric for evaluating the quality of AI-driven judgments. Twenty sequential prompts were developed to generate an artificial judgment which was then compared with the human decision with respect to structure, factual accuracy, use of legal authority, legal reasoning and clarity. ChatGPT-4o led across all five metrics, though formatting inadequacies were a weakness. This study did not assess the substantive legal correctness of either judgment. However, it is interesting that both judgments reached nearly identical outcomes on the main substantive legal issues, viz privacy, statutory compliance (Protection of Personal Information Act 4 of 2013) and admissibility of surveillance data. This may be attributable to their common reliance on similar submissions, suggesting that persuasive pleadings can channel both human and machine reasoning to similar ends. These findings support the potential of AI as a judicial aid but reinforce the need for robust human oversight.

Keywords: Artificial Intelligence, Judicial Drafting, AI Judgment, Legal drafting

1 Introduction

“If justice delayed is justice denied, then what of justice automated?”

In an era marked by digital acceleration and institutional strain, courts around the world are exploring how Artificial Intelligence (AI) might assist with the burdens of adjudication. Large language models (LLMs) like OpenAI’s Generative Pre-trained Transformer (ChatGPT)¹—trained on a vast corpora of legal and general texts—are now widely used for research, drafting documents, and summarising tasks in legal

¹ See OpenAI “ChatGPT Overview” <https://openai.com/chatgpt/overview/> 2025 (last accessed 2025-03-24).

practice.² But can such a model go further? Can it draft a legal judgment of sufficient quality to meet the standards of a real court?

This article takes up that challenge by empirically evaluating the performance of ChatGPT-4o—the flagship large language model (LLM) from OpenAI—³ on a real South African High Court case: *De Jager v Netcare Ltd.*⁴ While the scope of a single case study is narrow, it does provide some useful initial insights into AI-driven judicial work future expanded experiments can build upon. The court case entailed constitutional privacy rights, statutory data protection law, and evidentiary admissibility. We asked ChatGPT-4o to write a full judicial opinion using real court filings from the case and a carefully constructed 20-step prompting protocol. We then compared its output to the actual judgment handed down by the presiding officer Motha J.

The stakes of this inquiry are high. In South Africa, judicial authority is constitutionally vested in the courts, and no person or organ of state may interfere with their functioning.⁵ Any suggestion that a machine might perform judicial reasoning raises profound constitutional, ethical, and institutional questions. Dismissing the potential of generative AI altogether would ignore the reality that such tools are already entering legal practice—including behind the bench.⁶

This article makes three contributions. First, it offers the first documented South African experiment in using an LLM to write a full-length legal judgment based on real case materials. Second, it develops a value-driven, replicable methodology for prompting and evaluating AI-generated judgments, grounded in local jurisprudence, judgment-writing pedagogy, and constitutional norms. Third, it provides a comparative analysis of the strengths and limitations of AI-generated judicial writing, identifying areas where such tools may support judicial work—particularly in drafting, accessibility, and structural clarity—and where they fall short, particularly in legal reasoning, citation, and contextual understanding.

Our central argument is that while ChatGPT-4o can generate judgments that are facially plausible and structurally coherent, it cannot yet replace human judicial reasoning. However, it holds significant promise as a judicial drafting assistant—a tool to accelerate workflow, support multilingual access, and enhance the transparency and traceability of the judgment-writing process. We conclude that South African courts should treat AI-generated outputs analogously to research clerks' memos: useful, suggestive, but never dispositive.

In the sections that follow, we review the legal and technological context, describe the experimental methodology, outline our evaluation criteria, and analyse the results.

² See Galarreta, Francisco Javier Fernández, "The Impact of Artificial Intelligence Technologies on the Justice Administration and on the Judicial Office Personnel" *Zeitschrift für Rechtssoziologie*, vol. 45, no. 1, 2025, pp. 103-121.

³ See OpenAI "ChatGPT-4o" <https://platform.openai.com/docs/models/gpt-4o> 2025 (last accessed 2025-03-13).

⁴ *De Jager v Netcare Limited and Others* 2025 JDR 0793 (GP).

⁵ S165 (1) of The Constitution of the Republic of South Africa, 1996.

⁶ Galarreta, Francisco Javier Fernández, "The Impact of Artificial Intelligence Technologies on the Justice Administration and on the Judicial Office Personnel" *Zeitschrift für Rechtssoziologie*, vol. 45, no. 1, 2025, pp. 103-121.

An initial foray into judicial drafting: a comparative experiment in a constitutional privacy case

In doing so, we aim not to speculate, but to provide evidence-based guidance for how courts might pragmatically and constitutionally embrace AI as part of their evolving institutional toolkit.

2 Background and literature context

The integration of artificial intelligence (AI) into legal systems has shifted from theoretical curiosity to practical inevitability. Around the world, courts and legal practitioners are experimenting with large language models (LLMs) to streamline administrative tasks, conduct legal research, and in some instances, assist with decision-making. In Colombia, for example, a judge openly cited ChatGPT-3.5 in a judicial ruling, sparking international debate on transparency, accountability, and the limits of AI-generated reasoning.⁷ In South Africa, legal firms such as Schindlers have tested proprietary AI tools to simulate adjudication, and academic scholars have begun probing the normative boundaries of such use.⁸ Yet, the question remains: can generative AI ever satisfy the institutional and constitutional requirements of judicial authority?

The South African Constitution is unequivocal: judicial authority is vested in the courts and must be exercised independently and without fear, favour, or prejudice. Section 165(4) of the Constitution obliges the state to take “legislative and other measures” to support the independence and effectiveness of the judiciary. Whether generative AI may be considered such a “measure” raises urgent legal and ethical questions, particularly where automation could undermine impartiality, accountability, or the right to reasons. “To be democratically legitimate, rules and judgments should be the outcome of popular sovereignty, something over which the people exercise authorship.”⁹

Hicks et al. argue that generative LLMs are optimized for likelihood rather than likeness to reality and are trained to produce plausible-looking text, not to represent the world accurately so their errors (and even their correct answers) should be understood as “bullshit” outputs unconcerned with truth.¹⁰ This analogy underscores that accuracy is incidental under a truth-indifferent training regime because correctness is statistically emergent rather than epistemically grounded. This paper argues that assuming inaccuracy at the macro level may be a misstep if the goal is to guide responsible AI deployment. Such pessimism is, by definition, indifferent to AI as an assistive judicial tool—

⁷ See First Circuit Court of Cartagena, Case No. 11001-22-2023-00016.

⁸ See Schindlers Attorneys, ‘Sarah – Schindlers’ Adjudication AI Engine Hands Down Judgment in Favour of the Legalisation of Psilocybin (Psychedelic Mushrooms) in South Africa’ 22 April, <https://www.schindlers.co.za/sarah-schindlers-adjudication-ai-engine-hands-down-judgement-in-favour-of-the-legalisation-of-psilocybin-psychedelic-mushrooms-in-south-africa/> (last accessed 2025/05/15).

⁹ Kiel Brennan-Marquez and Stephen Henderson, *Artificial Intelligence and Role-Reversible Judgment*, 109 J. CRIM. L. & CRIMINOLOGY 137 (2019) p152.

¹⁰ Hicks, M.T., Humphries, J., & Slater, J. (2024). *ChatGPT is bullshit. Ethics and Information Technology*, 26, 38. <https://doi.org/10.1007/s10676-024-09775-5>.

one that would serve to support, rather than to make final binding decisions. While AI in law may on the surface appear epistemically hazardous, independent judicial authority requires that its outputs be validated by humans, just as one would evaluate the contributions of a research assistant.

Despite these high stakes, there is little empirical research in the South African context assessing how AI tools perform when applied to the task of judgment writing. Most local commentary to date has focused on AI in legal practice more broadly—including contract generation, predictive analytics, and document review. For instance Petse and Phindelo deemed ChatGPT a deceptive tool that threatens courts and the administration of justice mainly on the basis of three cases involving lawyers who submitted unverified AI-written submissions.¹¹ Their analysis treated a few instances as proof of systemic danger despite empirical studies proving that error-rates can be significantly reduced when users apply best practices¹²—such as supplying sources using Retrieval Augmented Generation (RAG).¹³ RAG entails prompting a large language model to generate an answer based only on supplied documents anchoring the response of the model to those documents.¹⁴ In the three cases, the courts were critical of the lawyers for copying ChatGPT in bad faith but did not condemn ChatGPT itself as an assistive tool.¹⁵ Thaldar also demonstrated that ChatGPT-4 can assist in drafting Data Transfer Agreements for health research but found its outputs uneven and in need of expert revision.¹⁶ Similarly, Lefakane explored the potential for AI to alleviate burdens on South Africa’s overextended judiciary, concluding that any implementation must be directed by constitutional values and subject to rigorous oversight—but did not test AI performance on a real case.¹⁷

Internationally, researchers have experimented with LLMs by posing legal problems from a variety of jurisdictions. Brojde tested ChatGPT-4’s ability to respond to complex halachic legal questions, finding that while the model could apply precedent and analyse texts, it struggled with contextual reasoning and ethical nuance.¹⁸ In a German study, Schweitzer and Conrads found that ChatGPT-4 could resolve basic business law

¹¹ I Petse & U Phindelo ‘The Influence of ChatGPT-generated Data on the Administration of Justice in South Africa’ (2025) 5 *Turf Law Journal* 1-14.

¹² B Tomlinson, AW Torrance & RW Black ‘ChatGPT and Works Scholarly: Best Practices and Legal Pitfalls in Writing with AI’ (2023) 76 *SMU L Rev Forum* 108.

¹³ See Reece Rogers ‘Reduce AI Hallucinations With This Neat Software Trick’ 2024 https://www.wired.com/story/reduce-ai-hallucinations-with-rag/?utm_source=chatgpt.com (last accessed 2025/06/12).

¹⁴ *Ibid.*

¹⁵ See *Mavundla v MEC: Department of Co-Operative Government and Traditional Affairs KwaZulu-Natal and Others* [2025] 2 ZAKZPHC para 37; *Parker v Forsyth* (2023)1 1585/20 ZAGPRD para 89-90; *Mata v Avianca Inc* 22-cv-1461 (PKC) United States District Court, S.D. New York (2023).

¹⁶ See DW Thaldar ‘How effectively can ChatGPT-4 draft data transfer agreements for health research?’ (2025) *Humanit Soc Sci Commun* 12, 524.

¹⁷ N Lefakane ‘The Use of Artificial Intelligence for Adjudication in South Africa’ LLM dissertation, North-West University, 2022.

¹⁸ See Brojde & Michael J ‘AI and Jewish Law: Seeing How ChatGPT 4.0 Looks at a Novel Issue’ (2023) Buchman Law School, Tel Aviv University, and Cardozo Law School, Yeshiva University, *CSLR Research Paper No. 12.2023-AFF*.

An initial foray into judicial drafting: a comparative experiment in a constitutional privacy case

disputes but lacked sophistication in edge cases.¹⁹ Collectively, these studies suggest that LLMs can simulate legal reasoning under constrained conditions but remain dependent on careful prompting, verification, and human correction.

African perspectives remain limited in this global discourse. While the 2025 Africa AI Privacy Report promotes AI use in transcribing rulings and managing caseloads, it does not substantively address whether AI can, or should, write judgments.²⁰ Nor does it engage with the jurisprudential traditions—such as ubuntu or restorative justice—that shape African legal reasoning and which may not be easily modelled by non-African AI systems trained primarily on Euro-American corpora.²¹

This article seeks to fill that gap. By subjecting ChatGPT-4o to a demanding, real-world test using South African case law, procedural rules, and constitutional values, we explore both its technical capabilities and its jurisprudential limitations. Our findings contribute to a grounded, context-sensitive understanding of AI's role in judicial writing—and offer practical insight for courts, academics, and policy-makers navigating the promises and perils of automation in the judiciary.

3 Methodology

This study employed an empirical case study approach to explore the capabilities of ChatGPT-4o, a large language model developed by OpenAI, in drafting a High Court judgment for a South African case, *De Jager v Netcare Ltd.*²² Selected for its doctrinal complexity and constitutional significance, the case concerned the admissibility of surveillance evidence in a delictual damages claim, raising issues under privacy provision Section 14 of the Constitution of the Republic of South Africa, 1996, and the Protection of Personal Information Act (POPIA).²³ By comparing the AI-generated judgment to the human-authored judgment by Motha J, the study assessed formal attributes—structure, factual accuracy, use of legal authority, legal reasoning, and clarity—as outlined in the evaluation framework (Section 4) below. The methodology prioritised transparency and alignment with South African judicial standards, focusing on the AI's ability to produce a plausible draft judgment while adhering to ethical and constitutional principles.

3.1 Case selection

The case, heard in the Gauteng Division of the High Court, Pretoria, involved a plaintiff claiming approximately R25.7 million in damages from a healthcare provider for negligent surgery. The central legal issue was whether the defendant's covert

¹⁹ S Schweitzer & M Conrads 'Correction to the digital transformation of jurisprudence: an evaluation of ChatGPT-4's applicability to solve cases in business law' (2024) *Artif Intell Law*.

²⁰ See Africa AI Policy Lab *Africa AI Privacy Report 2025* (2025) available at <https://www.aipolicy.africa/africa-intelligence-report> (last accessed 15 April 2025/04/15) 67, 73–74.

²¹ *Ibid.*

²² *De Jager v Netcare Limited and Others* 2025 JDR 0793 (GP).

²³ Act 4 of 2013.

surveillance—which incidentally captured the plaintiff’s minor grandchildren—violated privacy rights under Section 14 and POPIA, and whether the evidence was admissible. The case’s reliance on written submissions, including those from amici curiae (Professors Thaldar and Snail ka Mtuze), and its well-defined legal questions made it an ideal test case for AI judgment drafting. Two published judgments—a procedural ruling (2024)²⁴ and the substantive judgment (2025)²⁵—provided the basis for comparison, supporting the study’s exploratory aim to evaluate the AI’s formal drafting ability against a human judgment.

3.2 Large language model: ChatGPT-4o

Released in May 2024,²⁶ ChatGPT-4o was selected for its advanced multimodal capabilities and strong contextual coherence, despite lacking dedicated training in South African law.²⁷ ChatGPT-4o’s knowledge cut-off date is October 2023.²⁸ The experiment was conducted in a closed system using the RAG approach. The model relied exclusively on the provided case materials and prompts to generate its output whilst restricted from internet access. This simulated realistic judicial drafting conditions and minimised external influence. Newer models have since been released i.e. GPT-5. Methodology applied in this case study may not necessarily be replicable.

3.3 Prompting protocol

A structured 20-prompt and responses sequence (fully available in citation) was used to guide ChatGPT-4o in drafting the judgment.²⁹ This sequence emulated a High Court judge’s process of synthesising legal submissions to ensure effectiveness in accordance with effective prompting practice.³⁰ While the sequence included preparatory prompts (e.g., for document analysis), the central prompt instructed the model to produce a full-length judgment, using explicit parameters to ensure adherence to judicial standards in South Africa (discussed under subheading 4. Formulating evaluation criteria). The first and main prompt, read as follows:

“Prompt 1: Hello, ChatGPT. In the case of *De Jager v Netcare*, before the Pretoria High Court, a constitutional issue has been raised by the plaintiff concerning the surveillance of the plaintiff and some of his family members by private investigators appointed by the defendant. I would like you to draft a judgment on this constitutional

²⁴ *De Jager v Netcare* (2024) Case No: 42041/16 ZAGPPHC 503.

²⁵ *De Jager v Netcare Limited and Others* 2025 JDR 0793 (GP).

²⁶ See Kyle Wiggers ‘OpenAI debuts GPT-4o ‘omni’ model now powering ChatGPT’ *TechCrunch* 13 May 2024 <https://techcrunch.com/2024/05/13/openais-newest-model-is-gpt-4o/> (last accessed 12 June 2025/06/12).

²⁷ See OpenAI ‘What is the ChatGPT model selector?’ https://help.openai.com/en/articles/7864572-what-is-the-chatgpt-model-selector#h_80fc86948a (last accessed 2025/06/12).

²⁸ OpenAI ‘ChatGPT-4o’ <https://platform.openai.com/docs/models/gpt-4o> (last accessed 2025/03/13).

1. ²⁹ AI Judgment <https://chatgpt.com/share/68e29955-19ac-800d-9c63-cb0fd9f636cc> (last accessed 2025/10/5).

³⁰ See Z Lin ‘How to write effective prompts for large language models’ (2024) 8 *Nature Human Behaviour* 614.

An initial foray into judicial drafting: a comparative experiment in a constitutional privacy case

issue. I will provide you with all the relevant court papers which contain written arguments and the case law cited in those arguments. Your role will be to act as the presiding judge, Judge Motha of the Pretoria High Court. You are to draft a judgment in the typical structure (provide heading & numbered paragraphs) and style of South African jurisprudence as reflected in the Jura Judgment Style Guide which I have provided you in the attachment. You must provide full citations of law & cases you utilise using footnotes; if you cannot perform this function provide references in brackets within the text. This task will entail, most prominently, that you provide an overview of the facts, analyse the issue for determination—i.e., the constitutional issue—critically evaluate the arguments presented by the parties and the amici curiae, and decide the issue, with an appropriate cost order. You are required to make decisions based on papers provided and not your inherent knowledge. In writing your judgement you must consider judgment writing guidelines under South African law that are listed below: Write using a logical structure Provide reasons to justify your conclusions Be impartial and bias free Avoid undue technicality but maintain professionalism Utilise prescribed formats Succinctly deliberate and pronouns on all arguments, evidence, applicable laws Deliver judgement promptly Promote human rights subject to appropriate limitations We will now proceed as follows: I will provide you with ten of the relevant documents, one at a time. For each document, you will first provide me with a summary. Are you ready to begin?³¹

3.4 Inputs and exclusions

The AI was provided with the following documents: plaintiff's heads of argument, defendant's heads of argument, both parties' supplementary heads of argument, the amicus curiae submissions from Professors Thaldar and Snail ka Mtuze, the Rule 16A notice,³² relevant extracts from the Uniform Rules of Court (Rules 16A, 36, and provisions on admissibility),³³ and the *Jura Judgment Style Guide*.³⁴ These materials encompassed the legal and factual framework of the case. Excluded materials included oral submissions and transcripts, which were unavailable and less central to the study's aim of evaluating formal drafting. To prevent bias and regurgitation, the human-authored judgment was withheld during drafting.

³¹ AI Judgment <https://chatgpt.com/share/68e29955-19ac-800d-9c63-cb0fd9f636cc> (last accessed 2025/10/5).

³² Plaintiff's Notice in Terms of Rule 16A (24 May 2024).

³³ Supreme Court Act 59 of 1959 as amended. GNR 5124, G. 51056 of 16 August 2024 (with effect from 20 September 2024).

³⁴ See Jura 'Judgment Style Guide Jura' https://jura.co.za/media/filestore/2017/05/Judgment_Style_Guide_Jura.doc (last accessed 2025/01/29).

3.5 Ethical safeguards and constitutional alignment

To ensure constitutional alignment, the prompting protocol included explicit ethical directives. ChatGPT-4o was instructed to apply legal principles impartially,³⁵ justify conclusions logically, promoting constitutional values and rights (i.e. human dignity, equality, children's rights),³⁶ and acknowledge its analytical limitations. These directives were reiterated throughout the prompting sequence to align the AI's output with the expectations of S165(2) of the Constitution of the Republic of South Africa, 1996.³⁷ In embedding these safeguards, the study ensured that the AI remained an impartial support tool—capable of assisting and modelling after judicial reasoning norms such as effectiveness and accessibility.³⁸

4 Formulating Evaluation framework

To meaningfully assess ChatGPT-4o's performance in drafting a South African High Court judgment, this study developed a structured evaluation framework aligned with both doctrinal expectations and constitutional values. The framework draws on effective prompting mechanisms,³⁹ judgment-writing manuals,⁴⁰ ethical standards for legal professionals, and academic literature on AI-assisted legal reasoning. It comprises five interrelated criteria.

4.1 Structure

Judgments must be intelligible and adhere to recognised conventions. South African courts typically adopt a fluid format based on: introduction, factual background, issues for determination, legal framework, application of the law, and conclusion.⁴¹ Formal features such as paragraph numbering, proper citation, and meaningful headings contribute to navigability and professional integrity.⁴² The *Juta Judgment Style Guide*⁴³ and

³⁵ S165(2) of the Constitution of the Republic of South Africa, 1996: 'The courts are independent and subject only to the Constitution and the law, which they must apply impartially and without fear, favour or prejudice'.

³⁶ S28(2) of the Constitution of the Republic of South Africa, 1996 'A child's best interests are of paramount importance in every matter concerning the child.'

³⁷ S28(2) of the Constitution of the Republic of South Africa, 1996 'A child's best interests are of paramount importance in every matter concerning the child.'

³⁸ Section 165(4) of the Constitution of the Republic of South Africa, 1996.

³⁹ Z Lin 'How to write effective prompts for large language models' (2024) 8 *Nature Human Behaviour* 611–615.

⁴⁰ Gauteng Local Division: Pretoria 'Practice Manual of the Gauteng Division Pretoria' 25 July 2011 <https://johannesburgbar.co.za/practice-manual> (last accessed 2025/03/16).

⁴¹ L Theron 'Refreshing And Enhancing Judgment Writing Skills' *The South African Judicial Education Journal* Volume 5 Issue 1 2022 Pages 4–5, Section IV.

⁴² B Tomlinson, AW Torrance & RW Black 'ChatGPT and Works Scholarly: Best Practices and Legal Pitfalls in Writing with AI' (2023) 76 *SMU L Rev Forum* 117

⁴³ Juta 'Judgment Style Guide Juta' https://juta.co.za/media/filestore/2017/05/Judgment_Style_Guide_Juta.doc (last accessed 2025/01/29).

An initial foray into judicial drafting: a comparative experiment in a constitutional privacy case

the training materials of the South African Judicial Education Institute (SAJEI) were used as benchmarks.⁴⁴

This criterion evaluates whether ChatGPT-4o:

Organises content according to conventional judicial structure;

Applies paragraph numbering correctly;

Uses appropriate citation format;

Produces a document consistent with professional publication standards.

4.2 Factual accuracy

Judicial decisions must be grounded in a faithful representation of the factual matrix. Misstatements of fact, incorrect attribution of arguments, or distortion of the procedural history can compromise the integrity of a judgment.⁴⁵ In the context of generative AI, the risk of hallucinated or fabricated content must be carefully assessed and verified with independent reading of original sources.

This criterion evaluates whether ChatGPT-4o:

Accurately identifies and summarises material facts from the record;

Correctly attributes arguments and evidence to the appropriate party;

Avoids factual invention, distortion, or omission.

4.3 Use of legal authority

Value 6 of the Bangalore Principles places emphasis on making conscientious conclusions based on understanding of the law.⁴⁶ Thus citing relevant legislation, precedents, and scholarly sources is central to sound judgment writing.⁴⁷ Proper citation affirms adherence to the principle of legality, ensures transparency, reinforces the persuasive weight of judicial decisions and affirms faith in the integrity of the judicial decisions.⁴⁸ In this context, authority is assessed not only for correctness but also for contextual appropriateness.

This criterion evaluates whether ChatGPT-4o:

Correctly identifies and applies relevant statutory provisions (e.g. POPIA, Constitution);

Cites applicable judicial precedents from relevant South African courts;

Appropriately draws on legal texts and academic commentary;

Demonstrates an understanding of how and why legal authorities are invoked.

⁴⁴ Established in terms of South African Judicial Education Institute Act 14 of 2008.

⁴⁵ B Tomlinson, AW Torrance & RW Black 'ChatGPT and Works Scholarly: Best Practices and Legal Pitfalls in Writing with AI' (2023) 76 *SMU L Rev Forum* 108.

⁴⁶ United Nations Bangalore Principles and Commonwealth Latimer House Principles ECOSOC 2006/23.

⁴⁷ B Tomlinson, AW Torrance & RW Black 'ChatGPT and Works Scholarly: Best Practices and Legal Pitfalls in Writing with AI' (2023) 76 *SMU L Rev Forum* 117.

⁴⁸ United Nations Bangalore Principles and Commonwealth Latimer House Principles ECOSOC 2006/23 Value 6.

4.4 Legal reasoning

Legal reasoning constitutes the core of any judicial opinion and ‘...the primary purpose of the judgment is to communicate the decision as well as the reasoning of the court.’⁴⁹ Legal reasoning involves the interpretation and application of legal rules to the facts, the assessment of competing arguments, and the articulation of a logically coherent path to judgment.⁵⁰ The purpose here is not to determine whether the AI or the court reached the “correct” outcome, but to assess, in a formal sense, the internal consistency and logical progression of the reasoning presented.

This criterion evaluates whether ChatGPT-4o:

Clearly identifies and frames the legal questions in dispute;

Applies relevant legal principles and interpretive doctrines (e.g. subsidiarity, proportionality, evidentiary thresholds);

Engages with the core submissions of both parties and any amici curiae;

Demonstrates logical coherence, doctrinal fidelity, and analytical adequacy.

By isolating the structural quality of legal reasoning from its normative content, this criterion allows for an assessment of the model’s capacity to simulate formal judicial logic—not its capacity for judicial discretion or moral judgment.

4.5 Clarity

‘Brevity, simplicity, and clarity are the watchdogs for effective judicial writing.’⁵¹ Essentially judicial language must be clear, internally consistent, and appropriate to its legal and institutional function, courteous and averse to imploring inflammatory or disparaging language.⁵² Judgments are written primarily for legal professionals and future courts, and may appropriately employ legal terminology, including Latin maxims and technical phrases. The expectation is not simplicity but clarity—the avoidance of unnecessary convolution, ambiguity, or verbosity. This approach reflects the realities of South African legal writing, where technical language is expected and where clarity is judged not by popular accessibility, but by the internal communicative standards of the legal profession.

This criterion evaluates whether ChatGPT-4o:

⁴⁹ L Theron ‘Refreshing And Enhancing Judgment Writing Skills’ *The South African Judicial Education Journal* Volume 5 Issue 1 2022 Page 3, Section IV.

⁵⁰ L Theron ‘Refreshing And Enhancing Judgment Writing Skills’ *The South African Judicial Education Journal* Volume 5 Issue 1 2022 Pages 4–5, Section IV.

⁵¹ Theron citing address delivered by the former Chief Justice Ngcobo of the Constitutional Court of South Africa entitled ‘Judgment Writing’ at L Theron ‘Refreshing And Enhancing Judgment Writing Skills’ *The South African Judicial Education Journal* Volume 5 Issue 1 2022 9.

⁵² ‘Do be respectful and courteous to all parties, the legal representatives as well as colleagues’ at L Theron ‘Refreshing And Enhancing Judgment Writing Skills’ *The South African Judicial Education Journal* Volume 5 Issue 1 2022, 8.

An initial foray into judicial drafting: a comparative experiment in a constitutional privacy case

- Uses language that is appropriate to the judicial context;
- Presents its reasoning in a coherent and comprehensible manner;
- Maintains a tone of judicial professionalism and impartiality;
- Avoids stylistic inconsistency or rhetorical excess.

5 Results and comparative analysis

The 20-prompt experiment in ChatGPT-4o generated a full-length judgment (AI Judgment) in the Appendix—which was assessed against the actual judgment in *De Jager Case* by Motha J. Applying the evaluation criteria established in Section 4, the comparison confirmed the AI-generated judgment's superior technical performance across four core areas—structure, factual accuracy, legal reasoning, and clarity—while recognising the human-authored judgment's effective use of expert input and its engagement with broader contextual considerations.

5.1 Structure

The AI-generated judgment broadly followed the conventional structure of a South African High Court judgment, presenting distinct sections titled Introduction, Factual Background, Legal Framework, Arguments Presented, Evaluation of Arguments, and Conclusion. It employed consistent paragraph numbering, clearly defined headings, and an overall tone aligned with the Jura Judgment Style Guide. Case and statutory citations—such as *Bernstein v Bester*⁵³ and relevant POPIA provisions—were accurately cited. The inclusion of a formal order and a judge's signature block further contributed to its professional appearance. However, certain stylistic choices deviated from South African judicial norms. The use of bullet points (e.g., to enumerate legal questions and to structure the order), bold text for headings (e.g., "Introduction,"), and emphasis (e.g., "reasonable expectation of privacy,") are inconsistent with established High Court conventions, which favour numbered paragraphs and plain or underlined headings.⁵⁴ While these elements may enhance readability in general prose, they are inappropriate for formal judgments and detract from stylistic fidelity. Additionally, minor repetition of central issues—such as privacy and admissibility—slightly affected navigability.

The human-authored judgment presented a broadly recognisable structure, with sections titled Introduction, The Issue in Contention, Submissions by Amici Curiae and Counsel, Under POPIA, Conclusion, and Order.⁵⁵ It did not include a dedicated factual background section, and legal analysis was often blended with the legal framework, which reduced clarity. Transitions between key issues were occasionally abrupt or implicit. It correctly employed paragraph numbering, footnoting and conformed to basic

⁵³ 1996 (2) SA 751 (CC).

⁵⁴ *AI Judgment*, para 24; para 76; para 1; para 50.

⁵⁵ *De Jager v Netcare Ltd* para 1; para 5; para 11; para 20; para 31; para 32.

formatting expectations in line with SAJEI standards. Importantly, it avoided the use of bullets or bold text, adhering more closely to conventional judicial presentation.

Comparative Assessment: The AI judgment exhibited strong internal structure and logical sequencing, but its use of non-standard formatting elements—such as bullets and boldface—undermined its conformity to High Court conventions. The Human Judgment, though less well-structured and occasionally lacking in clarity, adhered more faithfully to accepted formatting norms. As a result, it holds a slight advantage in stylistic compliance, despite its structural limitations.

5.2 Factual accuracy

The AI-generated judgment presented a detailed and factually accurate account of the background relevant to the constitutional challenge. It correctly identified the surveillance conducted by Dion Pienaar, the plaintiff's R25.7 million damages claim, and the incidental inclusion of minor grandchildren in the surveillance footage.⁵⁶ The judgment appropriately attributed the plaintiff's reliance on Section 14 of the Constitution to support a claim of unlawful surveillance, and the defendant's counter-argument that the surveillance was justified for litigation purposes.⁵⁷ The procedural history was clearly and coherently summarised.⁵⁸ Although the judgment included extraneous factual details—such as the plaintiff's cataract surgery and the specific timeline of the surveillance (August–September 2020)—these were accurate, if not strictly necessary for resolving the constitutional and statutory questions.⁵⁹

The human-authored judgment also accurately referenced the central facts necessary for adjudicating the constitutional and statutory issues.⁶⁰ It noted the surveillance conducted by Dion Pienaar and the plaintiff's damages claim, expressed within a range of R24.8–25.7 million.⁶¹ While it did not contain a dedicated factual section or elaborate on the broader factual context (e.g., the surgery or precise surveillance timeline), these omissions were not material to the legal reasoning, which focused principally on the application of subsidiarity and the interpretation of POPIA. The limited engagement with party submissions reflects the procedural posture of the case: the defendant adopted the submissions of Professor Thaldar without elaboration, and the plaintiff only addressed POPIA in supplementary heads of argument, prompted by Thaldar's intervention. The judgment's reliance on *amici curiae*—Professors Thaldar and Snail ka Mtuze—provided the necessary interpretive structure, particularly in light of the parties' constrained submissions. No factual errors or misrepresentations were identified.

Comparative Assessment: Both judgments satisfied the accuracy criterion, presenting reliable and contextually appropriate factual accounts. The AI Judgment held a slight advantage due to its clearer attribution of arguments and fuller factual detail,

⁵⁶ *AI Judgment* para 7.

⁵⁷ *AI Judgment* para 14–16; para 53; para 75.

⁵⁸ *AI Judgment* para 1–12.

⁵⁹ *AI Judgment* para 6–8.

⁶⁰ *De Jager v Netcare Ltd* para 5.

⁶¹ *De Jager v Netcare Ltd* para 36.

An initial foray into judicial drafting: a comparative experiment in a constitutional privacy case

while the Human Judgment offered sufficient but more concise references, appropriately reflecting its doctrinal focus and reliance on amici for interpretive clarity.

5.3 Use of legal authority

The AI judgment cited an array of legal authorities, including constitutional provisions (sections 14, 28, and 36), various sections of POPIA (e.g., 11, 12, 18, 27, 34), and leading case law such as *Bernstein v Bester*,⁶² *Gaertner v Minister of Finance*,⁶³ *Key v Attorney-General*,⁶⁴ and *City of Cape Town v SANRAL*.⁶⁵ These citations were generally contextually appropriate and well integrated into the judgment's logical flow. For example, *Bernstein* was used to demarcate the scope of privacy rights, particularly for surveillance in private settings, while *Gaertner* supported the assessment of reasonable expectations of privacy in the home.⁶⁶ *Key* framed the admissibility of potentially unlawfully obtained evidence, and *City of Cape Town* supported the balance between privacy and open justice.⁶⁷ The AI judgment also referenced *Eskom Holdings SOC Ltd v Vaal River Development Association (Pty) Ltd and Others* to support the application of the subsidiarity doctrine.⁶⁸ The authorities were formatted correctly and used with superficial coherence. However, the judgment did not cite any academic literature, limiting its depth of engagement with broader theoretical discourse.

The human-authored judgment cited a wider selection of legal authorities. It relied on constitutional provisions (sections 14 and 36), POPIA sections (including 2, 6, 9, 11, 18, 26, 27), and precedent such as *Bernstein v Bester*, *Mazibuko v City of Johannesburg*,⁶⁹ *Eskom Holdings*, *Gaertner*, and *Smuts⁷⁰ v Botha*.⁷¹ The use of *Mazibuko* and *Eskom Holdings* appropriately grounded the subsidiarity argument, emphasising that POPIA, rather than a direct constitutional claim, was the proper basis for the plaintiff's challenge.⁷² *Smuts v Botha* was cited to assess the reasonableness of public surveillance, aligning with the case's focus.⁷³ Academic works by Cohen and Neethling, enhanced the judgment's doctrinal context, particularly in the analysis of personality rights.⁷⁴ Judge Motha referenced the European General Data Protection Regulation 2018 (GDPR) and traces its lineage from German models.⁷⁵ The amici

⁶² (1996 (2) SA 751 (CC);
AI Judgment para 26, 39, 50.

⁶³ 2014 (1) SA 442 (CC).

⁶⁴ 1996 (4) SA 187 (CC).

⁶⁵ 2015 (3) SA 386 (SCA).

⁶⁶ *AI Judgment* 15–16.

⁶⁷ *AI Judgment* para 47; 65.

⁶⁸ 2023 (4) SA 325 (CC).

⁶⁹ 2010 (4) SA 1 (CC).

⁷⁰ 2022 (2) SA 425 (SCA).

⁷¹ *De Jager v Netcare Ltd* para 2-3; para 6-8; para 11, para 14-15, 17-19; para 26-27; para 34.

⁷² *De Jager v Netcare Ltd* para 5-8.

⁷³ *De Jager v Netcare Ltd* para 34.

⁷⁴ *De Jager v Netcare Ltd* para 2; para 31.

⁷⁵ *De Jager v Netcare Ltd* para 38.

curiae—Professors Thaldar and Snail ka Mtuze—provided significant interpretive contributions, particularly in clarifying the applicability of POPIA exemptions and reinforcing the importance of subsidiarity. Their input was procedurally significant, as the defendant adopted Thaldar’s submissions and the plaintiff addressed POPIA only belatedly. However, the judgment did not reference certain evidentiary precedents such as *S v Naidoo*, which might have strengthened the admissibility analysis. The judgment also included a brief discussion on decolonisation and Afro-centric law that, while socially relevant, lacked legal specificity and did not contribute to doctrinal analysis.⁷⁶

Comparative Assessment: The AI judgment demonstrated imprecise and informal citation of relevant legal authorities, offering structurally sound integration of constitutional principles and precedent. Its omission of academic commentary limited its interpretive richness. The Human Judgment, while citing fewer cases overall, enriched its analysis through scholarly references and the strategic and procedurally justified use of amici submissions. Its citation of *Smuts v Botha* was appropriate and well-contextualised. However, the absence of additional evidentiary case law and the inclusion of a rhetorically framed decolonisation discussion somewhat diluted the legal focus. On balance, the AI judgment showed greater formal breadth in its citation practices, while the Human Judgment demonstrated complementary intellectual and procedural value. This evaluation is limited to the formal application of legal authorities and does not assess their substantive correctness.

5.4 Legal reasoning

The AI judgment clearly articulated the legal questions: whether the surveillance infringed the plaintiff’s constitutional right to privacy under section 14, whether it complied with POPIA, and whether the evidence was admissible.⁷⁷ It demonstrated a formal understanding of subsidiarity, citing *Eskom Holdings SOC Ltd* to support the conclusion that POPIA, not the Constitution, should have been the primary basis for the plaintiff’s claim.⁷⁸ Its proportionality analysis under section 36 was logically structured, evaluating the necessity and reasonableness of the surveillance.⁷⁹ The AI judgment concluded that surveillance in public spaces was permissible due to a diminished expectation of privacy, whereas surveillance at the plaintiff’s home and involving minors was not, noting less intrusive alternatives like Rule 36 medical examinations.⁸⁰ It engaged with the parties’ submissions and incorporated amici perspectives, citing relevant precedent including *Bernstein v Bester*, *Gaertner*, and *Key v Attorney-General*, showing a superficially coherent application of legal principles. Minor repetition did not detract significantly from its analytical structure.

⁷⁶ *De Jager v Netcare Ltd* para 29; para 30.

⁷⁷ *AI Judgment* para 74.

⁷⁸ *AI Judgment* para 75.

⁷⁹ *AI Judgment* para 75.

⁸⁰ *AI Judgment* para 10; para 33; para 69; para 75.

An initial foray into judicial drafting: a comparative experiment in a constitutional privacy case

The human-authored judgment focused on whether the plaintiff's section 14 claim was sustainable in light of the subsidiarity principle.⁸¹ Citing *Mazibuko*⁸² v *City of Johannesburg* and *Eskom Holdings*, it concluded that POPIA should have been relied on instead of the Constitution—a procedurally sound finding given that POPIA was only introduced via Professor Thaldar's amicus submission.⁸³ The defendant adopted these submissions without elaboration, and the plaintiff responded only in supplementary heads. The judgment's proportionality analysis under section 36 was brief but formally correct, referencing *Gaertner* and *Smuts v Botha*⁸⁴ to assess the legitimacy of the surveillance.⁸⁵ The reliance on amici curiae (Professors Thaldar and Snail ka Mtuze) was contextually appropriate and provided necessary doctrinal clarity. The conclusion to admit the surveillance under POPIA's litigation exemption was supported by a logical rationale. A discussion of decolonisation and Afro-centric legal theory, though present, was tangential and did not substantively contribute to the legal reasoning. Transitions between issues were less clearly structured, and the engagement with party submissions was minimal, but consistent with the case's procedural dynamics.

Comparative Assessment: The AI judgment exhibited greater structural coherence, clearer issue framing, and broader engagement with submissions and precedent, providing an advantage in formal reasoning. It demonstrated a superficially coherent application of key doctrines and maintained analytical adequacy throughout. The Human Judgment, while less systematically developed, correctly applied subsidiarity and POPIA principles and effectively relied on amici to address procedural deficiencies. However, its transitions were less clear, and the inclusion of unrelated socio-legal commentary somewhat diluted its analytical focus. On balance, the AI judgment reflected greater technical rigour, while the Human Judgment demonstrated sufficient doctrinal understanding within a narrower scope. This assessment is limited to the formal application of legal principles, without commenting on the substantive merits of either judgment's outcome.

5.5 Clarity

The AI judgment maintained a formal, professional, and impartial tone throughout. It used legal terminology with precision—for example, 'reasonable expectation of privacy', 'fruit of the poisonous tree', and 'proportionality under section 36'—in correct doctrinal contexts.⁸⁶ Complex constitutional and statutory issues were clearly articulated, with logical transitions between sections. Sentences were concise and unambiguous, and the final order was phrased with clarity and formality. The AI judgment's language use aligned well with formal judicial communication and would be readily intelligible to legal professionals.

⁸¹ *De Jager v Netcare Ltd* para 6–7.

⁸² 2010 (4) SA 1 (CC).

⁸³ *De Jager v Netcare Ltd* para 5–9.

⁸⁴ 2022 (2) SA 425 (SCA).

⁸⁵ *De Jager v Netcare Ltd* para 19–24.

⁸⁶ *AI Judgment* para 10–11; para 19; para 11.

The human-authored judgment communicated key legal issues effectively and used doctrinal language accurately. It correctly employed terms like ‘subsidiarity’, ‘*boni mores*’, ‘data subject’, and ‘minimality principle’ within their relevant contexts.⁸⁷ Key findings, particularly those addressing amici contributions and the final order, were expressed with clarity. However, the judgment’s writing style was uneven at times. Informal expressions such as “just what the doctor ordered,” “packs a devastating punch,” and “Kasi-flavour” disrupted the otherwise judicial tone.⁸⁸ The discussion of decolonisation and Afro-centric law had potential relevance to South Africa’s socio-legal context, but it was not connected to the legal analysis of privacy, POPIA, or admissibility.⁸⁹ As a result, it came across as rhetorical rather than substantive, diminishing cohesion. Some transitions between sections were less fluid, occasionally requiring the reader to infer logical connections. Despite these stylistic issues, the judgment successfully conveyed its core reasoning and conclusions.

Comparative Assessment: The AI judgment exhibited consistent formality, simpler legal terminology, and a professional tone suited to High Court writing. The human judgment, while generally clear in its analysis, suffered from uneven stylistic execution and rhetorical asides that detracted from coherence. Although the decolonisation commentary could have added socio-legal insight, its lack of actionable development rendered it a distraction. On balance, the AI judgment better conformed to the communicative expectations of judicial writing, while the human-authored judgment—though effective in parts—was less consistent in tone and structure. This assessment concerns the clarity and suitability of language, without evaluating the substantive merits of either judgment.

5.6 Overall assessment

This evaluation, focuses on formal logic and conformity to judicial standards, does not assess the substantive correctness of either outcome. Across the five evaluative criteria, the AI-generated judgment outperformed the human-authored judgment in structural coherence, factual accuracy, formal legal reasoning, and linguistic clarity, demonstrating strong technical alignment with the standards expected of a South African High Court judgment. Its structured layout, accurate use of legal authorities—such as *Bernstein v Bester* and *Key v Attorney-General*—and clear articulation of doctrines like subsidiarity and proportionality underscored its potential as an effective drafting tool. However, its use of non-standard formatting elements—namely bullets and bold text (e.g., headings)—departed from accepted judicial publishing norms, slightly diminishing stylistic fidelity. Occasional repetition also detracted from overall concision.⁹⁰

The human-authored judgment, while less technically refined, effectively leveraged the input of amici curiae—Professors Thaldar and Snail ka Mtuze—to address procedural shortcomings, particularly in light of the plaintiff’s delayed engagement with

⁸⁷ *De Jager v Netcare Ltd* para 7; para 33; para 35; para 34.

⁸⁸ *De Jager* para 1; para 5; para 38.

⁸⁹ *De Jager v Netcare Ltd* para 37.

⁹⁰ *AI Judgment* para 24; para 76; para 50.

An initial foray into judicial drafting: a comparative experiment in a constitutional privacy case

POPIA and the defendant's adoption of Thaldar's submissions. Its citation of academic sources (e.g., Cohen, Neethling) and contextually appropriate case law (e.g. *Smuts v Botha*) added doctrinal depth. Nonetheless, structural inconsistencies, stylistic variability, and an underdeveloped decolonisation discussion reduced its overall formal coherence. While the commentary on decolonisation was potentially relevant to broader socio-legal concerns—such as colonial influences on privacy law—it was not substantively integrated into the legal analysis and appeared more rhetorical than doctrinal.

These findings highlight the significant promise of generative AI in producing structured, legally plausible draft judgments, particularly in technical execution. However, formatting irregularities illustrate the need for closer adherence to judicial conventions. The human judgment's strategic use of expert contributions and engagement with South Africa's procedural realities underscores the irreplaceable interpretive role of human judges in constitutional adjudication. AI-generated drafts may serve as valuable foundations for effective judicial writing, but rigorous human oversight remains essential to ensure stylistic compliance, doctrinal soundness and that judicial authority and independence remains with the courts in alignment with constitutional values, as required by section 165 of the Constitution.

6 Discussion

This study represents an initial exploration of generative AI, specifically the model ChatGPT-4o, in the context of judicial drafting, comparing its output to a human-authored High Court judgment in *De Jager v Netcare Ltd*. The evaluation focused on formal criteria—structure, factual accuracy, use of legal authority, legal reasoning, and clarity of language—and found that the AI-generated judgment outperformed the human-authored judgment in technical execution. Notably, however, both judgments exhibited striking analytical and conclusory similarities. This shared reasoning trajectory—driven by a common reliance on party and amici submissions—raises important questions about the role of legal argumentation in shaping both AI and human decision-making. This section explores the significance of these findings, their implications for legal technology, and directions for future research, while emphasising the exploratory nature of this inquiry.

6.1 Surprising convergence: Analytical and conclusory similarity

The most notable finding was the similarity between the AI and Human Judgments in both analytical structure and legal conclusions. Both judgments framed the key legal questions—constitutional privacy (Section 14), POPIA compliance, and evidence admissibility—through the doctrines of subsidiarity (*Eskom Holdings SOC Ltd; Mazibuko v City of Johannesburg* and proportionality (Section 36), ultimately finding that surveillance in public spaces was admissible while surveillance of the home or minors was more problematic. Although their orders diverged slightly—the AI judgment explicitly excluded home/minor footage, while the Human Judgment admitted surveillance more

broadly under POPIA exemptions—both prioritised litigation imperatives over privacy claims.

This convergence is especially significant given the evaluation’s methodological focus on formal qualities rather than substantive outcomes. As detailed in the *Results*, both judgments were shaped by the same argumentative inputs: the plaintiff’s delayed invocation of POPIA, the defendant’s adoption of Professor Thaldar’s arguments, and the amici submissions of Professors Thaldar and Snail ka Mtuze. The *Legal reasoning* analysis further showed that both judgments drew from a shared corpus of precedents, such as *Bernstein v Bester* and *Gaertner v Minister of Finance*, reinforcing the analytical symmetry.

6.2 Why the Similarity? Submissions as a framing force

The similarity in reasoning likely reflects the central role of submissions in framing the analysis. As noted in *Factual accuracy*, the plaintiff focused initially on constitutional rights, while amici supplied the doctrinal framework of POPIA exemptions and subsidiarity. The defendant adopted Thaldar’s submissions in full. The AI judgment systematically presented these arguments, while the Human Judgment relied on amici to clarify the legal framework.⁹¹

This convergence invites further reflection: do AI and human judges process persuasive legal arguments in comparable ways? In South African legal practice, well-developed submissions often shape outcomes. The AI’s capacity to mirror this process—producing a formally coherent judgment grounded in the same inputs—suggests that AI may emulate core patterns of human legal reasoning when prompted with structured material. Alternatively, it may reflect a deeper alignment between the way large language models and judges weigh legal inputs. While these findings are preliminary, they raise meaningful questions about the nature of legal reasoning across human and artificial decision-makers.

6.3 Technical strengths and stylistic flaws of AI

The *Results* confirmed the AI judgment’s technical strengths, including superior structural coherence (Section 5.1), detailed factual precision (Section 5.2), broad use of authority (Section 5.3), clear legal reasoning (Section 5.4), and consistent tone and clarity (Section 5.5). These traits mark generative AI as a promising tool for producing polished, legally plausible draft judgments. However, formatting deviations—particularly the use of bullets and bold text—departed from the *Juta Judgment Style Guide* and detracted from stylistic fidelity.⁹²

Additionally, the AI’s lack of scholarly engagement (Section 5.3) highlights a limitation in interpretive richness. While its use of terminology such as "reasonable expectation of privacy" was contextually accurate, it offered limited insight into normative

⁹¹ *De Jager v Netcare Ltd* para 11–13.

⁹² *AI Judgment* para 3; para 8; para 10–13; para 17; para 21; para 31; para 33; para 42; para 75.

An initial foray into judicial drafting: a comparative experiment in a constitutional privacy case

or academic debates.⁹³ This suggests that AI excels in formal replication of legal language but still requires human input to achieve interpretive depth.

6.4 Human Judgment’s procedural and contextual nuances

Despite weaker technical execution, the Human Judgment demonstrated contextually informed decision-making. Its reliance on amici was appropriate, given the plaintiff’s delayed POPIA submissions and the defendant’s reliance on Thaldar. It effectively integrated academic commentary (e.g., Neethling, Cohen) and cited *Smuts v Botha* to assess surveillance reasonableness (Section 5.3).⁹⁴

However, structural inconsistency (Section 5.1), variable tone (Section 5), and the inclusion of rhetorical devices (e.g., “Kasi-flavour”) reduced formal coherence.⁹⁵ The decolonisation discussion, while potentially relevant to critiques of colonial legacies in privacy law, was not tied to the legal analysis and functioned more as a rhetorical aside than a substantive argument (Section 5). This illustrates the human capacity to introduce socio-legal commentary, even when imperfectly executed—something AI cannot replicate without explicit prompting.

6.5 Implications and future directions

The analytical alignment between AI and human judgments in this case suggests that AI can draft judgments that are, at least formally, comparable to human output—particularly when prompted with high-quality submissions and quality prompts. While AI’s formatting errors and shallow normative engagement remain significant limitations, the findings support the notion of AI as a valuable drafting tool.

Future studies should explore how generative AI performs across different legal contexts, including other jurisdictions, levels of court, and types of legal disputes. Comparative research on large language models and prompt engineering may further clarify whether AI genuinely mirrors human reasoning or merely assembles plausible text. The Human Judgment’s use of amici also suggests a future role for AI in organising and structuring expert input for judicial review.

7 Conclusion

This study demonstrates that ChatGPT-4o can generate technically proficient, legally plausible judicial drafts. The AI judgment displayed strong formal structure, doctrinal accuracy, and clarity, though it was weakened by formatting inconsistencies and limited contextual depth. Strikingly, its analysis and conclusions closely mirrored those of the Human Judgment—likely a result of shared reliance on party and amici submissions. The Human Judgment, for its part, illustrated the enduring value of human discretion,

⁹³ *AI Judgment* para 10–11.

⁹⁴ *De Jager v Netcare Ltd* para 2; para 34.

⁹⁵ *De Jager* para 1; para 5; para 38.

expert engagement, and contextual awareness, despite structural and stylistic limitations. These findings support the potential of AI as a judicial aid but reinforce the need for robust human oversight to ensure doctrinal integrity and constitutional fidelity, as required by section 165(4) of the Constitution. The evaluation focuses on form, not substantive correctness; whether both judgments' conclusions were legally sound remains a question for appellate review.

8 References

1. Africa AI Policy Lab Africa AI Privacy Report 2025 (2025) available at <https://www.aipolicy.africa/africa-intelligence-report> (last accessed 2025/04/15).
2. Bernstein v Bester 1996 (2) SA 751 (CC).
3. Broyde, M & Michael, J 'AI and Jewish Law: Seeing How ChatGPT 4.0 Looks at a Novel Issue' (2023) Buchman Law School, Tel Aviv University, and Cardozo Law School, Yeshiva University, CSLR Research Paper No. 12.2023-AFF.
4. Hicks, M.T., Humphries, J., & Slater, J. (2024). *ChatGPT is bullshit. Ethics and Information Technology*, 26, 38. <https://doi.org/10.1007/s10676-024-09775-5>.
5. City of Cape Town v SANRAL 2015 (3) SA 386 (SCA).
6. De Jager v Netcare (2024) Case No: 42041/16 ZAGPPHC 503.
7. De Jager v Netcare Limited and Others 2025 JDR 0793 (GP)
8. Eskom Holdings SOC Ltd v Vaal River Development Association (Pty) Ltd and Others 2023 (4) SA 325 (CC).
9. European General Data Protection Regulation 2018 (GDPR).
10. First Circuit Court of Cartagena, Case No. 11001-22-2023-00016.
11. AI Judgment <https://chatgpt.com/share/68e29955-19ac-800d-9c63-cb0fd9f636cc> (last accessed 2025/10/5).
12. Gaertner v Minister of Finance 2014 (1) SA 442 (CC).
13. Galarreta, Francisco Javier Fernández, "The Impact of Artificial Intelligence Technologies on the Justice Administration and on the Judicial Office Personnel" *Zeitschrift für Rechtssoziologie*, vol. 45, no. 1, 2025.
14. Gauteng Local Division: Pretoria 'Practice Manual of the Gauteng Division Pretoria' 25 July 2011 <https://johannesburgbar.co.za/practice-manual> (last accessed 2025/03/16).
15. Hicks, M.T., Humphries, J. & Slater, J. ChatGPT is bullshit. *Ethics Inf Technol* 26, 38 (2024). <https://doi.org/10.1007/s10676-024-09775-5>
16. Juta 'Judgment Style Guide Juta' https://juta.co.za/media/filestore/2017/05/Judgment_Style_Guide_Juta.doc (last accessed 2025/01/29).
17. Key v Attorney-General 1996 (4) SA 187 (CC)).
18. Kyle Wiggers 'OpenAI debuts GPT-4o 'omni' model now powering ChatGPT' TechCrunch 13 May 2024 <https://techcrunch.com/2024/05/13/openais-newest-model-is-gpt-4o/> (last accessed 2025/06/12)
19. KwaZulu-Natal and Others [2025] 2 ZAKZPHC.
20. Lin Z 'How to write effective prompts for large language models'(2024) 8 *Nature Human Behaviour* 614.
21. Mata v Avianca Inc 22-cv-1461 (PKC) United States District Court, S.D. New York (2023).
22. Mavundla v MEC: Department of Co-Operative Government and Traditional Affairs [2025] 2 ZAKZPHC.
23. Mazibuko v City of Johannesburg 2010 (4) SA 1 (CC)

An initial foray into judicial drafting: a comparative experiment in a constitutional privacy case

24. N Lefakane 'The Use of Artificial Intelligence for Adjudication in South Africa' LLM dissertation, North-West University, 2022.
25. Open AI 'ChatGPT Traceability Document - AI Experiment' <https://chatgpt.com/canvas/shared/67cd52d06db08191a3507752cfc022e1> (last accessed 2025/04/19).
26. OpenAI "ChatGPT Overview" <https://openai.com/chatgpt/overview/> 2025 (last accessed 2025-03-24).
27. OpenAI "ChatGPT-4o" <https://platform.openai.com/docs/models/gpt-4o> 2025 (last accessed 2025-03-13).
28. OpenAI 'ChatGPT-4o' <https://platform.openai.com/docs/models/gpt-4o> (last accessed 2025/03/13).
29. OpenAI 'What is the ChatGPT model selector?' https://help.openai.com/en/articles/7864572-what-is-the-chatgpt-model-selector#h_80fc86948a (last accessed 2025/06/12).
30. Parker v Forsyth (2023)1 1585/20 ZAGPRD.
31. Petse, I & Phindelo, U 'The Influence of ChatGPT-generated Data on the Administration of Justice in South Africa' (2025) 5 Turf Law Journal 1-14.
32. Protection of Personal Information Act 4 of 2013.
33. Reece Rogers 'Reduce AI Hallucinations With This Neat Software Trick' 2024 https://www.wired.com/story/reduce-ai-hallucinations-with-rag/?utm_source=chatgpt.com (last accessed 2025/06/12).
34. S v Naidoo 1998 (1) SACR 479 (NPD)
35. Schindlers Attorneys, 'Sarah – Schindlers' Adjudication AI Engine Hands Down Judgment in Favour of the Legalisation of Psilocybin (Psychedelic Mushrooms) in South Africa' 22 April 2024 <https://www.schindlers.co.za/sarah-schindlers-adjudication-ai-engine-hands-down-judgement-in-favour-of-the-legalisation-of-psilocybin-psychedelic-mushrooms-in-south-africa/> (last accessed 2025/05/15)
36. Schweitzer, S & Conrads, M 'Correction to the digital transformation of jurisprudence: an evaluation of ChatGPT-4's applicability to solve cases in business law' (2024) Artif Intell Law.
37. Smuts v Botha 2022 (2) SA 425 (SCA).
38. Supreme Court Act 59 of 1959 as amended. GNR 5124, G. 51056 of 16 August 2024 (with effect from 20 September 2024).
39. Thaldar, DW 'How effectively can ChatGPT-4 draft data transfer agreements for health research?' (2025) Humanit Soc Sci Commun 12, 524.
40. The Constitution of the Republic of South Africa, 1996.
41. The South African Judicial Education Institute Act 14 of 2008.
42. Theron, L 'Refreshing And Enhancing Judgment Writing Skills' The South African Judicial Education Journal Volume 5 Issue 1 2022 Pages 4–5, Section IV.
43. Tomlinson, B, Torrance, AW, & Black, RW 'ChatGPT and Works Scholarly: Best Practices and Legal Pitfalls in Writing with AI' (2023) 76 SMU L Rev Forum 108.
44. United Nations Bangalore Principles and Commonwealth Latimer House Principles ECOSOC 2006/23.