REVIEW Open access

Generative AI and Foundation Models for Energy Systems: A Critical Review of Emerging Applications, Governance Risks, and Conceptual Research Opportunities

William Taylor[1*], Emily Davis[1], Thomas Wilson[2]

Abstract

Generative AI and foundation models are rapidly entering energy systems, offering novel capabilities for data generation, interaction, and analysis. Their transformative potential is matched by significant uncertainty about applications, risks, and research direction. The literature lacks a systematic synthesis of how these models are being used, what governance risks they introduce, and how the research community can conceptually structure this emerging field. This gap is especially important because energy systems are critical infrastructures where errors, opacity, and automation failures may have public consequences. This article provides a systematic review of emerging applications and governance risks of generative AI and foundation models in energy systems, and develops a novel conceptual framework to identify and organise research opportunities. The review connects application domains with governance challenges and future conceptual inquiry. A systematic review methodology following PRISMA principles was applied to 36 peer-reviewed articles from 2017–2025. A structured search, screening, and data extraction process was employed. Findings were synthesised into a conceptual framework. The review identifies four clusters of emerging applications: operational, design, scientific, and consumer-facing applications. It also catalogues five categories of governance risks and maps conceptual research opportunities across technical, ethical, and institutional dimensions. Two tables present the PRISMA flow and the conceptual framework. Generative AI represents a pivotal opportunity for energy systems research, but its trajectory must be guided by robust conceptual frameworks and proactive risk governance. This article provides both a systematic state-of-the-art and a research roadmap.

Keywords Systematic review, Conceptual framework, Foundation models, Generative AI, Governance risks, Energy systems

*Correspondence:
William Taylor
william.taylor@auckland.ac.nz
[1] Department of Artificial Intelligence for Grid Resilience and Geothermal Energy, Faculty of Engineering, University of Auckland, Auckland, New Zealand
[2] Department of Intelligent Energy Infrastructure and Data Science, Faculty of Engineering, University of Canterbury, Christchurch, New Zealand

Introduction

Generative AI and foundation models have begun to reshape energy systems research by introducing flexible capabilities for language reasoning, synthetic data generation, scenario construction, design support, and human–machine interaction. Lin *et al*. describe generative artificial intelligence as a new paradigm for research and education in smart energy systems, while Majumder *et al.* show that large language models already possess relevant but uneven capabilities

This is an open access article distributed under the terms of the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License (CC BY-NC-SA 4.0), which permits non-commercial use, distribution, reproduction, and adaptation in any medium or format, provided that appropriate credit is given to the original author(s) and source, a link to the license is provided, and any changes made are indicated. Adapted material must be distributed under the same license.

© 2026 The Author(s). | Published by Institute of Smart Energy Systems Press

for electric-sector reasoning [1, 2]. The field has therefore moved beyond generic enthusiasm toward early exploration of how models might assist dispatch, forecasting, policy analysis, building simulation, and operational documentation [3, 4]. However, this literature remains emergent, unevenly distributed across application areas, and insufficiently connected to the governance needs of critical infrastructure.

A structured review is needed because generative AI in energy is not simply an extension of conventional predictive analytics. Unlike task-specific machine learning, foundation models can generate text, code, scenarios, explanations, and synthetic data across multiple contexts, creating opportunities for power dispatch, building energy modelling, wind forecasting, and renewable scenario generation [4-7]. At the same time, general AI governance research has warned that foundation models raise distinctive risks related to hallucination, security, privacy, opacity, bias, and accountability [8-10]. A systematic synthesis is therefore necessary to connect the technical energy literature with the broader governance literature on high-stakes generative AI.

This article makes a dual contribution by combining a systematic review with conceptual framework development. The systematic component identifies and organises peer-reviewed studies on generative AI, foundation models, synthetic data generation, diffusion models, and governance risks relevant to energy systems from 2017 to 2025 [11, 12]. The conceptual component then translates this evidence into a framework linking application domains, model capabilities, governance vulnerabilities, and research opportunities. In doing so, the article treats generative AI for energy systems as an emerging socio-technical field rather than a narrow set of computational tools.

Research Questions

This review is guided by three research questions. First, what emerging applications of generative AI and foundation models have been reported or conceptually proposed for energy systems, including grid operation, forecasting, building energy management, policy analysis, and scenario generation [3, 13, 14]? Second, what governance risks are associated with these applications, particularly hallucination, privacy leakage, security vulnerability, opacity, liability diffusion, and sustainability cost [9, 10, 15]? Third, how can these applications and risks be synthesised into a conceptual framework that identifies research opportunities for responsible innovation in energy systems?

The rationale for the first question is that the field is expanding across operational and design contexts but lacks a consolidated map of application domains. Studies on advanced power dispatch, building energy modelling, HVAC fault diagnosis, wind speed forecasting, and automated energy data mining demonstrate that generative and foundation models are being tested across heterogeneous technical environments [4, 6, 16, 17]. The rationale for the second question is that energy systems involve public reliability, consumer protection, infrastructure security, and environmental responsibility, making general risks of generative AI especially consequential [18, 19]. The rationale for the third question is that conceptual fragmentation can slow responsible development unless applications and risks are integrated into a research agenda.

The questions also determine the methodological design of the review. The search strategy was designed to capture both energy-specific generative AI applications and broader foundation-model governance literature, because sector-specific studies rarely address the full range of ethical, legal, and safety issues [8, 20]. The screening criteria prioritised peer-reviewed journal articles from 2017–2025 to ensure that both earlier generative scenario modelling and recent large language model studies were included [7, 21]. Data extraction was then organised around application domain, generative capability, energy-system function, risk type, mitigation approach, and conceptual opportunity.

Systematic Review Method

Search strategy

© 2026 The Author(s). | Published by Institute of Smart Energy Systems Press

The review applied a structured search strategy across major academic databases, including Scopus, Web of Science, IEEE Xplore, ScienceDirect, ACM Digital Library, SpringerLink, and publisher databases relevant to energy and AI. Search strings combined terms such as “generative AI,” “foundation model,” “large language model,” “energy system,” “grid,” “utility,” “synthetic data,” “diffusion model,” “governance risk,” “hallucination,” and “critical infrastructure,” reflecting both the application literature and the governance literature [1, 10, 11]. The search covered publications from 2017 to 2025 because renewable scenario generation using generative adversarial networks predates the recent foundation-model wave [7]. Searches were complemented by backward and forward citation checks to capture adjacent studies on building energy applications, synthetic load profiles, and AI ethics frameworks [12, 18, 22].

Study selection criteria

Studies were included if they were peer-reviewed journal articles published between 2017 and 2025 and addressed generative AI, foundation models, large language models, diffusion models, generative adversarial networks, variational generative models, or governance risks directly relevant to energy systems or high-stakes AI. Energy-focused studies were eligible when they examined applications such as power dispatch, wind forecasting, building energy management, HVAC fault diagnosis, renewable scenario generation, load synthesis, transport energy demand forecasting, or EV charging scenarios [4, 6, 17, 23]. Governance studies were included when they addressed hallucination, privacy, security, AI ethics, foundation-model risk, climate impact, or responsible implementation in ways that could inform energy-system deployment [9, 10, 15]. Books, standards, reports, theses, websites, preprints without peer-reviewed journal publication, and conference papers were excluded to maintain a journal-based review corpus.

Screening process and PRISMA flow

The screening process followed PRISMA principles by moving from database identification to duplicate removal, title and abstract screening, full-text eligibility assessment, and final inclusion. An initial pool of 214 records was identified, and 47 duplicates were removed, leaving 167 records for title and abstract screening. After excluding records outside the scope of generative AI, foundation models, energy systems, or governance risk, 72 full-text articles were assessed for eligibility. The final corpus comprised 36 studies that collectively covered energy applications, synthetic data generation, foundation-model opportunities, and governance risks [1, 2, 11, 24].

Table 1 summarizes the PRISMA-based study selection process, outlining the flow from 214 identified records through duplicate removal, screening, eligibility assessment, and the final inclusion of 36 studies in the systematic review.

Table 1. PRISMA Flow of Study Selection for Systematic Review on Generative AI, Foundation Models, and Energy Systems Research

PRISMA stage	Records/articles retained or excluded	Rationale
Records identified from database searching	214	Searches covered generative AI, foundation models, energy systems, synthetic data, diffusion models, LLMs, and governance risks
Duplicates removed	47	Duplicate records across Scopus, Web of Science, IEEE Xplore, ScienceDirect, ACM Digital Library, and SpringerLink were removed
Records after duplicates removed	167	Unique records entered title and abstract screening
Records excluded during title and abstract screening	95	Excluded for being unrelated to energy, not focused on generative/foundation models, not peer-reviewed journal articles, or not within 2017–2025

© 2026 The Author(s). | Published by Institute of Smart Energy Systems Press

Full-text articles assessed for eligibility	72	Full texts were reviewed against application, governance, and methodological relevance criteria
Full-text articles excluded	36	Excluded because of insufficient relevance to energy systems or governance, conference-only status, missing DOI, non-peer-reviewed status, or purely empirical scope without conceptual relevance
Studies included in systematic review	36	Final corpus used for synthesis and conceptual framework development

Data extraction and coding

Data extraction recorded bibliographic details, energy-system domain, model type, generative capability, task function, governance concern, and conceptual contribution. Studies on building energy modelling, automated data mining, HVAC diagnosis, and load prediction were coded as design and operational applications, while renewable scenario generation, load-profile synthesis, and EV charging scenario generation were coded as synthetic data and scenario applications [13, 14, 23, 25]. Studies on hallucination, privacy, security, ethics, foundation models, and carbon footprint were coded as governance and cross-sector risk sources rather than direct energy applications [9, 10, 15, 26]. The coding process was interpretive but systematic, allowing the review to identify recurring application clusters and translate them into conceptual research opportunities.

Overview of Included Studies

Bibliometric characteristics

The included studies span 2017–2025, with a visible shift from generative adversarial approaches for renewable scenarios and energy data synthesis toward large language models and foundation models after 2023. Earlier studies focused on model-free renewable scenario generation and privacy-preserving synthetic energy data, while recent work examines LLMs for building energy modelling, power dispatch, energy policy research, and smart energy education [3, 4, 7, 21]. The corpus is concentrated in journals such as Energy and AI, Applied Energy, Energy and Buildings, Energy, IEEE Transactions on Smart Grid, IEEE Transactions on Power Systems, and broader AI governance venues [1, 6, 19, 23]. This distribution suggests that the field is developing at the intersection of energy informatics, power systems, building science, AI safety, and responsible innovation.

Thematic distribution

The thematic distribution of the corpus can be organised into three broad clusters. The first cluster focuses on energy-specific generative applications, including power dispatch, wind forecasting, building energy management, HVAC fault diagnosis, automated building model development, and energy load prediction [4, 14, 17, 27]. The second cluster focuses on synthetic data and scenario generation, including renewable generation scenarios, residential load profiles, consumer-level load profiles, EV charging scenarios, and diffusion-based demand forecasting [7, 28-30]. The third cluster addresses governance and conceptual foundations, drawing on studies of foundation-model risks, hallucination, AI ethics principles, privacy and security, climate impact, and cross-domain responsible AI frameworks [9, 15, 18, 24].

Emerging Applications

Operational and grid management applications

Operational applications are emerging around decision support, forecasting, fault diagnosis, and dispatch assistance. Cheng *et al*. examine a large language model for advanced power dispatch, suggesting that foundation models may support interpretation, planning, and operator-facing reasoning in grid control contexts [4]. Wu and Ling extend pre-

© 2026 The Author(s). | Published by Institute of Smart Energy Systems Press

trained language-model logic to spatio-temporal wind speed forecasting, while Zhang *et al*. show that domain-specific LLM fine-tuning may support HVAC fault diagnosis in building systems [6, 17]. These studies indicate that generative AI may contribute not only to prediction but also to explanation, workflow assistance, and operational communication.

Design, planning, and scientific discovery

Design and planning applications are especially visible in building energy research, where LLMs have been explored for automated data mining, energy conservation analysis, load prediction, model development, and debugging. Zhang *et al.* show that GPT-based workflows can support building energy management and conservation-oriented data mining, while Zhang, Ford, and Chen demonstrate how LLMs may assist building energy model development through exploratory case studies [5, 13, 16]. Agentic workflows for automatic building energy model development and debugging further indicate that foundation models may become design collaborators rather than merely analysis tools [27]. In scientific and engineering contexts, broader foundation-model research suggests that such systems could eventually support cross-domain energy materials discovery, simulation preparation, and knowledge synthesis, although this remains less mature in the energy-specific literature [24, 31].

Consumer-facing and market applications

Consumer-facing and market applications are less developed than operational and design applications, but the included studies point to plausible pathways through conversational AI, synthetic household data, smart-home optimisation, and policy analysis. Buster *et al.* demonstrate that LLMs can support energy policy research through wind energy siting ordinance analysis, which suggests future use in utility communication, regulatory interpretation, and local energy planning [3]. Razghandi *et al.* combine synthetic dataset generation with smart-home energy management, while studies of residential load profiles and consumer-level load generation show how generative methods can model customer demand patterns without relying exclusively on raw household data [28, 29, 32]. These applications create opportunities for customer service, demand-response communication, tariff explanation, and privacy-preserving market analysis, but they also require careful governance because consumer-facing systems can amplify misinformation, bias, and trust failures [10, 20].

Governance Risks

Risk categories and their manifestations

The reviewed literature indicates that governance risks can be organised into five interrelated categories: hallucination and factual unreliability, data sensitivity and privacy leakage, opacity and explainability deficits, security manipulation, and accountability diffusion. Hallucination is especially relevant because LLMs may produce plausible but false operational explanations, policy interpretations, or diagnostic narratives, and general studies of natural language generation show that hallucination is a persistent structural weakness rather than a minor error mode [9]. Security and privacy risks are equally central, since foundation models can expose sensitive information, be manipulated through prompts, or generate unsafe code-like outputs in high-stakes settings [10]. These risks are intensified in energy systems because generated content may influence infrastructure decisions, customer communication, planning assumptions, or regulatory interpretation.

Severity and specificity to energy systems

The severity of generative AI risk in energy systems arises from the sector's dependence on reliability, continuity, safety, and public trust. Foundation-model risks identified by Bommasani *et al*. are particularly consequential when translated into grid control, building operations, and utility-facing decision support, because model errors may affect not only information quality but also operational resilience [24]. The energy domain also introduces distinctive sustainability concerns, since large-scale model training and inference can produce environmental burdens that conflict with

© 2026 The Author(s). | Published by Institute of Smart Energy Systems Press

decarbonisation goals [15]. Nordgren's analysis of AI and climate change further suggests that responsible energy AI must evaluate not only what AI can optimise, but also whether its computational and institutional costs are justified [33].

Current mitigation measures and their limitations

Current mitigation approaches include domain-specific fine-tuning, human review, retrieval augmentation, structured workflows, privacy-preserving generation, and ethics-oriented governance principles. Domain-specific LLM work in HVAC fault diagnosis and building energy applications suggests that supervised adaptation may improve relevance, but it does not eliminate hallucination, opacity, or liability concerns [17, 22]. AI ethics tools and principles can support responsible implementation, yet Morley *et al*. show that translating principles into practice remains difficult without operational methods, accountability structures, and institutional incentives [34]. The current literature therefore provides useful safeguards, but it has not yet produced a mature governance model tailored to generative AI in energy infrastructure.

Conceptual Research Opportunities

Technical research gaps and opportunities

Technical research opportunities include domain-grounded evaluation, reliability benchmarking, uncertainty communication, secure prompting, and model architectures designed for energy-specific constraints. Diffusion-based approaches for load forecasting and transport energy demand show that generative models can support probabilistic scenario construction, but future research must clarify when generated scenarios are operationally meaningful rather than merely statistically plausible [30, 35]. Synthetic data studies in renewable energy, residential load profiles, and EV charging also create opportunities for privacy-preserving experimentation, yet they require stronger validation against physical constraints and grid realities [23, 25, 28]. A central technical agenda is therefore to develop evaluation methods that assess factuality, physical consistency, and operational usefulness together.

Governance and ethical research gaps

Governance research must move beyond broad AI ethics principles toward energy-specific accountability arrangements. Jobin *et al.* show that AI ethics guidelines often converge around values such as transparency, fairness, privacy, and accountability, but energy systems require those values to be translated into sectoral obligations involving reliability, consumer protection, and infrastructure security [18]. Mittelstadt argues that principles alone cannot guarantee ethical AI, which is especially relevant where utility operators, regulators, vendors, and model providers share responsibility for generated outputs [19]. Future governance research should therefore examine liability allocation, auditability, documentation standards, and procedural oversight for generative AI in energy organisations.

Integrative and transdisciplinary opportunities

Integrative research opportunities lie at the intersection of energy engineering, AI safety, law, ethics, organisational studies, and sustainability science. Kasirzadeh and Gabriel argue that aligning language models with human values requires attention to interaction, interpretation, and social context, which implies that energy-oriented generative AI should be studied as a socio-technical system rather than a stand-alone model [20]. Floridi and Chiriatti's analysis of GPT-3 highlights the broad scope and limits of large language models, suggesting that energy researchers should distinguish between useful assistance, unsupported automation, and misleading competence [8]. This creates a transdisciplinary agenda focused on how foundation models reshape expertise, organisational responsibility, user trust, and institutional decision-making in energy systems.

Conceptual Framework Development

Framework architecture and dimensions

© 2026 The Author(s). | Published by Institute of Smart Energy Systems Press

The conceptual framework developed from this review is structured around four linked dimensions: application domain, generative capability, governance risk, and research opportunity. Application domains include operational grid management, building and design support, synthetic scenario generation, consumer-facing interaction, policy analysis, and scientific discovery [1, 3, 5]. Generative capabilities include language reasoning, code assistance, data synthesis, diffusion-based forecasting, multi-modal design support, and explanation generation [6, 11, 35]. Governance risks and research opportunities are then mapped onto these capabilities to show where technical innovation must be paired with institutional safeguards.

Figure 1 visualises the proposed conceptual framework by showing how generative AI applications in energy systems are shaped by model capabilities, governance risks, and corresponding research opportunities.

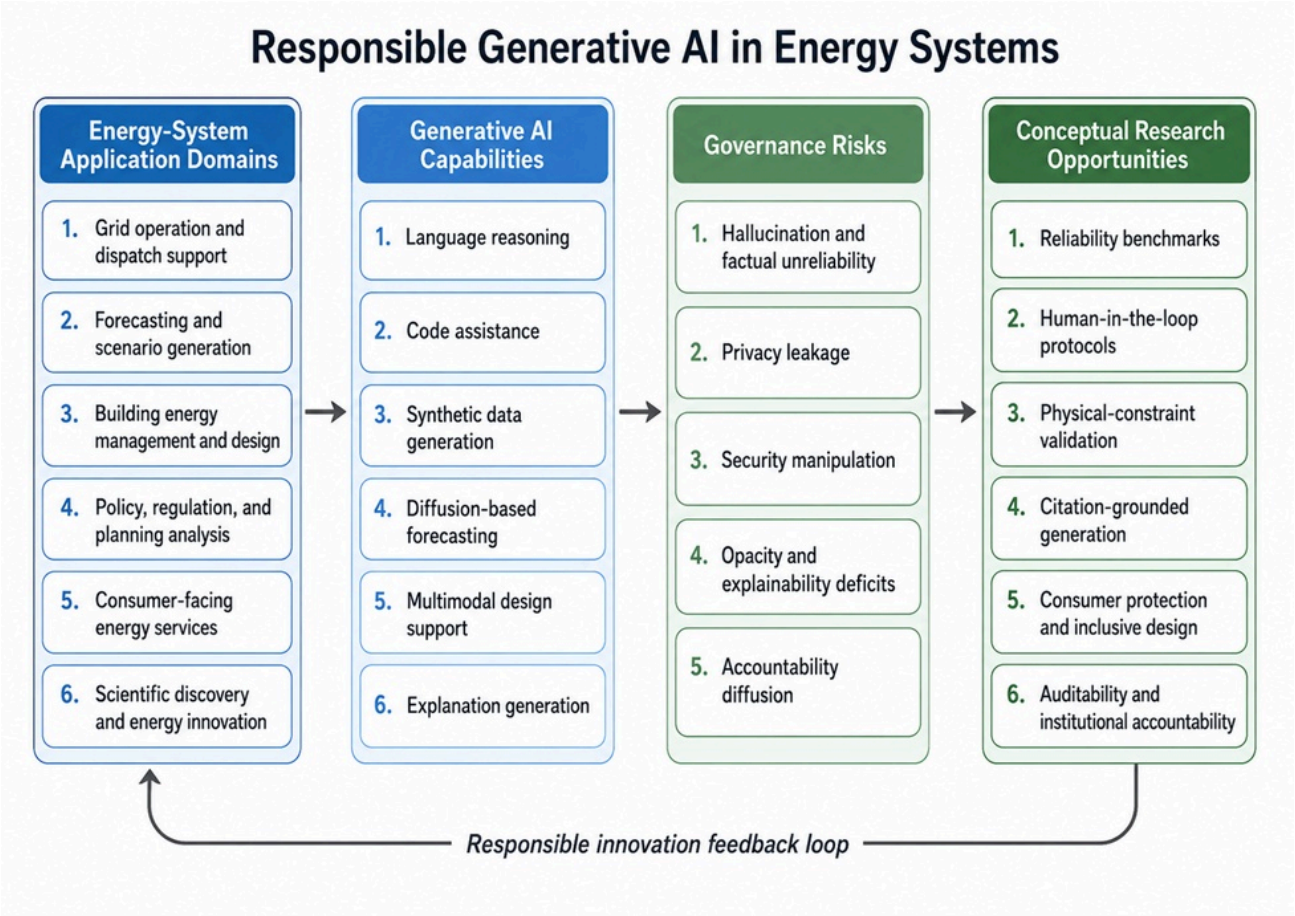

Figure 1. Responsible Generative AI in Energy Systems: Applications, Risks, and Research Opportunities

Mapping applications, risks, and opportunities

The framework treats applications, risks, and opportunities as co-produced rather than sequential. For example, LLM-assisted power dispatch may create opportunities for operator support, but it also raises risks of hallucinated recommendations, unclear accountability, and overreliance in high-pressure settings [4]. Synthetic load and renewable scenarios may support privacy-preserving planning, but they also raise risks of distributional distortion, model collapse, and false confidence in generated data [7, 21, 29]. This mapping approach helps clarify why responsible generative AI in energy requires simultaneous attention to capability, context, and consequence.

The framework as a research compass

© 2026 The Author(s). | Published by Institute of Smart Energy Systems Press

The framework functions as a research compass by directing scholars toward questions that connect technical feasibility with governance maturity. It does not claim that generative AI should be adopted wherever it is possible; rather, it asks whether each application has adequate evidence, safeguards, accountability, and public-value justification [26]. Floridi and Cowls' five-principle approach supports this orientation by linking beneficence, non-maleficence, autonomy, justice, and explicability to practical AI governance [26]. **Table 2** synthesises the framework by connecting major application domains with generative capabilities, governance risks, and conceptual research opportunities.

Table 2. Generative AI Applications in Energy Systems: Capabilities, Governance Risks, and Research Opportunities

Application Domain	Key Generative AI Capability	Governance Risk	Conceptual Research Opportunity
Grid operation and dispatch support	Language-based reasoning, operator assistance, code-like recommendation generation	Hallucinated instructions, automation bias, unclear liability during operational decisions	Develop energy-specific reliability benchmarks, human-in-the-loop protocols, and accountability models for generated operational advice
Forecasting and scenario generation	GAN-based, diffusion-based, and foundation-model-supported synthetic data generation	Physically unrealistic scenarios, privacy leakage, false confidence, model collapse	Build validation methods that combine statistical realism, physical constraints, privacy protection, and decision relevance
Building energy management and design	Automated data mining, model debugging, design support, diagnostic explanation	Erroneous diagnostics, opaque reasoning, unsafe optimisation recommendations	Create auditable LLM workflows for building simulation, fault diagnosis, and energy conservation planning
Policy, regulation, and planning analysis	Text generation, document analysis, ordinance interpretation, regulatory summarisation	Misinterpretation of legal or policy text, bias amplification, public misinformation	Establish governance protocols for verified retrieval, citation-grounded generation, and regulator-facing transparency
Consumer-facing energy services	Conversational explanation, smart-home support, tariff communication, demand-response interaction	Misleading advice, unequal access, privacy exposure, erosion of trust	Study inclusive conversational design, consumer protection requirements, and explainable utility communication
Scientific discovery and energy innovation	Cross-domain knowledge synthesis, hypothesis generation, multi-modal design support	Unsupported claims, opacity of generated hypotheses, sustainability cost of model use	Develop responsible discovery frameworks linking model outputs, expert validation, computational footprint, and reproducibility

Future Research Directions

Short-term research priorities

Short-term research should prioritise reliability evaluation, domain-grounded benchmarking, and deployment protocols for specific energy tasks. Large language model studies in power dispatch, energy policy, and building energy modelling provide early evidence of promise, but they also show the need for task-specific evaluation rather than generic demonstrations of capability [3-5]. Building energy surveys similarly indicate that opportunities are expanding faster than

© 2026 The Author(s). | Published by Institute of Smart Energy Systems Press

validated methods for safe operational integration [12, 22]. Immediate research should therefore focus on controlled feasibility studies, failure-mode analysis, and transparent documentation of model boundaries.

Long-term research horizons

Long-term research should examine how generative AI could reshape energy-system knowledge production, infrastructure governance, and institutional responsibility. The movement from task-specific models to AI large models for power systems suggests that future energy platforms may integrate forecasting, simulation, explanation, control support, and documentation within shared foundation-model environments [11]. Generalist foundation-model research in other domains indicates that such systems can consolidate multiple capabilities, but this consolidation may also create new dependencies and systemic vulnerabilities [31]. Energy research should therefore investigate not only model performance but also how foundation models reorganise expertise, accountability, and public oversight.

Methodological recommendations for future studies

Future studies should combine systematic review, case analysis, technical evaluation, governance assessment, and participatory inquiry. Energy-focused generative AI studies should report model assumptions, data provenance, validation procedures, human oversight mechanisms, and known limitations rather than presenting outputs as self-evident improvements [16, 27]. Governance studies should examine how ethical principles are operationalised through procurement, audit, documentation, incident reporting, and regulatory review [18, 34]. Methodologically, the field would benefit from mixed conceptual and practice-oriented designs that connect technical artefacts with organisational and societal consequences.

Limitations of the Review

Search and selection limitations

This review is limited by its reliance on peer-reviewed journal articles and by the fast-moving nature of generative AI research. Some relevant work on foundation models for energy systems may appear first in conference papers, technical reports, preprints, or industry demonstrations, but these sources were excluded to preserve the review's peer-reviewed evidence base [2, 11]. The 2017–2025 window captures both earlier generative scenario modelling and recent LLM applications, yet the field may evolve rapidly as new utility pilots and regulatory responses emerge [1, 3]. As a result, the review should be read as a systematic synthesis of published journal literature rather than an exhaustive account of all ongoing technical development.

Synthesis and framework subjectivity

The conceptual framework reflects an interpretive synthesis of heterogeneous studies rather than a statistical aggregation of comparable empirical findings. This is appropriate because the corpus spans GANs, diffusion models, LLMs, energy policy analysis, building energy modelling, ethics guidelines, hallucination research, and foundation-model governance [7, 9, 24, 35]. However, the mapping from applications to risks and opportunities necessarily involves judgement about how general AI governance insights apply to energy systems [10, 19]. Future reviews may refine this framework as more empirical deployments, regulatory cases, and sector-specific evaluation studies become available.

Conclusion

Generative AI and foundation models are becoming increasingly relevant to energy systems through applications in grid operations, building energy management, scenario generation, policy analysis, consumer interaction, and scientific discovery. The systematic review shows that the field is promising but still fragmented, with technical innovation advancing faster than governance capacity.

© 2026 The Author(s). | Published by Institute of Smart Energy Systems Press

The conceptual framework developed in this article responds to that fragmentation by linking application domains, generative capabilities, governance risks, and research opportunities. It encourages scholars to treat generative AI in energy as a socio-technical research field rather than a collection of isolated computational tools.

The central implication is that responsible innovation in generative AI for energy systems requires proactive conceptual work, rigorous evaluation, and governance design from the earliest stages of research. The research community should use the framework to guide future studies that are technically credible, ethically defensible, institutionally accountable, and aligned with public energy-system goals.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 15 Jun 2024 Revised: 20 Aug 2024 Accepted: 06 Sep 2024
Published online: 15 June 2026

Rights and permissions

Open Access The author(s) retain copyright. This article is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. It may be shared and adapted for non-commercial purposes with appropriate attribution, an indication of changes, and distribution of adaptations under the same license. Third-party material may be subject to separate terms identified in its credit line. View the license at https://creativecommons.org/licenses/by-nc-sa/4.0/.

References

1. Lin X, Luo Z, Du-Ikonen L, Lin X, Mao Y, Jiang H, et al. Generative artificial intelligence: pioneering a new paradigm for research and education in smart energy systems. Energy AI. 2025;22:100610. https://doi.org/10.1016/j.egyai.2025.100610.

2. Majumder S, Dong L, Doudi F, Cai Y, Tian C, Kalathil D, et al. Exploring the capabilities and limitations of large language models in the electric energy sector. Joule. 2024;8(6):1544-9. https://doi.org/10.1016/j.joule.2024.05.009.

3. Buster G, Pinchuk P, Barrons J, McKeever R, Levine A, Lopez A. Supporting energy policy research with large language models: a case study in wind energy siting ordinances. Energy AI. 2024;18:100431. https://doi.org/10.1016/j.egyai.2024.100431.

4. Cheng Y, Zhao H, Zhou X, Zhao J, Cao Y, Yang C, et al. A large language model for advanced power dispatch. Sci Rep. 2025;15(1):8925. https://doi.org/10.1038/s41598-025-91940-x.

5. Zhang L, Chen Z, Ford V. Advancing building energy modeling with large language models: exploration and case studies. Energy Build. 2024;323:114788.

6. Wu T, Ling Q. STELLM: spatio-temporal enhanced pre-trained large language model for wind speed forecasting. Appl Energy. 2024;375:124034. https://doi.org/10.1016/j.apenergy.2024.124034.

7. Chen Y, Wang Y, Kirschen D, Zhang B. Model-free renewable scenario generation using generative adversarial networks. IEEE Trans Power Syst. 2018;33(3):3265-75.

8. Floridi L, Chiriatti M. GPT-3: its nature, scope, limits, and consequences. Minds Mach. 2020;30(4):681-94.

9. Ji Z, Lee N, Frieske R, Yu T, Su D, Xu Y, et al. Survey of hallucination in natural language generation. ACM Comput Surv. 2023;55(12):1-38. https://doi.org/10.1145/3571730.

10. Yao Y, Duan J, Xu K, Cai Y, Sun Z, Zhang Y. A survey on large language model (LLM) security and privacy: the good, the bad, and the ugly. High-Confid Comput. 2024;4(2):100211. https://doi.org/10.1016/j.hcc.2024.100211.

11. Yao Q, Fang F, Chen Y, Liu J, Mo H, Ao Y. AI large models for power system: a survey and outlook. IET Smart Energy Syst. 2025;1(1):3-21. https://doi.org/10.1049/ses2.70000.

12. Arslan M, Munawar S. Large language models in building energy applications: a survey. Energy Build. 2026;352:116800. https://doi.org/10.1016/j.enbuild.2025.116800.

13. Zhang C, Lu J, Zhao Y. Generative pre-trained transformers (GPT)-based automated data mining for building energy management: advantages, limitations and the future. Energy Built Environ. 2024;5(1):143-69. https://doi.org/10.1016/j.enbenv.2023.06.005.

14. Zhang C, Zhang J, Zhao Y, Lu J. Automated data-driven building energy load prediction method based on generative pre-trained transformers (GPT). Energy. 2025;318:134824. https://doi.org/10.1016/j.energy.2025.134824.

15. Luccioni AS, Viguier S, Ligozat AL. Estimating the carbon footprint of BLOOM, a 176B parameter language model. J Mach Learn Res. 2023;24(253):1-15.

16. Zhang C, Zhang J, Zhao Y, Lu J. Automated data mining framework for building energy conservation aided by generative pre-trained transformers (GPT). Energy Build. 2024;305:113877. https://doi.org/10.1016/j.enbuild.2023.113877.

17. Zhang J, Zhang C, Lu J, Zhao Y. Domain-specific large language models for fault diagnosis of heating, ventilation, and air conditioning systems by labeled-data-supervised fine-tuning. Appl Energy. 2025;377:124378. https://doi.org/10.1016/j.apenergy.2024.124378.

© 2026 The Author(s). | Published by Institute of Smart Energy Systems Press

18. Jobin A, Ienca M, Vayena E. The global landscape of AI ethics guidelines. Nat Mach Intell. 2019;1(9):389-99.

19. Mittelstadt B. Principles alone cannot guarantee ethical AI. Nat Mach Intell. 2019;1(11):501-7.

20. Kasirzadeh A, Gabriel I. In conversation with artificial intelligence: aligning language models with human values. Philos Technol. 2023;36(2):27. https://doi.org/10.1007/s13347-023-00606-x.

21. Li Y, Li J, Wang Y. Privacy-preserving spatiotemporal scenario generation of renewable energies: a federated deep generative learning approach. IEEE Trans Ind Inform. 2022;18(4):2310-20. https://doi.org/10.1109/TII.2021.3098259.

22. Liu M, Zhang L, Chen J, Chen WA, Yang Z, Lo LJ, et al. Large language models for building energy applications: opportunities and challenges. Build Simul. 2025;18(2):225-34. https://doi.org/10.1007/s12273-025-1235-9.

23. Li S, Xiong H, Chen Y. DiffCharge: generating EV charging scenarios via a denoising diffusion model. IEEE Trans Smart Grid. 2024;15(4):3936-49. https://doi.org/10.1109/TSG.2024.3360874.

24. Bommasani R, Hudson DA, Adeli E, Altman R, Arora S, von Arx S, et al. On the opportunities and risks of foundation models. arXiv [Preprint]. 2021. arXiv:2108.07258.

25. Li J, Chen Z, Cheng L, Liu X. Energy data generation with Wasserstein deep convolutional generative adversarial networks. Energy. 2022;257:124694.

26. Floridi L, Cowls J. A unified framework of five principles for AI in society. In: Carta S, editor. Machine learning and the city: applications in architecture and urban design. Hoboken (NJ): Wiley; 2022. p.535-45.

27. Zhang L, Ford V, Chen Z, Chen J. Automatic building energy model development and debugging using large language models agentic workflow. Energy Build. 2025;327:115116. https://doi.org/10.1016/j.enbuild.2024.115116.

28. Claeys R, Cleenwerck R, Knockaert J, Desmet J. Capturing multiscale temporal dynamics in synthetic residential load profiles through generative adversarial networks (GANs). Appl Energy. 2024;360:122831. https://doi.org/10.1016/j.apenergy.2024.122831.

29. Xia W, Huang H, Salazar Duque EM, Hou S, Palensky P, Vergara PP. Comparative assessment of generative models for transformer- and consumer-level load profiles generation. Sustain Energy Grids Netw. 2024;38:101338. https://doi.org/10.1016/j.segan.2024.101338.

30. Jiang S, Tran CQ, Keyvan-Ekbatani M. A diffusion-model-based approach for forecasting energy demand in New Zealand's transport sector. Appl Energy. 2025;400:126617.

31. Moor M, Banerjee O, Abad ZS, Krumholz HM, Leskovec J, Topol EJ, et al. Foundation models for generalist medical artificial intelligence. Nature. 2023;616(7956):259-65. https://doi.org/10.1038/s41586-023-05881-4.

32. Razghandi M, Zhou H, Erol-Kantarci M, Turgut D. Smart home energy management: VAE-GAN synthetic dataset generator and Q-learning. IEEE Trans Smart Grid. 2024;15(2):1562-73. https://doi.org/10.1109/TSG.2023.3288824.

33. Nordgren A. Artificial intelligence and climate change: ethical issues. J Inf Commun Ethics Soc. 2023;21(1):1-5.

© 2026 The Author(s). | Published by Institute of Smart Energy Systems Press

34. Morley J, Floridi L, Kinsey L, Elhalal A. From what to how: an initial review of publicly available AI ethics tools, methods and research to translate principles into practices. In: Ethics, governance, and policies in artificial intelligence. Cham: Springer; 2021. p.153-83.

35. Zhang L, Jiang Z, Ji T, Chen Z. Diffusion-based inpainting approach for multifunctional short-term load forecasting. Appl Energy. 2025;377:124442. https://doi.org/10.1016/j.apenergy.2024.124442.

© 2026 The Author(s). | Published by Institute of Smart Energy Systems Press