Proceedings of the Design Society, Volume 6: DESIGN 2026
https://doi.org/10.1017/pds.2026.10593

Can large language models understand engineering design patents? An exploratory study

Pingfei Jiang✉, Yuxuan Wang and and Ji Han

INDEX, University of Exeter, United Kingdom

✉ P.Jiang@exeter.ac.uk

ABSTRACT: Patents contain valuable design insights, yet manual analysis remains time-consuming and complex. This study explores Large Language Models' capacity to automate patent analysis for engineering design. GPT-5 and Gemini 2.5 Pro were evaluated across Motivation, Novelty, and Key Invention Features using three patents and expert evaluators assessed outputs through Accuracy & Fidelity, Comprehensiveness, and Analytical Depth. Results indicate LLMs demonstrate proficiency in feature synthesis but exhibit inferential limitations in motivation analysis, underscoring the necessity for human oversight.

KEYWORDS: artificial intelligence (AI), engineering design, design knowledge, patent analysis, large language model (LLM)

1. Introduction

Intellectual Property (IP) rights, particularly patents, constitute a substantial component of modern enterprises' intangible asset portfolios. When organisations develop novel technologies and secure patent protection, these legal instruments prevent unauthorised use while promoting further innovation and market exploitation (Trappey et al., 2013). The World Intellectual Property Organization (WIPO) estimates that 90-95% of global technological innovations are documented within patent databases (Chen, 2009). Recent data indicate that worldwide patent applications exceeded 3.55 million in 2023, representing continued growth in intellectual property activity (WIPO, 2025). Patent documents contain extensive technical information unavailable through other sources, making them valuable resources for both academic research and industrial applications (Fantoni et al., 2013). Patent analysis provides organisations with multiple strategic benefits, including identification of technological trends, forecasting of developmental trajectories, and mitigation of infringement risks (Jiang et al., 2023). However, the consistent growth in patent applications has increased the complexity of analysis while simultaneously increasing intellectual property disputes, emphasising the critical need for enhanced prior art awareness during design processes (Sorce et al., 2018).

Despite their importance, designers inadequately integrate patent analysis into their product development workflows. This limitation stems primarily from patents being drafted by legal professionals using specialised terminology that makes them difficult and time-consuming to comprehend (Jiang et al., 2018). Consequently, patent analysis within engineering design contexts remains a vital research area, aimed at improving designer comprehension and providing competitive intelligence while preventing intellectual property conflicts. The substantial volume and consistent growth of patent documentation make manual analysis impractical at scale (Jiang et al., 2021). Critical design information is often fragmented across abstracts, background information, detailed descriptions, and claims, necessitating a resource-intensive interpretation to construct a comprehensive technical understanding.

Recent advances in artificial intelligence, particularly Large Language Models (LLMs), offer promising solutions to these persistent challenges. Models from the Generative Pre-trained Transformer (GPT) family demonstrate remarkable capabilities in processing, understanding, and generating sophisticated human language, achieving performance comparable to human experts across various benchmarks

https://doi.org/10.1017/pds.2026.10593 Published online by Cambridge University Press

(Brown et al., 2020). Their capacity for semantic comprehension and analysis of complex textual relationships has driven rapid adoption across specialised professional fields, from medical research to legal case analysis (Zhao et al., 2023). Within engineering design domains, LLMs have demonstrated value in data analysis, concept generation, and design evaluation (Regenwetter et al., 2022).

Large Language Models (LLMs) are increasingly transforming patent-related workflows by automating document drafting, enhancing search capabilities, and optimising analytical processes for intellectual property professionals (Bui, 2025). Legal practices are adopting these technologies due to their ability to process, analyse, and generate high-quality patent documentation efficiently and accurately (Shomee et al., 2025). As LLMs systems become embedded in design and innovation processes, the question is no longer whether they can process technical documents, but how in-depth they can understand and interpret design intent. Understanding this capacity is crucial for integrating LLMs into design reasoning and innovation strategy. Therefore, this study aims to present an exploratory investigation of commercially available LLMs' capabilities in engineering design patent document understanding and analysis. The focus centres on three design aspects: Motivation, Novelty, and Key Invention Features, rather than statistical metrics such as patent classification numbers or citations.

2. Related work

This section reviews the intersection of patent analysis, engineering design, and artificial intelligence. The literature was identified through a systematic search of databases, including Google Scholar, Scopus and Web of Science. Keywords utilised included 'Large Language Models', 'Engineering Design', and 'Patent Analysis'.

2.1. Patent analysis for engineering design

While patents function primarily as legal instruments for Intellectual Property protection, they simultaneously constitute a valuable repository of design case studies, providing valuable resources for innovation, education, and competitive analysis (Ernst, 2003). For engineering designers, the ability to systematically navigate and interpret this information is a critical competency. It provides direct insight into the state-of-the-art, prevents the redundant reinvention of existing solutions, and reveals pathways for future technological advancement (Jiang et al., 2022). However, superficial keyword searches or cursory document reviews are insufficient for extracting meaningful design insights, necessitating more structured and methodical analytical approaches (Cascini et al., 2004).

Early efforts to bridge this gap focused on creating systematic, human-driven frameworks to guide designers. These initial methodologies provided structured processes for generating novel design concepts that represent variations of existing patents while carefully avoiding infringement (Hsu et al., 2010). This line of inquiry evolved to include more sophisticated strategies for "design-around" activities, which relied on functional analysis to systematically circumvent existing intellectual property (Cheng et al., 2016). A key objective of this time was to develop formal representation schemes, particularly for engineering design patents, that helped designers deconstruct the working principles of an invention to better understand its core function and potential areas of conflict (Jiang et al., 2017).

A significant evolution in the field occurred with the shift from process-centric frameworks to more automated, computational approaches aimed at extracting deeper semantic meaning from patent text. A critical step in this direction involved creating functional representations of geometric interactions directly from patent claims, producing novel semantic annotations to identify emerging design conflicts (Atherton et al., 2017). This leads the way to the broader application of Natural Language Processing (NLP) techniques. Initial integrations of NLP focused on extracting linguistic features to improve patent classification (Chiarello et al., 2018) and were extended to interpret an inventor's underlying design intentions through the lens of established frameworks like axiomatic design (Li & Tate, 2019). Later, the development of large-scale semantic networks based on NLP started to emerge, such as TechNet, which represents engineering design knowledge derived from patent data in a structured, machine-readable format (Sarica et al., 2020), enabling the automated generation of design insights (Jiang et al., 2021).

More recent work has focused on integrated methodologies that synthesise computational approaches with holistic design processes. Such methods combine techniques like reverse engineering and patent circumvention into multi-phase workflows that encompass problem analysis, redesign, and implementation (Akerdad et al., 2022). The current frontier of this research seeks to automate the

https://doi.org/10.1017/pds.2026.10593 Published online by Cambridge University Press

identification of an invention's core motivation and explicitly connect it to the technical specifications and structures detailed in the patent (Jiang et al., 2023).

2.2. Large language models for patent analysis

The integration of Large Language Models (LLMs) in patent analysis represents an emerging and rapidly evolving field of research, opening new possibilities in areas such as novelty assessment, prior art comparison, and automated document drafting.

Initial research seeks to overcome the inherent complexity of patent language by integrating structured knowledge. Zhang et al. (2023), for instance, combined LLMs with knowledge graphs to enhance the semantic retrieval and classification of patent documents. This foundational work has been extended to more specialised tasks, such as multi-label hierarchical patent classification, where LLM-leveraging architectures have demonstrated substantial performance gains over previous models (Rafieian & Vázquez, 2025). The quality of input data has also been identified as a critical factor, where Yoshikawa and Krestel (2025) found that leveraging AI-generated summaries can significantly improve classification accuracy, particularly in challenging multilingual contexts.

Beyond analysing existing corpora, a substantial body of research has explored the generative capabilities of LLMs for creating patent content. A key development in this area is the creation of domain-specific models, such as PatentGPT, which was engineered to generate novel patent concepts and text, demonstrating superior performance on intellectual property benchmarks compared to general-purpose models (Ren et al., 2025). This has enabled more targeted applications, including the automated drafting of technically complex patent claims. While fine-tuned LLMs can replicate correct writing styles and technical dependencies, studies emphasise that human oversight remains indispensable for ensuring legal and technical precision (Chen & Pan, 2025). This necessity for expert validation has prompted the development of Human-in-the-Loop (HITL) systems, which integrate generative AI into professional workflows to enhance efficiency while maintaining rigorous legal and ethical standards (Bui, 2025). For less legal-sensitive tasks such as generating patent abstracts, state-of-the-art models have proven capable of producing outputs comparable in quality to those written by human experts (Shomee et al., 2025).

Another popular application involves leveraging LLMs for assessing patentability, particularly novelty. Researchers have framed this application as a textual entailment problem, creating specialised datasets to benchmark model performance in this area (Lee et al., 2024). Others have designed experiments to test if LLMs can mimic the judgment of patent examiners by comparing claims against prior art, with findings indicating that generative models can capture these complex relationships with reasonable accuracy (Ikoma & Mitamura, 2025). To ensure these automated judgments are not "black boxes", Jang et al. (2023) developed an explainable AI framework that incorporates patent claim structures to provide transparent justifications for its novelty assessments. Recognising the limitations of standard text-generation metrics, Yoo et al. (2025) introduced PatentScore, a multi-dimensional framework that assesses the structural, legal, and semantic quality of LLM-generated claims, establishing a more reliable standard for evaluating AI-generated patent content.

Collectively, these advancements underscore a research trajectory focused on utilising the patent as a legal document. The primary objectives have been to enhance the efficiency of search, drafting, and novelty assessment from a legal perspective. However, as Jiang and Goetz (2025) stated, the full potential of LLMs in the patent domain remains underexplored. The current body of work has largely overlooked the patent's role as an artefact of engineering design. While some studies have explored tangential connections between product features and inventive value (Ali & Carlo, 2023), a systematic exploration of engineering design patents reflecting invention motivation, technical problem framing, and the creative reasoning embodied in the design remains limited. While these advancements have transformed patent informatics, their orientation remains predominantly legal and textual. Engineering design research, on the other hand, seeks to extract functional, motivational, and conceptual insights from patents: dimensions largely overlooked by current LLM-driven studies.

2.3. Research gap and opportunity

A review of the literature reveals two parallel and disconnected streams of research. Engineering design scholarship has progressively emphasised the need for deep, conceptual analysis of patents, prioritising understanding of the core inventive rationale, specifically the motivation driving an invention and the technical novelty of its solution, as prerequisites for meaningful innovation. Conversely, Large Language

https://doi.org/10.1017/pds.2026.10593 Published online by Cambridge University Press

Model applications in patent informatics have been overwhelmingly driven by automation of legal and textual tasks, such as prior art retrieval, document classification, and claim generation. While these applications have demonstrated significant value, their focus remains on processing patents as legal instruments rather than as sources of design insight.

A significant research gap, therefore, exists at the intersection of these fields. It remains an open question whether current LLMs, which excel at text-based pattern recognition and generation, possess the capacity for deeper inferential reasoning required to analyse patents as artefacts of engineering design. This study directly addresses this gap through an exploratory investigation that assesses the capability of modern LLMs to perform holistic design analysis of engineering design patents.

Specifically, this paper investigates the extent to which these models can: (1) identify an invention's motivation, (2) articulate its core novelty, and (3) establish key invention features and their interrelationships. This investigation aims to determine whether LLMs can serve as effective tools for design-oriented patent analysis, potentially transforming how designers interact with patent repositories as sources of technical knowledge and innovation inspiration.

3. Research methodology

The research methodology adopted in this paper is illustrated in Figure 1. The methodology employed in this study centres on comparative analysis between LLM-generated patent analysis and expert-derived evaluations. This approach enables quantitative assessment of LLM performance in patent analysis tasks that mirror the cognitive processes employed by design engineers.

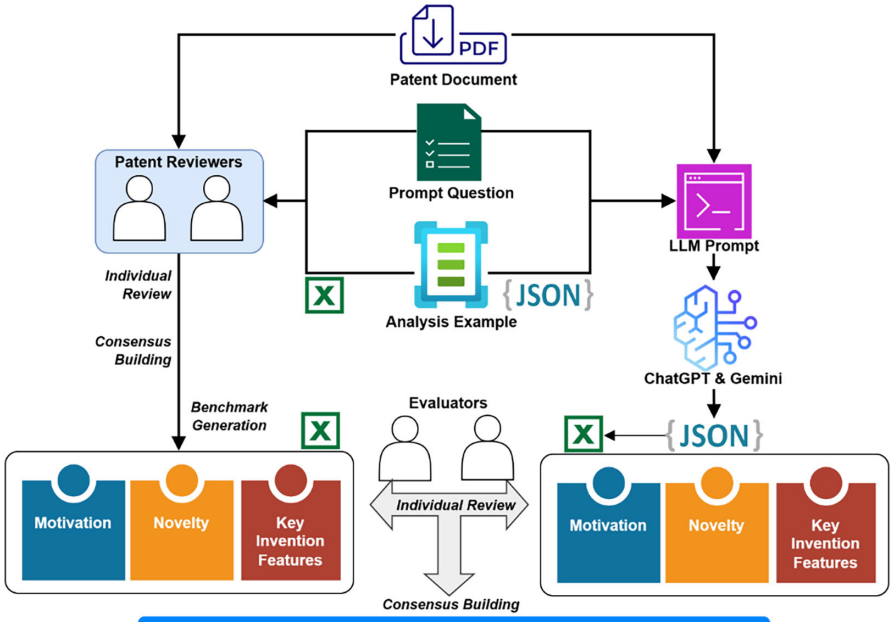


Figure 1. Research methodology for the comparative analysis

To ensure consistent analytical conditions, patent documents were downloaded and stored locally, providing standardised access for both experts and LLMs. A patent analysis example was developed to maintain consistency between expert and LLM outputs, with format adaptations to suit each application: spreadsheet format for human experts and JSON structure for LLM processing. The JSON template (see Figure 2) was integrated into the LLM prompt, with the analytical outputs derived from experts on a separate patent.

To minimise potential biases while maintaining analytical feasibility, three patents spanning distinct engineering design applications were selected. The selection criteria prioritised granted US patents

https://doi.org/10.1017/pds.2026.10593 Published online by Cambridge University Press

representing a spectrum from predominantly mechanical systems with minimal electronics to sophisticated mechatronic products:

- US11647866 – Grinder/Doser device for coffee beans
- US10631685 – Toaster
- US8529203 – Fan assembly

An expert evaluation was conducted by two engineering designers, each with over ten years of experience and specialised expertise in engineering design and patent analysis. Each expert independently analysed the selected patents, generating individual assessments across the three design aspects. A consensus-building approach was subsequently employed, wherein the experts engaged in structured discussions to reconcile differences and establish agreed-upon outcomes for each patent. All expert evaluations were systematically recorded in spreadsheet format to enable direct comparison with LLM-generated analyses.

The LLM analysis protocol paralleled the expert protocol. Complete patent documents were provided to the LLMs in PDF format to ensure that LLMs did not rely on internal prior knowledge or web-search hallucinations. To mitigate model-specific biases, two state-of-the-art LLMs were employed: OpenAI's GPT-5 and Google's Gemini 2.5 Pro.

To maintain response consistency, the LLMs were guided by a structured prompt featuring JSON output constraints. This prompt included specific instructions for deconstructing the 'Motivation', 'Novelty', and 'Key Invention Features' of the patents. Due to space constraints, the full prompt specification, JSON schema, and the primary patent corpus are available in an open-access repository at https://github.com/PingfeiJiang/DESIGN2026. LLM-generated JSON outputs were manually processed and transferred to spreadsheet format to enable systematic comparison with expert analysis.

While exact time-on-task was not recorded via a controlled stopwatch, the experts reported an average of 30-45 minutes per patent for full manual deconstruction and consensus building. In contrast, the LLMs (GPT-5 and Gemini 2.5 Pro) generated the structured JSON outputs in under 15 seconds per document. This represents a potential timesaving of over 90%, assuming human oversight is maintained for verification. A cross-comparison was conducted between reconciled expert results and LLM responses by the same two experts across all three patents. Evaluation employed a five-level scale applied to three analytical aspects, with standardised metrics developed for each design aspect. The evaluation framework incorporated three assessment metrics for each analytical aspect: Accuracy & Fidelity, Comprehensiveness, and Analytical Depth. Metric-specific guide questions were formulated to ensure appropriate assessment protocols (see Table 1). This systematic approach enables quantitative assessment of LLM performance relative to expert-level patent analysis across multiple dimensions. Initial independent analysis by the two experts showed an agreement rate of approximately 80%, particularly regarding factually verifiable technical features. Differences were primarily limited to the interpretative 'Analytical Depth' of the motivation sections, which were subsequently resolved through structured reconciliation to establish the final benchmark.

Table 1. Guide questions and scaled ratings for comparative analysis between expert analysis and LLM responses

Accuracy & Fidelity	Motivation	**To what extent is the motivation information factually correct and verifiable within the "Background" or "Summary" sections of the patent?** *1: Contains significant factual errors or fabrications (hallucinations).* *5: All statements are factually correct and directly supported by the source text.*
	Novelty	**To what extent is the articulated novelty factually correct and supported by the patent text?** *1: Contains significant errors or misinterpretations of the novelty.* *5: The described novelty is a faithful and accurate representation of the patent.*
	Key Invention Features	**To what extent are the identified features and their descriptions factually correct and verifiable in the patent text?** *1: Contains significant errors in identifying parts or their functions.* *5: All identified features are verifiably correct.*

(*Continued*)

https://doi.org/10.1017/pds.2026.10593 Published online by Cambridge University Press

Table 1. Continued

Compre -hensiveness	Motivation	**How comprehensively did the LLM capture the full spectrum of motivations identified by the expert analysis?** *1: Missed most of the key points and only captured the most obvious one.* *5: Captured all distinct points identified in the expert analysis.*
	Novelty	**How comprehensively did the LLM capture the full scope of the invention novelty?** *1: Captured only the primary novelty and missed key supporting inventive concepts.* *5: Captured the full hierarchy of novel concepts.*
	Key Invention Features	**How comprehensively did the LLM identify the full set of distinct elements and limitations listed within the primary independent claim?** *1: Missed a majority of the claim's essential elements.* *5: Identified all or nearly all of the distinct elements and limitations.*
Analytical Depth	Motivation	**Does the analysis just list problems, or does it frame them within a clear context?** *1: A simple, flat list of problems with no context.* *5: Clear contextualisation of problems that are easy to understand.*
	Novelty	**Does the analysis just restate patent claims or summary, or does it synthesize them into a coherent and understandable inventive concept?** *1: A literal restatement of patent claims or summary.* *5. A clear, well-synthesized explanation of the core inventive principle that are easy to understand.*
	Key Invention Features	**Does the analysis just list the nominal elements from the claim, or does it correctly interpret the structural and functional relationships that define the invention?** *1: The output is a simple list of the nouns in the claim, ignoring the relationships that connect them.* *5: The output accurately describes the defined structure, capturing how the elements are connected to each other and what their specified functions are, as defined by the claim language.*

4. Results

The comparative evaluation results for GPT-5 and Gemini 2.5 Pro across all three patents are presented in Table 2. Both models demonstrated substantial proficiency in patent analysis tasks, with performance patterns revealing distinct strengths and limitations relative to expert benchmarks.

Both models achieved exceptional performance in technical feature identification, obtaining perfect scores across all analytical metrics for every patent examined. This includes the conceptually complex Bladeless Fan patent, indicating robust capability in accurately identifying and synthesising complete feature sets that constitute patent inventions. The models demonstrated particular strength in interpreting structural and functional relationships between claim elements, producing integrated system descriptions rather than just listing component parts stated in the independent claim.

Near-perfect scores were consistently achieved for Analytical Depth, demonstrating that the models perform sophisticated information synthesis rather than simple text extraction or keyword matching. This performance suggests LLMs' capacity for technical reasoning that approaches genuine comprehension of inventive principles.

The most significant performance variation was observed in Motivation identification, where both models, particularly GPT-5, received notably lower Accuracy & Fidelity scores. GPT-5 scored 2 points for both Coffee Grinder and Bladeless Fan patents, while Gemini 2.5 Pro achieved marginally higher scores (4, 5, 3, respectively). This pattern indicates several possible phenomena: hallucination, where models generate plausible but unsubstantiated content, inferential reasoning working backwards from

https://doi.org/10.1017/pds.2026.10593 Published online by Cambridge University Press

solutions to deduce motivations, or more broadly, extended thinking that goes beyond explicitly stated facts in the patent documents. The exact mechanism remains unclear, though the consistent deviation from source material suggests these models may be supplementing documented information with generated content that, while potentially logical, compromises factual fidelity when it diverges from inventors' actual stated purposes.

Table 2. Ratings for GPT-5 and Gemini 2.5 Pro for all three patents analysed, compared to expert analysis

	Motivation					
	Accuracy & Fidelity		**Comprehensiveness**		**Analytical Depth**	
	GPT-5	Gemini 2.5 Pro	GPT-5	Gemini 2.5 Pro	GPT-5	Gemini 2.5 Pro
Coffee Grinder	2	4	5	5	5	5
Toaster	4	5	5	5	5	5
Bladeless Fan	2	3	4	3	5	5
	Novelty					
	Accuracy & Fidelity		**Comprehensiveness**		**Analytical Depth**	
	GPT-5	Gemini 2.5 Pro	GPT-5	Gemini 2.5 Pro	GPT-5	Gemini 2.5 Pro
Coffee Grinder	5	5	5	5	5	5
Toaster	5	5	4	5	4	5
Bladeless Fan	4	4	4	3	5	5
	Key Invention Features					
	Accuracy & Fidelity		**Comprehensiveness**		**Analytical Depth**	
	GPT-5	Gemini 2.5 Pro	GPT-5	Gemini 2.5 Pro	GPT-5	Gemini 2.5 Pro
Coffee Grinder	5	5	5	5	5	5
Toaster	5	5	5	5	5	5
Bladeless Fan	5	5	5	5	5	5

LLMs' performance remained largely consistent across patents of varying conceptual complexity. The Bladeless Fan, representing the most technically challenging patent, yielded comparable scores to simpler patents. However, subtle performance degradation was observed in Novelty analysis for the Bladeless Fan, with Accuracy & Fidelity and Comprehensiveness scores dropping to 4 and 3, respectively. This suggests that increased conceptual complexity may cause models to capture primary inventive concepts while potentially overlooking secondary or more nuanced technical details identified by human experts.

While both models performed at state-of-the-art levels, systematic differences emerged in their analytical approaches. Gemini 2.5 Pro consistently demonstrated superior fidelity to source documents, achieving higher Accuracy & Fidelity scores across multiple evaluation criteria. Conversely, both models exhibited comparable analytical depth capabilities, suggesting this represents a mature competency in current large language model architectures.

5. Discussion

The evaluation results demonstrate that contemporary LLMs exhibit substantial capability for engineering design patent analysis, particularly regarding technical feature identification and analytical synthesis. Consistently high performance in articulating novelty and interpreting interrelationships between key invention features suggests that these models transition beyond basic text processing toward a functional understanding of patented designs. This capability suggests a potential paradigm shift in engineering design practice, wherein LLMs serve as analytical partners rather than mere search or summarisation tools. Such analytical capacity could significantly accelerate critical design activities such as competitor analysis, technology landscaping, and domain knowledge acquisition. By providing rapid, coherent interpretations of complex documentation, these models allow human experts to bypass the time-intensive task of initial document comprehension and redirect cognitive resources toward strategic decision-making and creative ideation.

https://doi.org/10.1017/pds.2026.10593 Published online by Cambridge University Press

However, the performance degradation observed in motivation analysis represents the study's most significant finding, revealing a fundamental limitation in current LLM approaches to patent understanding. These errors suggest a failure mode termed 'Inferential Fabrication,' in which models prioritise logical narrative over strict source grounding. This weakness likely stems not from insufficient capability but from the models' default operational mode when confronted with ambiguous or sparsely documented information. Patent background and summary sections, from which motivation is typically derived, often represent the least structured portions of patent documents. When faced with incomplete information, LLMs appear to employ backward reasoning from documented solutions to infer probable motivations. This behaviour demonstrates sophisticated cognitive processing but fundamentally compromises the factual fidelity essential for rigorous patent analysis. The tendency suggests that current LLM architectures prioritise narrative coherence and logical consistency over strict adherence to source material. This operational characteristic poses a critical challenge for LLM deployment in domains that require high factual grounding, as outputs may appear plausible and well-reasoned yet contain unsubstantiated content that requires human verification.

As an exploratory investigation, this study has several limitations. The sample size of three patents, while enabling detailed quantitative analysis, precludes statistically significant conclusions about the performance of LLMs across the entirety of the engineering design domain. Additionally, the expert panel composition of two individuals represents a limited benchmark, and the incorporation of legal expertise alongside engineering perspectives could strengthen the evaluative framework. Current assessments utilise expert-assigned Likert scales, suggesting that future studies should integrate computational measures, such as semantic similarity or automated grounding checks via specific claim citations, to provide more granular accuracy metrics.

The findings suggest several promising future research directions. Large-scale validation studies employing diverse patent portfolios across multiple engineering domains and expanded expert panels would provide statistical robustness to these preliminary observations. The identified deficiencies in motivation extraction offer opportunities for targeted model refinement, particularly through fine-tuning that emphasises factual extraction over inferential reasoning. To mitigate fabrication, future protocols should implement grounding requirements that task models with citing specific paragraph or claim spans. Additionally, interactive conversational AI systems represent a compelling direction for human-in-the-loop workflows, enabling designers to pursue follow-up queries and verify claims in real-time.

From a practical standpoint, this framework enables the rapid prototyping of design knowledge. Practitioners can utilise LLM-generated JSON features to map competitor product architectures against their own during early-stage ideation. By synthesising motivation and novelty data, design teams can identify market 'white spaces', areas where existing intellectual property fails to address specific user needs, without the traditional friction of interpreting dense legal terminology. Our initial metrics demonstrate that LLMs can generate these structured outputs in under 15 seconds, representing a potential efficiency gain of over 90% compared to manual deconstruction. Future work will transition from this conceptual evaluation to a comprehensive workflow study to measure time-savings, decision quality, and analyst trust in professional competitive landscaping environments.

6. Conclusion

This study presents an exploratory investigation of state-of-the-art large language model capabilities in engineering design patent analysis. The evaluation framework examined three critical aspects to understand a patented invention: Motivation, Novelty, and Key Invention Features, through comparative assessment between expert analysis and LLM performance (GPT-5 and Gemini 2.5 Pro) across metrics of Accuracy & Fidelity, Comprehensiveness, and Analytical Depth. The results reveal a foundational insight into the operational logic of LLMs in this domain: they function not as passive text extractors but as active analytical synthesisers. This is evidenced by their remarkable proficiency in deconstructing a patent's technical novelty and interpreting the complex relationships within its claims, a capability that points toward a genuine comprehension of inventive principles. This study addresses the disconnection between the prevailing legalistic applications of LLMs in patent informatics and the conceptual needs of engineering designers.

Nevertheless, a critical limitation emerged in motivation analysis, where models exhibited tendencies toward inference rather than extraction, generating plausible but potentially inaccurate content. This behaviour indicates that current LLM architectures prioritise narrative coherence over strict textual fidelity, necessitating human oversight to ensure analytical rigour and source document grounding.

https://doi.org/10.1017/pds.2026.10593 Published online by Cambridge University Press

Future research should address the identified limitations through larger, more diverse patent datasets and expanded expert panels to achieve statistical significance. The integration of human-AI collaborative workflows represents a promising direction for leveraging LLM analytical capabilities while maintaining the factual rigour essential for engineering design applications.

References

Akerdad, M., Aboutajeddine, A., & Elmajdoubi, M. (2022). Development of a redesign process based on reverse engineering and patent circumvention. *International Journal on Interactive Design and Manufacturing (IJIDeM)*, 17, 261–277. https://doi.org/10.1007/s12008-022-00918-w

Ali, Y., & Carlo, L. (2023, 24–28 July). Measuring Patent Novelty Using Natural Language Processing International Conference on Engineering Design, ICED23, Bordeaux, France.

Atherton, M., Jiang, P., Harrison, D., & Malizia, A. (2017). Design for invention: annotation of functional geometry interaction for representing novel working principles. *Res Eng Des*, 29(2), 245–262. https://doi.org/10.1007/s00163-017-0267-2

Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., . . . Amodei, D. (2020). Language Models are Few-Shot Learners 34th Conference on Neural Information Processing Systems (NeurIPS 2020), Vancouver, Canada.

Bui, L. V. (2025). Advancing patent law with generative AI: Human-in-the-loop systems for AI-assisted drafting, prior art search, and multimodal IP protection. *World Patent Information*, 80. https://doi.org/10.1016/j.wpi.2025.102341

Cascini, G., Fantechi, A., & Spinicci, E. (2004). Natural Language Processing of Patents and Technical Documentation. International Workshop on Document Analysis Systems, DAS 2004: Document Analysis Systems VI,

Chen, F.-C., & Pan, C.-L. (2025). Evaluating application of large language models to biomedical patent claim generation. *World Patent Information*, 80. https://doi.org/10.1016/j.wpi.2025.102339

Chen, R. (2009). Design patent map visualization display. *Expert Systems with Applications*, 36(10), 12362–12374. https://doi.org/10.1016/j.eswa.2009.04.049

Cheng, S., Mi, J., Zhao, R., & Yang, X. (2016). Patent Product Innovation Re-Design Based on Function Analysis. *Recent Patents on Mechanical Engineering*, 9(1).

Chiarello, F., Cimino, A., Fantoni, G., & Dell'Orletta, F. (2018). Automatic users extraction from patents. *World Patent Information*, 54, 28–38. https://doi.org/10.1016/j.wpi.2018.07.006

Ernst, H. (2003). Patent information for strategic technology management. *World Patent Information*, 25, 233–242. https://doi.org/10.1016/S0172-2190(03)00077-2

Fantoni, G., Apreda, R., Dell'Orletta, F., & Monge, M. (2013). Automatic extraction of function–behaviour–state information from patents. *Advanced Engineering Informatics*, 27(3), 317–334. https://doi.org/10.1016/j.aei.2013.04.004

Hsu, Y. L., Hsu, P. E., Hung, Y. C., & Xiao, Y. D. (2010, August 15–18). Development and Application of a Patent-based Design Around Process. Proceedings of the ASME 2010 International Design Engineering Technical Conferences & Computers and Information in Engineering Conference, Montreal, Quebec, Canada.

Ikoma, H., & Mitamura, T. (2025). Can AI Examine Novelty of Patents?: Novelty Evaluation Based on the Correspondence between Patent Claim and Prior Art. arXiv:2502.06316v1 https://doi.org/10.48550/arXiv.2502.06316

Jang, H., Kim, S., & Yoon, B. (2023). An eXplainable AI (XAI) model for text-based patent novelty analysis. *Expert Systems with Applications*, 231. https://doi.org/10.1016/j.eswa.2023.120839

Jiang, L., & Goetz, S. M. (2025). Natural language processing in the patent domain: a survey. *Artificial Intelligence Review*, 58(214). https://doi.org/10.1007/s10462-025-11168-z

Jiang, P., Atherton, M., Harrison, D., & Malizia, A. (2017, 21–25 August). Framework of mechanical design knowledge representations for avoiding patent infringement. 21st International Conference on Engineering Design, ICED 17, The University of British Columbia, Vancouver, Canada.

Jiang, P., Atherton, M., & Sorce, S. (2021, 16–20 August). Automated Functional Analysis of Patents for Producing Design Insight. Proceedings of the Design Society: International Conference on Engineering Design, Gothenburg, Sweden.

Jiang, P., Atherton, M., & Sorce, S. (2023). Extraction and linking of motivation, specification and structure of inventions for early design use. *Journal of Engineering Design*, 34(5-6), 411–436. https://doi.org/10.1080/09544828.2023.2227934

Jiang, P., Atherton, M., Sorce, S., Harrison, D., & Malizia, A. (2018). Design for invention: a framework for identifying emerging design–prior art conflict. *Journal of Engineering Design*, 29(10), 596–615. https://doi.org/10.1080/09544828.2018.1520204

https://doi.org/10.1017/pds.2026.10593 Published online by Cambridge University Press

Jiang, S., Sarica, S., Song, B., Hu, J., & Luo, J. (2022). Patent Data for Engineering Design: A Critical Review and Future Directions. *Journal of Computing and Information Science in Engineering*, 22(6). https://doi.org/10.1115/1.4054802

Lee, R., Spangher, A., & Ma, X. (2024). PatentEdits: Framing Patent Novelty as Textual Entailment. arXiv:2411.13477v1. https://doi.org/10.48550/arXiv.2411.13477

Li, Z., & Tate, D. (2019, August 15–18). Automatic Function Interpretation: Using Natural Language Processing on Patents to Understand Design Purposes. Proceedings of the ASME 2010 International Design Engineering Technical Conferences & Computers and Information in Engineering Conference Montreal, Quebec, Canada.

Rafieian, B., & Vázquez, P.-P. (2025). Improved multi-label hierarchical patent classification using LLMs. World Patent Information, 81. https://doi.org/10.1016/j.wpi.2025.102356

Regenwetter, L., Nobari, A. H., & Ahmed, F. (2022). Deep Generative Models in Engineering Design: A Review. arXiv:2110.10863v4. https://doi.org/10.48550/arXiv.2110.10863

Ren, R., Ma, J., & Luo, J. (2025). Large language model for patent concept generation. *Advanced Engineering Informatics*, 65. https://doi.org/10.1016/j.aei.2025.103301

Sarica, S., Luo, J., & Wood, K. L. (2020). TechNet: Technology semantic network based on patent data. *Expert Systems with Applications*, 142. https://doi.org/10.1016/j.eswa.2019.112995

Shomee, H. H., Maity, S. K., & Medya, S. (2025). PATENTWRITER: A Benchmarking Study for Patent Drafting with LLMs. arXiv:2507.22387v1 https://doi.org/10.48550/arXiv.2507.22387

Sorce, S., Malizia, A., Jiang, P., Atherton, M., & Harrison, D. (2018). A Novel Visual Interface to Foster Innovation in Mechanical Engineering and Protect from Patent Infringement. IOP Conf. Series: Journal of Physics: Conf. Series 1004. https://doi.org/10.1088/1742-6596/1004/1/012024

Trappey, A. J. C., Trappey, C. V., Wu, C.-Y., Fan, C. Y., & Lin, Y.-L. (2013). Intelligent patent recommendation system for innovative design collaboration. *Journal of Network and Computer Applications*, 36(6), 1441–1450. https://doi.org/10.1016/j.jnca.2013.02.035

WIPO. (2025). WIPO IP Statistics Data Center. Retrieved 08/09 from https://www3.wipo.int/ipstats/key-search/search-result?type=KEY&key=201

Yoo, Y., Xu, Q., & Cao, L. (2025). PatentScore: Multi-dimensional Evaluation of LLM-Generated Patent Claims. arXiv:2505.19345. https://doi.org/10.48550/arXiv.2505.19345

Yoshikawa, N., & Krestel, R. (2025). Do large language models understand patents? Enhancing patent classification through AI-generated summaries. World Patent Information, 81. https://doi.org/10.1016/j.wpi.2025.102353

Zhang, C., Jian, S., Chao, L., Fan, C., Bo, L., & Chen, L. (2023, 15–18 December). A Semantic Understanding Method for Patent Text Based on Large Language Model 2023 IEEE the 7th Conference on Energy Internet and Energy System Integration, Hangzhou, China.

Zhao, W. X., Zhou, K., Li, J., Tang, T., Wang, X., Hou, Y., Min, Y., Zhang, B., Zhang, J., Dong, Z., Du, Y., Yang, C., Chen, Y., Chen, Z., Jiang, J., Ren, R., Li, Y., Tang, X., Liu, Z., . . . Wen, J.-R. (2023). A Survey of Large Language Models. arXiv:2303.18223v13

https://doi.org/10.1017/pds.2026.10593 Published online by Cambridge University Press