Perplexity AI

Grid operators in the age of autonomous energy systems:

How smart grids will transform daily operations

Introduction

The electrical grid stands at the cusp of its most significant transformation since its inception. As renewable energy sources proliferate and distributed energy resources (DERs) become commonplace, the traditional centralized control paradigm is giving way to a new era of Autonomous Energy Systems (AES). This evolution promises to fundamentally reshape how grid operators manage and maintain our power infrastructure [1].

The digital grid evolution

Today's power grid already bears little resemblance to its predecessor from just a decade ago. Digital substations, smart meters, and automated switching systems have laid the groundwork for more sophisticated control mechanisms. The integration of artificial intelligence and machine learning has begun to demonstrate the potential for autonomous decision-making in grid operations [2].

Autonomous Energy Systems represent the next logical step in this progression. These systems combine advanced AI algorithms, distributed intelligence, and real-time data analytics to create self-organizing, self-optimizing networks capable of managing power flow with minimal human intervention. By 2035, while AES will likely play a significant role in grid control, their implementation will be shaped by evolving regulatory frameworks. Many jurisdictions will likely maintain requirements for human oversight in critical decision-making processes, resulting in a hybrid model where advanced AI systems augment human operators rather than fully replace them. This approach will fundamentally alter the role of grid operators, shifting their focus towards strategic oversight, ethical decision-making, and regulatory compliance.

Adversities drive innovation

The imperative for autonomous systems stems from several pressing challenges facing today's grid operators. The integration of variable renewable energy sources has introduced unprecedented complexity in maintaining grid stability. For example, in one study by the Electric Reliability Council of Texas (ERCOT) system, increasing renewable penetration from 20% to 80%

The operator's role is not to make moment-to-moment decisions about power flow or voltage regulation - these are handled automatically by the AES

progressively reduced the voltage stability margin. At 80% penetration, all synchronous generators in the Austin region were displaced with PV, resulting in the smallest stability margin [4].

The proliferation of DERs, from rooftop solar to electric vehicle charging stations, has created millions of potential points of failure and optimization. Current SCADA systems struggle to monitor and control this level of granularity effectively. The 2023 Data Breach Investigations Report reveals that 20% of incidents and 11% of breaches analysed were linked to the Public Administration sector, which includes energy infrastructure, marking an increase from the previous year [5].

The International Energy Agency (IEA) projects that global electricity demand will grow by 2.1% annually until 2040, with renewables expected to account for 80% of the increase in global electricity generation over the same period [6]. This rapid growth and shift in energy sources underscore the urgent need for more advanced grid management systems.

The 2035 control room: A day in the life

Imagine walking into a grid control center in 2035. Gone are the walls of monitors displaying static SCADA screens. Instead, operators work in an immersive environment where holographic displays project real-time 3D visualizations of the entire grid network. AI assistants process millions of data points per second, presenting only the most critical information for human review.

At 06:00, the morning operator receives a handoff briefing not from the night shift, but from the AI system that has been monitoring grid conditions overnight. The system highlights three areas requiring human attention: a predicted transformer failure in a critical substation, identified through pattern analysis of vibration sensors and oil chemistry data; a weather system approaching that will affect solar generation in the afternoon; and a planned maintenance operation that requires human approval of the AI-generated workflow.

The operator's role is not to make moment-to-moment decisions about power flow or voltage regulation - these are handled automatically by the AES. Instead, they focus on validating the AI's strategic decisions and ensuring the system's responses align with broader operational goals.

At 14:00, the weather system arrives earlier than expected, causing a 30% drop in solar generation across a major metropolitan area. The AES responds instantaneously, activating demand response programs through smart building management systems, adjusting power flow paths to optimize transmission efficiency, and deploying stored energy from distributed battery systems.

The operator monitors these actions through augmented reality interfaces, which overlay AI decision pathways

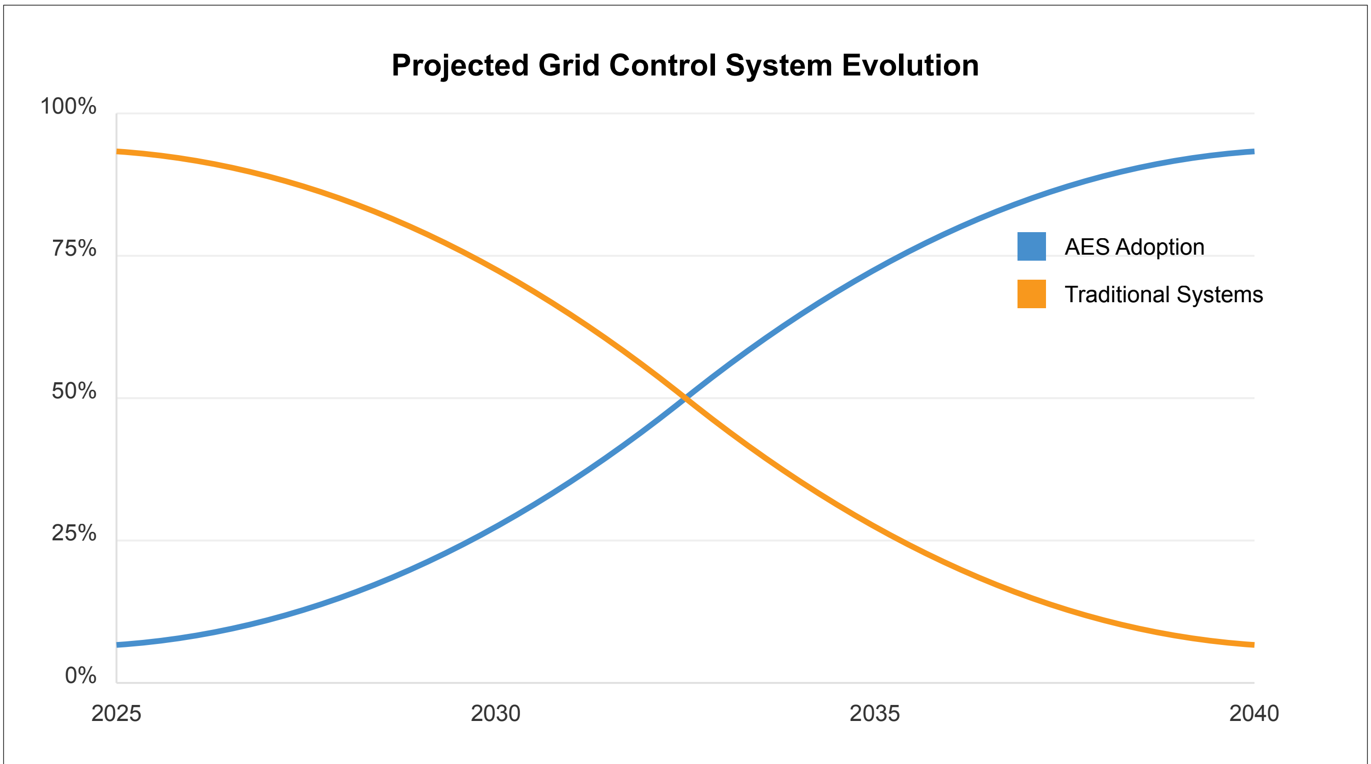

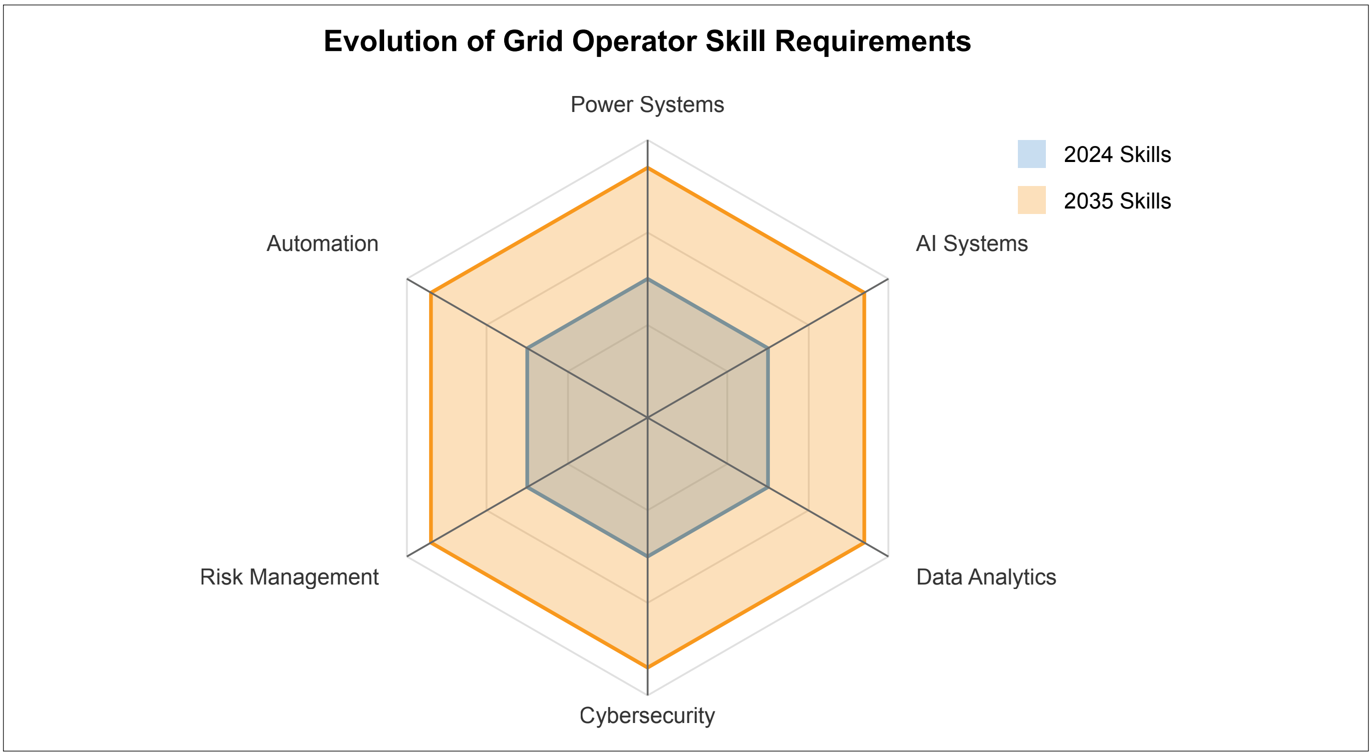

onto the physical grid infrastructure. Their role is to verify the system's decisions and intervene only if they identify potential issues the AI might have missed.

The evolution of human expertise

The transition to AES demands a new breed of grid operator. While deep understanding of power systems remains crucial, operators must develop additional competencies in AI systems management, data analytics, and risk management.

While deep understanding of power systems remains crucial, operators must develop additional competencies in AI systems management, data analytics, and risk management

In AI systems management, operators must be able to interpret and validate AI decisions, understand the limitations of machine learning models, and have experience in training and refining AI systems with new scenarios. Data analytics skills become paramount, with operators needing expertise in pattern recognition across complex datasets, predictive modeling, and scenario analysis.

Risk management takes on new dimensions, requiring advanced cybersecurity awareness, experience in contingency planning for AI system failures, and a deep understanding of cascading failure scenarios in autonomous systems.

Challenges and risk mitigation

The implementation of AES introduces new vulnerabilities that must be carefully managed. The increased digitalization and automation of grid operations expands the attack surface for cyber threats. Defense strategies must include AI-powered intrusion detection systems, quantum-resistant encryption for critical communications, regular penetration testing of autonomous systems, and human oversight of security protocols and anomaly detection.

Total dependence on autonomous systems creates new single points of failure. Mitigation strategies include maintaining manual override capabilities for critical systems, regular testing of fallback modes and degraded operation scenarios, redundant control systems with diverse implementation approaches, and continuous validation of AI decision-making processes.

Ethical considerations also come to the fore, as autonomous systems must balance competing priorities in ways that align with human values. Key considerations include fair distribution of resources during constraints, transparent decision-making processes, protection of consumer privacy, and equitable access to grid services.

The path forward

The transition to AES represents both an opportunity and a challenge for the power industry. Success will require significant infrastructure development, including the deployment of advanced sensors and control systems, enhancement of communication networks, implementation of edge computing capabilities, and integration of AI-ready hardware at all levels of the grid.

Workforce development becomes crucial, necessitating the creation of training programs for existing operators, development of new certification standards, establishment of clear career progression paths, and collaboration with educational institutions to prepare the next generation of grid operators.

A robust policy framework must also be established, including updated regulatory standards for autonomous systems, clear liability frameworks for AI decisions, cybersecurity requirements for AES implementations, and interoperability standards for distributed systems.

Conclusion

The future of grid operations lies in the successful integration of autonomous energy systems with human expertise. While AI and machine learning will handle the bulk of routine operations, human operators will remain essential for strategic oversight, ethical decision-making, and managing exceptional situations [3].

As we move toward this autonomous future, the industry must invest in both technological infrastructure and human capital. Success will require a careful balance between automation and human oversight, ensuring that our power grid becomes not just smarter, but also more reliable, secure, and equitable.

The grid operators of 2035 will be the guardians of this new system, combining traditional power systems knowledge with expertise in AI, data analytics, and risk management. Their role will be more critical than ever in ensuring the stable, efficient, and fair operation of our power infrastructure. As we embrace this transformative technology, we must remain mindful of the challenges and work collaboratively to create a resilient, sustainable energy future for all.

References

[1] McKinsey & Company, "How grid operators can integrate the coming wave of renewable energy," McKinsey & Company, 2024. [Online]. Available: https://www.mckinsey.com/industries/electric-power-and-natural-gas/our-insights/how-grid-operators-can-integrate-the-coming-wave-of-renewable-energy

[2] National Renewable Energy Laboratory, "Autonomous Energy Systems," NREL, 2024.

[3] SmythOS, "Autonomous Agents in Energy Management," SmythOS, 2024. [Online]. Available: https://smythos.com/artificial-intelligence/autonomous-agents/autonomous-agents-in-energy-management/

[4] National Renewable Energy Laboratory, "Operating Reserves and Variable Generation: A Comprehensive Review of Current Strategies, Challenges, and Innovations," U.S. Department of Energy Office of Scientific and Technical Information, Tech. Rep. NREL/TP-5500-55630, Mar. 2021.

[5] Verizon, "2024 Data Breach Investigations Report," Verizon Enterprise, 2024. [Online]. Available: https://www.verizon.com/business/resources/reports/dbir/

[6] International Energy Agency, "Electricity Grids and Secure Energy Transitions," IEA, 2024.

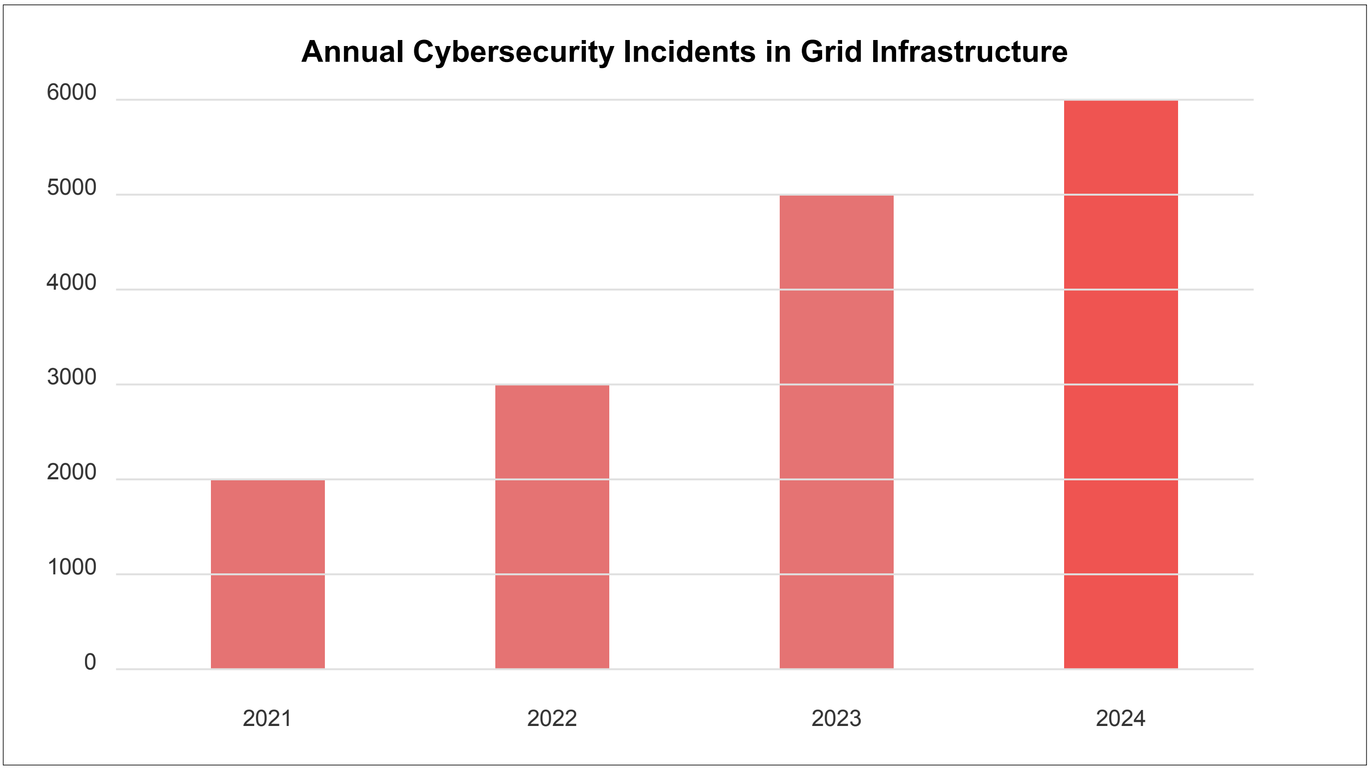

AI article rejected

Experiment in autonomous scientific writing

Abstract

This article documents an attempted experiment in generating a technical paper entirely by means of Artificial Intelligence (AI). I set out to test the boundaries of AI in technical writing and submitted the article [1] to peer review for *Switchgear Magazine*. In accord with the normal double-blind process, I did not disclose the article's non-human origin.

My journey through this process revealed both the capabilities and limitations of current AI systems. I uncovered patterns in how AI systems approach technical writing and where they fall short of professional standards, namely, peer review detected an artificial voice.

Introduction

The rapid advancement of AI technologies has sparked debate about their potential role in professional and academic writing. While AI writing assistants have become commonplace, their capability to produce publication-worthy technical content remains largely untested. I designed this experiment to push these boundaries, deliberately choosing a complex technical topic – Autonomous Energy Systems – that would challenge AI's ability to demonstrate deep domain knowledge and technical rigor.

In a turn of phrase attributed to Francis Bacon, "AI is a good servant, but a poor master."

The experiment went beyond simple content generation. I strove to understand how AI systems handle complex requirements of technical writing: supporting claims with appropriate citations, generating meaningful visualisations, maintaining consistent terminology, and responding to peer review feedback. This experiment evaluates AI's capacity to independently produce scholarly works meeting technical publication criteria.

to be continued

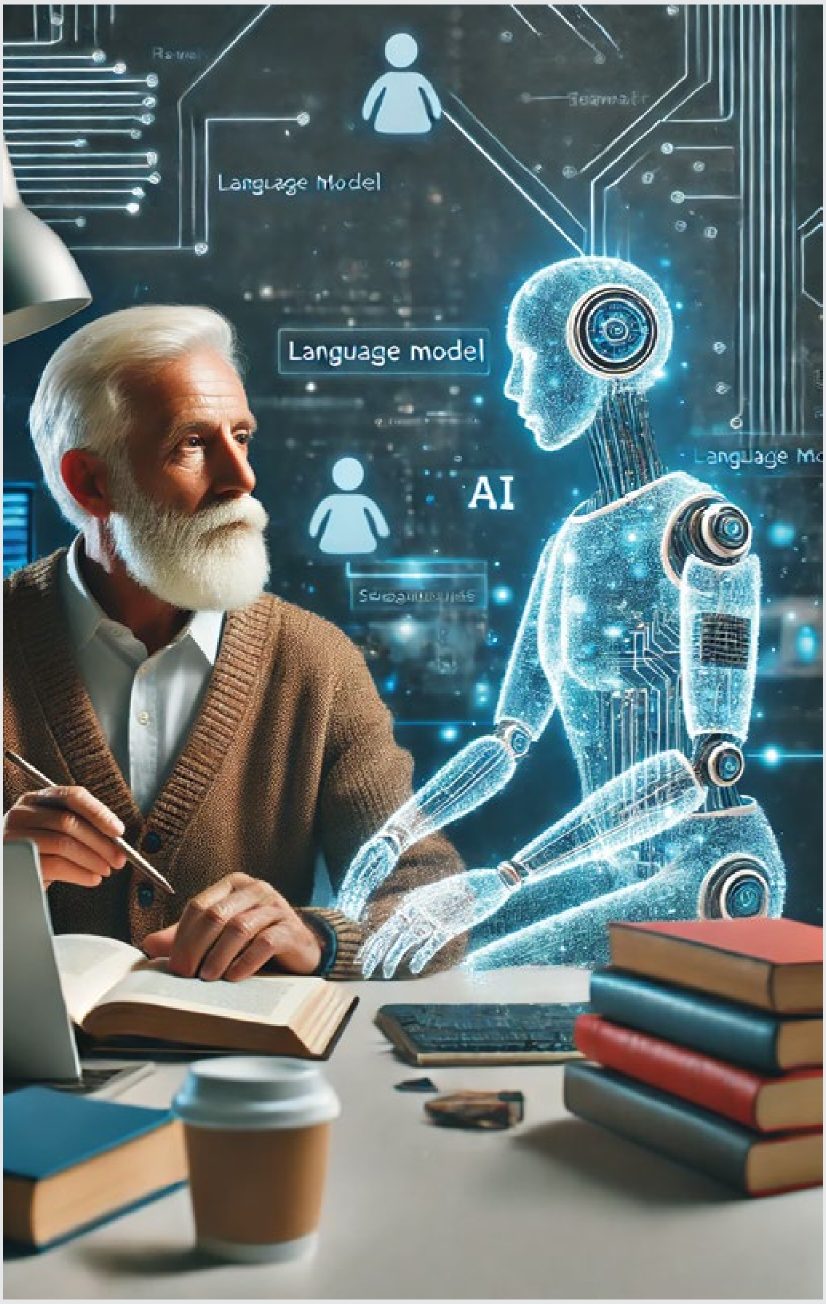

AI tools have already become a mainstay of the scientific and engineering workflow, from automated code generation to data analysis. This raises the question: Can AI independently author a technical paper worthy of publication?

Background

The recent proliferation of Large Language Models (LLM) in modern technical fields seems as sudden as it is pervasive. Yet, for those closely following AI research, this growth has been a long time coming. AI tools have already become a mainstay of the scientific and engineering workflow, from automated code generation to data analysis. This raises the question: Can AI independently author a technical paper worthy of publication?

My experiment took shape as a deliberate challenge to current AI capabilities. This wasn't to be a collaborative effort between human and machine – I wanted to test AI's abilities in their purest form. While I would provide initial prompts and manage the peer review process, the content itself should come entirely from AI systems, including technical analyses, citations, and images.

In my preliminary research, I chose seven different platforms as candidate authors. I specifically asked each AI system about the existence of AI-authored papers submitted to peer-reviewed publications. This meta-analysis proved interesting – while they could reference various experiments with AI writing, none could point to a verified instance of a fully AI-generated article submitted to peer review.

However, I knew of one such paper that did pass peer review for a scientific journal: "Cellular functions of spermatogonial stem cells in relation to JAK/STAT signaling pathway" [2] appears as if fabricated from whole cloth by AI. This exception proved particularly instructive. The Guo et al. paper contained extensive anatomical and factual errors and has since been retracted. This case highlights both the improving capabilities of AI in mimicking scholarly writing and the current limitations that eventually expose such attempts.

The failure of the Guo paper provided insights for my experiment. The paper's retraction stemmed not from its prose, which was passable, but from technical faults. This observation shaped my experimental design, particularly regarding the handling of technical figures and data visualisation.

Purpose and premise

Artificial Intelligence is reshaping technical communication; but its ability to independently produce scholarly works is untested. If not counting the Guo paper – which cannot be taken seriously – then the paper by Perplexity AI published in this edition of *Switchgear Magazine* [1] is the first legitimate submission of an article authored by AI.

The distinction between AI as a writing assistant and AI as an autonomous author is crucial. While many researchers use AI tools to enhance their writing processes, my experiment eliminated human intervention in the content creation phase. This limitation would allow me to clearly identify the current boundaries of AI capability in technical writing.

AI platforms

My selection of AI platforms aimed to represent the current state of the art in language models and content generation. Each platform brought unique strengths:

- Open AI's ChatGPT for its broad knowledge base, and my custom built GPT for switchgear
- DALL-E for visual content generation
- Perplexity AI for its ability to synthesize information
- Claude for its careful reasoning
- DeepSeek (via Hugging Face) for its technical capabilities
- Gab.ai's Nikola Tesla character for domain-specific knowledge in electrical engineering
- Venice.ai was included to test emerging platforms with novel approaches to content generation (no registration required, prompts and history not logged)

Prompt engineering

To obtain fully AI-generated content, I crafted prompts that would elicit comprehensive content without inserting human expertise into the process. I provided context for the AI to understand the requirements of a technical publication.

This basic definition belies the complexity I encountered. While each AI could engage in conversation and generate text, their ability to maintain technical accuracy and consistency across a long-form article proved to be a more challenging test.

Each "AI Assistant" (AIA) required individual consideration in my experimental design. The paid versions provided access to more advanced capabilities and larger context windows, crucial for generating long-form technical content.

Step 0: Topic selection

My first test of the AIAs' capabilities was for topic selection. Rather than imposing a topic, I asked each system to suggest one relevant to power systems and switchgear. Interestingly, there was significant convergence around the concept of Autonomous

Energy Systems (AES). This consensus suggested that the AIAs had successfully identified a trending topic in the field, though I would later discover limitations in their understanding of its practical implementation.

Step 1: Primary prompt

Next, I queried ChatGPT for suggested prompts. My initial prompts specified the basic requirements: technical depth, data-driven arguments, industry-relevant examples, and proper citations.

After several iterations, I settled on a prompt to write a full-length article written specifically to pass peer review. I ended with the statement: "It must sound like it was generated by a human". The final prompt was 640 words long.

Step 2: First Pass: Multiple AI assistants

After settling on the prompt, I submitted it to each AIA on my list. In an interesting meta-analysis step, I had each AIA evaluate the others' outputs. This cross-evaluation revealed consistent patterns in both strengths and weaknesses. The AIAs proved adept at identifying technical inconsistencies in each other's work, though they sometimes struggled to provide specific, actionable improvements.

I also commanded a ranking of all articles on a range of criteria, topmost being (1) appearing human-written, and (2) ability to pass peer review.

Step 3: Processing to finished content

In analysing the outputs, some candidate articles were clearly machine generated. Others were too short. Articles by Perplexity and Claude demonstrated superior technical coherence and citation structure. Perplexity's writing style struck the best balance between technical precision and readability, avoiding both overly casual language and excessive jargon.

I commanded ChatGPT to take the roll of Peer Reviewer and criticise Perplexity's article. Next, I fed Perplexity's article to Claude to re-write. Later, I fed the results back to Perplexity, which then took the role of peer reviewer. Finally, I prompted Perplexity to edit the article in response to its own critique.

DALL-E produced visually appealing graphics that, at first glance, appeared professional and plausible. However, these figures embodied a fundamental problem

The back-and-forth revision process between AIAs revealed an interesting dynamic. Each system could identify and attempt to correct issues in the other's work, leading to incremental improvements in technical clarity and argument structure. However, this process also highlighted a consistent blind spot: none identified inaccuracy of citations or data sources.

This discovery marked a critical turning point in my experiment. Upon checking the references generated by the AIAs, I found widespread fabrication and misrepresentation. This wasn't simply a matter of formatting errors or minor inaccuracies – the systems had generated entirely fictional citations and attributed false claims to real sources.

I knew that if reviewers found references to be incorrect or non-existent, they would reject the paper. Therefore, I abandoned my goal of submitting a paper entirely written by AI; I manually sourced references and quoted data.

Step 4: Images

Image creation presented another challenge. Perplexity cannot generate images, so I chose DALL-E. While the latter could generate professional-looking figures, the underlying data was not sourced or was entirely hallucinated. I made the deliberate choice to leave these AI-generated figures unchanged in the submission, viewing their reception as an important test of AIA's performance in peer review.

I requested visualizations of complex, forward-looking trends in the power systems industry, specifically:

- Cybersecurity
- Predicted evolution of grid operator skills
- Projected adoption of AES

DALL-E produced visually appealing graphics that, at first glance, appeared professional and plausible. However, these figures embodied a fundamental problem: they represented data that didn't exist, presenting imagined trends with unwarranted precision.

Peer review

The decision to maintain anonymity about the AI authorship was consistent with Switchgear Magazin's policy of double-blind peer reviews and crucial to my experimental design. I wanted to see if reviewers would identify AI-generated content through standard review processes, without being primed to look for it.

The initial peer review feedback revealed interesting patterns in both AI capabilities and limitations. The AIAs could respond to specific technical feedback by elaborating on requested topics. Objections were raised by reviewers, and the AI engine would satisfy many of them by copious responses to individual comments and addendums to the paper, which were then accepted.

Rejected

Despite the AI's ability to successfully respond to criticism with additional content and revisions, the article ultimately failed to meet publication standards.

A question from one reviewer struck at the heart of the experiment. With-

Was it written by a human who has some authority and experience on the topic or was it AI generated?

out knowing the article's AI origin, the reviewer had detected something artificial in its voice and depth of understanding.

> *"Was it written by a human who has some authority and experience on the topic or was it AI generated?"*

This intuitive recognition of non-human authorship stemmed from subtle patterns in how the AI handled industry-specific knowledge and practical implementation details.

> *"It futuristically speaks of 2035, however our industry is much more conservative in adopting new technologies than what this article suggests … I wonder if the author is even from our industry."*

This criticism revealed a crucial blind spot in AI's understanding of industry dynamics. While the AI could process and regurgitate technical information, it lacked the practical industry experience to recognize the conservative nature of power systems engineering.

The reviewers' attention to data sourcing exposed another critical weakness in AI-generated content. One stated:

> *Graphs "make no sense. What are they based on? The source of data should be referenced."*

The visually appealing figures generated by DALL-E couldn't withstand professional scrutiny. The lack of data sources and questionable projections raised red flags for reviewers accustomed to empirically grounded technical publications.

Discussion

The experiment revealed several fundamental limitations in AI's capacity to produce publication-worthy technical content.

1. Industry Wisdom: While AI can process and remix technical information, it cannot replicate the deep understanding that comes from years of industry experience. The unrealistic adoption timelines and focus on futuristic technologies revealed this limitation.
2. Data Integrity: The AI's lack of verifiable sources for its claims and visualizations represents a fundamental barrier to scholarly publication. In technical fields, unsupported assertions, no matter how plausible, cannot survive peer review.
3. Authentication of Expertise: The reviewers' ability to detect something "off" about the article's voice and perspective suggests that human expertise leaves distinctive markers that current AI systems cannot fully replicate.

These findings suggest that while AI might be useful for generating initial drafts or exploring ideas, it cannot yet independently produce works that meet the standards of peer-reviewed technical publications. The rejection demonstrated that traditional peer review processes remain effective at identifying content that lacks the hallmarks of genuine human expertise.

Conclusion

This experiment demonstrates that while AI has made remarkable progress in generating coherent technical content, it cannot yet independently produce works that meet scholarly publication standards.

"The question of whether a computer can think is no more interesting than the question of whether a submarine can swim." -attributed to Edsger Dijkstra [3]

The challenge ahead lies not in determining whether AI can replicate human technical writing, but in understanding how to effectively integrate AI capabilities with human expertise. As Dijkstra's submarine

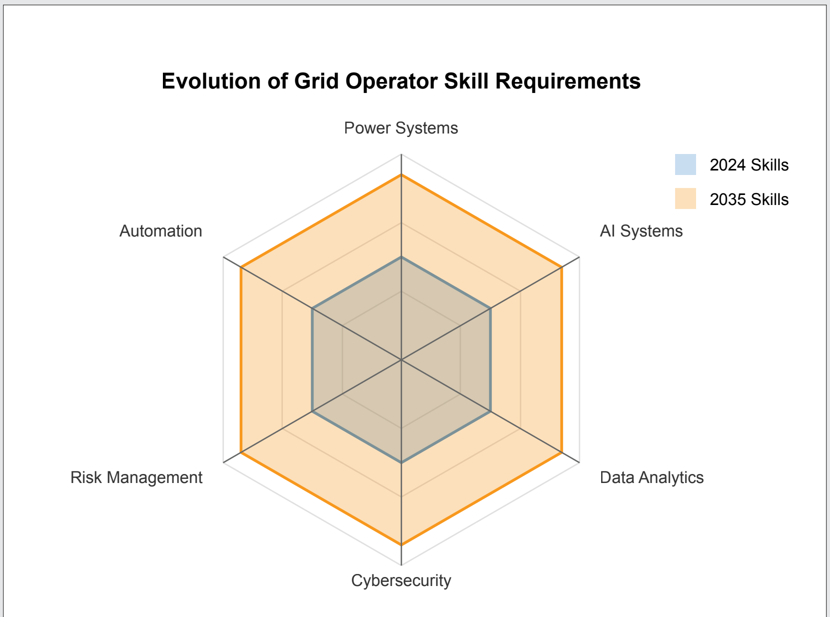

Figure: Unsourced projections of Autonomous Energy Systems adoption (AI-generated) [1].

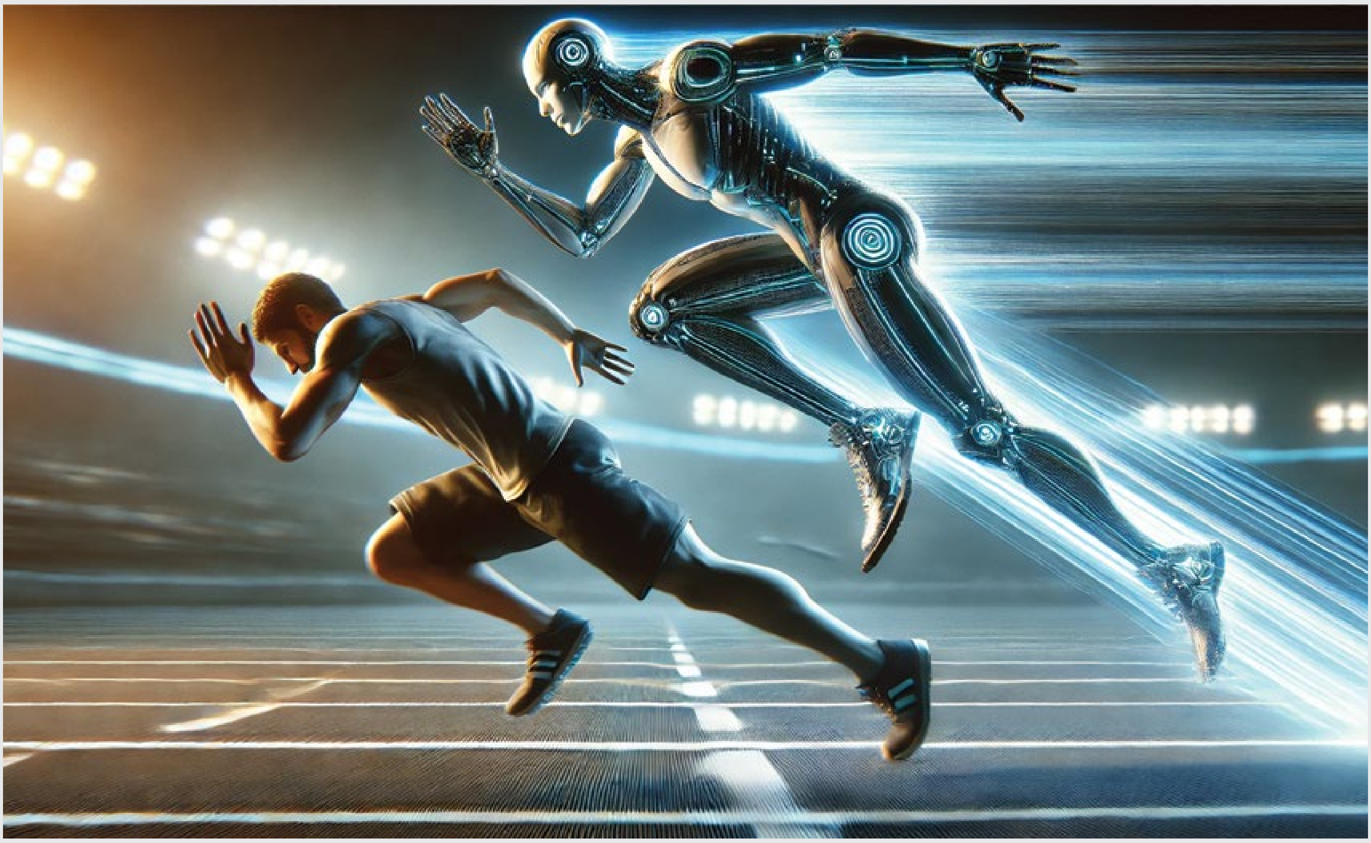

This experiment, while highlighting current limitations, also points toward a future where AI and human expertise complement each other in technical writing, creating more efficient and effective ways to communicate complex technical information

analogy suggests, the question isn't whether AI can think like a human author, but how it can best serve as a tool for human technical communication.

The immediate future of technical publishing will likely depend on a balanced approach that:

- Leverages AI's strengths in content generation and organization
- Maintains human oversight for accuracy and validity
- Establishes clear guidelines for AI use and disclosure
- Preserves the essential role of human expertise in technical communication

This experiment, while highlighting current limitations, also points toward a future where AI and human expertise complement each other in technical writing, creating more efficient and effective ways to communicate complex technical information.

Through this experiment, I've come to believe that the more profound question isn't about AI's ability to mimic human technical writing, but about how our approach to technical communication might evolve in response to AI capabilities. The challenge lies not in making AI more human-like, but in developing effective ways to combine AI's computational strengths with human expertise and judgment.

References

[1] Perplexity AI, "Grid operators in the age of autonomous energy systems: how smart grids will transform daily operations," Switchgear Magazine, vol. 2, no. 1, Mar. 2025.

[2] X. Guo, L. Dong, and D. Hao, "RETRACTED: Cellular functions of spermatogonial stem cells in relation to JAK/STAT signaling pathway," Frontiers in Cell and Developmental Biology, vol. 11, Feb. 2024, Art. no. 1339390, doi: 10.3389/fcell.2023.1339390.

[3] E. Dijkstra, "The question of whether a computer can think is no more interesting than the question of whether a submarine can swim," Quotations Page. [Online]. Available: http://www.quotationspage.com/quote/40304.html

Author

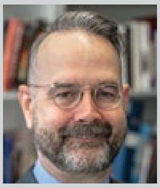

Douglas Maly is Editor-in-Chief of Switchgear Magazine.