

Generative AI in Manufacturing and Industrial Contexts: A Systematic Review of Applications, Challenges, and Future Directions

Galina Ilieva ^{1,*}  and Yuliy Iliev ²

¹ Department of Management and Quantitative Methods in Economics, University of Plovdiv Paisii Hilendarski, 4000 Plovdiv, Bulgaria

² Teletek Electronics, 1220 Sofia, Bulgaria; yuliy.iliev@teletek-electronics.bg

* Correspondence: galili@uni-plovdiv.bg

Abstract

Generative artificial intelligence (GAI) is expanding from model-centered research into engineering and manufacturing activities, but its scope and maturity remain uneven. This PRISMA-guided bibliometric and abstract-level thematic review maps peer-reviewed industrial GAI research published from 2022 to 4 June 2026. Searches of Scopus, Web of Science, and the ACM Digital Library identified 492 records; 119 duplicates and 121 ineligible records were removed, leaving 252 studies. Keyword normalization, co-occurrence analysis, dominant and secondary thematic coding, and an abstract-reported evidence characterization were applied. The corpus shows two connected trajectories: engineering generation based on generative models for design, topology, materials, and electronics, and knowledge-intensive industrial intelligence based on large language models, retrieval-augmented generation, knowledge graphs, agents, and human–AI collaboration. Most studies report empirical or computational evaluation (72.2%), but 84.5% remain research-stage; only 0.8% indicate operational industrial evidence in their abstracts. The findings, therefore, distinguish publication activity from deployment maturity. Priority requirements for adoption include domain-grounded data, verification, manufacturability checks, traceability, cybersecurity, intellectual property protection, system integration, workforce preparation, and human accountability. This review contributes a reproducible cross-domain map, an overlap-aware synthesis, and stakeholder-specific guidance for trustworthy industrial GAI.

Keywords: generative artificial intelligence; industrial AI; manufacturing; large language models; diffusion models; generative adversarial networks; engineering design; electronic design automation; smart manufacturing; bibliometric review



Academic Editor: Chunping Li

Received: 29 June 2026

Revised: 2 August 2026

Accepted: 4 August 2026

Published: 6 August 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

Manufacturing and industrial systems must increasingly combine productivity, flexibility, sustainability, and resilience. Conventional analytical AI supports classification, forecasting, detection, and optimization, whereas generative artificial intelligence (GAI) can create or transform technical artifacts, including text, code, geometry, process plans, material candidates, synthetic data, and decision support content. This capability is beginning to reshape engineering knowledge and work across industrial lifecycles [1,2]. In industrial applications, generated outputs must satisfy physical, economic, safety, and regulatory requirements in addition to being technically plausible. Their value, therefore,

depends on reliable verification, integration with established engineering tools, and use within accountable organizational processes.

Industrial GAI encompasses several learned model families and system architectures. Large language models (LLMs), retrieval-augmented generation (RAG), knowledge graphs, and agents support knowledge-intensive tasks, while generative adversarial networks (GANs), variational autoencoders (VAEs), and diffusion models generate data, structures, and candidate solutions. These technologies are increasingly combined with computer-aided design (CAD), simulation, digital twins, and industrial information systems [3–6]. Their roles are complementary rather than interchangeable: language-centered systems organize and transform technical knowledge, whereas deep generative models often explore high-dimensional design, material, image, signal, or process spaces. Hybrid architectures may connect these capabilities by grounding generation in domain repositories and submitting proposed outputs to simulation, optimization, or rule-based verification.

In this review, GAI refers to learned models that generate new content, representations, candidate solutions, or executable artifacts. Classical generative design based solely on rule-based search, topology optimization, evolutionary computation, or parametric optimization is excluded unless it incorporates a learned generative model. This distinction limits false-positive retrieval and keeps this review focused on contemporary GAI.

Applications span engineering and product development, materials discovery, additive manufacturing, process planning, electronics design automation (EDA), robotics, smart manufacturing, and industrial decision support. Learned generative models support manufacturable design and product ideation [5,7,8], mechanical engineering, CAD, digital twins, and process route generation [9–12], as well as EDA and human–AI co-design [13,14]. Knowledge-augmented automation, sustainable product development, platform personalization, and complex product design further extend this scope [15–18]. The generated artifacts range from component geometries and material candidates to circuit layouts, production instructions, explanations, and coordinated action plans. Across these settings, usefulness depends not only on generative performance but also on constraint satisfaction, verification, traceability, and integration with existing workflows. A plausible output may still be unmanufacturable, unsupported by source evidence, incompatible with production equipment, or unsafe under actual operating conditions.

GAI builds on the digital infrastructure associated with Industry 4.0 and the human-centered, resilient, and sustainable orientation of Industry 5.0, but it is not synonymous with either. Its industrial value lies in functions such as design space exploration, knowledge reuse, multimodal assistance, synthetic data generation, and decision support. These functions can shorten selected engineering activities and improve access to distributed technical knowledge, but they also introduce risks involving data quality, hallucination, intellectual property, cybersecurity, accountability, and human oversight [2,4,6,15]. Effective adoption, therefore, requires more than model selection. It also requires domain-grounded data, interoperable interfaces, documented approval points, monitoring procedures, and a clear allocation of responsibility between developers, engineers, managers, and end users.

Existing reviews provide valuable depth but generally emphasize a particular model family, application domain, or lifecycle stage. Manufacturing-wide syntheses examine applications and barriers [1,2], whereas other reviews focus on diffusion-based engineering [3], industrial LLMs [4,6], generative design [5], EDA [13], or knowledge-enhanced product development [18]. The literature, therefore, lacks a common cross-domain account of how these research streams relate and whether publication growth is accompanied by stronger validation, reproducibility, and operational adoption. This distinction is important because increasing research activity does not necessarily indicate technical maturity, reproducible

performance, or sustained industrial use. A cross-domain synthesis can reveal both technological convergence and the gap between research activity and deployment readiness.

This study addresses four research questions:

RQ1: How have the publication volume and normalized vocabulary of industrial GAI research developed?

RQ2: Which application domains and technological relationships structure the literature?

RQ3: What levels of validation, reproducibility reporting, and industrial adoption are reported in titles and abstracts?

RQ4: Which technical and organizational priorities emerge for researchers, managers, engineers, and policymakers?

This study makes four contributions. First, it provides a reproducible three-database corpus supported by a record-level screening audit. Second, it combines normalized keyword mapping with dominant and secondary thematic coding. Third, it characterizes the validation, reproducibility, and adoption signals reported in titles and abstracts, thereby distinguishing publication activity from deployment maturity. Fourth, it translates the synthesis into stakeholder-specific research and implementation priorities.

The remainder of this review is organized as follows. Section 2 describes the review design, information sources, eligibility criteria, coding procedure, keyword normalization, and analytical methods. Section 3 positions this study relative to representative reviews and clarifies its cross-domain contribution. Section 4 presents the corpus profile, keyword network, thematic structure, and abstract-reported evidence and adoption maturity. Section 5 interprets the convergence of engineering generation and industrial intelligence, discusses system integration challenges, presents stakeholder-oriented recommendations, and outlines this study's limitations. Section 6 summarizes the conclusions and identifies priorities for reproducible research and operational industrial validation.

2. Materials and Methods

This section describes the identification, screening, coding, normalization, and analysis of the review corpus. The methodological sequence maintains traceability between database retrieval, eligibility decisions, bibliometric and thematic analyses, and the abstract-reported evidence characterization. Particular attention is given to distinguishing learned GAI from classical generative design methods.

2.1. Review Design and Scope

This study follows PRISMA 2020 reporting principles [19] and is conducted as a PRISMA-guided bibliometric and abstract-level thematic review. It does not involve systematic full-text outcome extraction or a formal risk-of-bias assessment of every included publication. The primary unit of analysis is the bibliographic record, including its indexed title, abstract, author keywords, publication year, source, and document type.

Eligible records were English-language peer-reviewed articles, reviews, and conference papers published between 2022 and the search date of 4 June 2026. Each record had to combine a learned generative approach with an industrial, manufacturing, engineering, materials, electronics/EDA, production, or closely related artifact-development context. Biomedical structure and built environment studies were included only when their principal contributions concerned transferable engineering design, materials, fabrication, or infrastructure methods rather than clinical applications. General healthcare, pharmaceutical, educational, tourism, and non-industrial business applications were excluded.

In accordance with the definition established in the Introduction, studies described as generative design were eligible only when they explicitly applied a learned generative model. Rule-based, parametric, topology optimization, evolutionary computation, and

numerical search methods were not classified as GAI solely because they used the term “generative design.”

2.2. Information Sources and Search Strategy

Scopus, Web of Science Core Collection, and the ACM Digital Library were selected to provide complementary coverage of engineering journals, manufacturing and materials research, electronics and computing venues, and relevant conference proceedings. Database-specific search syntax and fields differed, while the queries retained a common conceptual structure combining learned generative technologies with industrial, manufacturing, engineering, and artifact development contexts. The final queries required an explicit learned-generative concept together with an industrial or engineering context. Broad standalone uses of “generative design” were removed because they retrieved classical optimization studies without a learned generative model. Database-specific syntax, search fields, limits, retrieved counts, and the execution date are reproduced in Supplementary Material Section S1 and the public dataset.

The search strategy combined terms describing learned generative technologies, including GAI, LLMs, foundation models, GANs, VAEs, diffusion models, RAG, and generative agents, with terms related to manufacturing, industry, production, engineering, products, materials, electronics, semiconductors, EDA, CAD, process planning, robotics, and factories. Search fields were adapted to the syntax of each database while preserving this conceptual structure.

Broad standalone uses of “generative design” were not used as sufficient evidence of eligibility because they retrieved numerous studies applying classical optimization without a learned generative model. Synthetic data generation was included when the data were produced by an eligible learned model and applied in an industrial or engineering context. Final inclusion decisions were made through author screening rather than an automated classifier.

2.3. Screening and Eligibility

The database searches intentionally used broad generative design terminology to maximize sensitivity. After deduplication, 373 unique records underwent title, abstract, and keyword screening. Ninety-two records were excluded because generative design referred only to rule-based, parametric, evolutionary, or conventional topology optimization methods without a learned generative model, and 29 records were outside the industrial or manufacturing scope. The resulting analytical corpus comprised 252 records.

Figure 1 presents the identification, deduplication, screening, exclusion, and inclusion decisions according to the PRISMA sequence.

As Figure 1 shows, deduplication reduced the initial set from 492 to 373 records, after which substantive screening removed approximately one-third of the unique records. The large number of exclusions related to classical generative design demonstrates the importance of applying a clear learned model eligibility criterion.

2.4. Coding and Keyword Normalization

Two authors screened the records and resolved ambiguous eligibility and dominant-topic decisions through discussion. A preliminary codebook was refined iteratively using titles, abstracts, author keywords, and source context. Each retained study was assigned one dominant topic for mutually exclusive counting and any applicable secondary tags for overlap analysis. Dominant categories support concise corpus and annual comparisons, whereas secondary tags preserve cross-domain relationships.




Figure 1. Simplified flow diagram based on PRISMA 2020 [19] for identification, screening, and inclusion of the strict industrial GAI corpus.

No formal independent double-coding exercise or inter-rater reliability coefficient was completed. Consensus-based resolution is not equivalent to independently quantified agreement, and this limitation is acknowledged in Section 5.4. The codebook, record-level screening decisions, dominant categories, and secondary tags are provided to support reproduction and alternative classification.

Author keywords were normalized using an instance-preserving thesaurus. The normalization rules covered capitalization, punctuation, whitespace, singular and plural forms, British and American spelling, unambiguous abbreviations, and clearly synonymous model names. Conceptually related but non-equivalent terms were retained separately. For example, “machine learning” was not merged with “deep learning,” and “generative design” was not merged with “topology optimization.” Ambiguous variants were also kept separate. Sensitivity to normalization was examined by comparing the normalized top-20 keyword list with a minimally cleaned raw list.

2.5. Bibliometric, Network, and Evidence Characterization

The bibliometric analysis reports annual publication counts and shares, year-over-year changes for complete years, document types, normalized keyword frequencies, and binary document-level keyword co-occurrence. The network comprises the 20 most frequent normalized author keywords. Node size represents the number of records containing a keyword, while edge weight represents the number containing both connected terms. To preserve readability, the weighted co-occurrence network of the 20 most frequent normalized author keywords displays the 42 strongest edges, while the complete set of 94 nonzero edges is provided in the supplementary analysis workbook.

Calculations were performed in Python 3.12.4 (Python Software Foundation, Wilmington, DE, USA) using pandas and openpyxl, and the figures were generated with Matplotlib 3.9.1. Temporal development is described using publication counts, shares, and year-over-year changes. Inferential trend regression was not applied because the dataset contains only four complete annual observations, while the 2026 value covers only the period up to 4 June.

A rule-based evidence characterization was applied to the indexed titles and abstracts. It records the study type, reported validation context, mention of publicly available code or data, and abstract-reported stage of industrial adoption. These variables characterize the evidence reported in the available metadata. They do not constitute a validated quality score, a formal full-text risk-of-bias assessment, or independent confirmation of operational effectiveness.

The methodology combines explicit eligibility criteria, record-level screening, transparent keyword normalization, mutually exclusive and overlapping thematic classifications, bibliometric and network analyses, and a bounded characterization of reported evidence and adoption maturity. This approach supports cross-domain mapping while preserving the inferential limitations of an abstract-level review. Section 3 uses this methodological basis to position the study relative to existing reviews.

3. Positioning Against Existing Reviews

A cross-domain synthesis should complement rather than duplicate the specialized reviews already available. This section, therefore, compares the coverage of representative review streams and identifies the additional analytical value provided by the present study.

Manufacturing-wide reviews summarize applications, cases, barriers, and future directions [1,2]. Diffusion-focused reviews examine generative modeling under engineering and physical constraints [3]. Product lifecycle and smart-manufacturing reviews emphasize LLM adaptation, RAG, knowledge integration, and industrial assistants [4,6]. EDA reviews analyze circuit and layout automation [13], while materials reviews address inverse design, discovery, and structure–property relationships [20–26]. These studies provide domain depth but use different boundaries, evidence bases, and taxonomies.

Table 1 summarizes this comparison. Its purpose is not to rank previous reviews, but to show that their main strengths are concentrated in particular technological or application layers.

Table 1. Positioning of the present review relative to representative review streams.

Review Emphasis	Primary Coverage	Boundary	Source
Manufacturing-wide GAI	Applications, cases, and challenges	Broad synthesis; limited cross-domain bibliometric integration	[1,2]
Diffusion in design/manufacturing	Diffusion models and physical design	Strong method focus; narrower model family	[3]
Industrial LLMs/smart manufacturing	LLM adaptation, RAG, and lifecycle support	Knowledge-centered; limited engineering generation coverage	[4,6]
EDA and complex product design	LLMs, agents, and knowledge graphs	Domain-specific electronics/product focus	[13,18]
Materials and inverse design	Molecules, alloys, structures, and properties	Deep materials focus; limited organizational synthesis	[20–26]
Present review	252-study cross-domain map	Overlap, evidence maturity, and stakeholder guidance	This study

Note: The comparison describes analytical emphasis and scope rather than methodological quality.

As Table 1 indicates, the present contribution lies in using one screened corpus and one analytical procedure to connect engineering generation, industrial knowledge systems, EDA, materials, process applications, and organizational readiness. The secondary-tag analysis and evidence characterization add two dimensions that are rarely combined in the earlier reviews: cross-domain technological overlap and the distinction between research visibility and reported deployment maturity.

This positioning establishes the gap addressed by the Results section. This review does not claim greater domain depth than each specialist synthesis; instead, it provides

a common systems-level map through which their otherwise separate findings can be interpreted together.

4. Results

The results are presented from general corpus structure to more interpretive analyses. First, annual output and document types describe the evidence base. Second, normalized keywords and their co-occurrences reveal the conceptual structure. Third, dominant and secondary codes show application domains and their overlap. Finally, the evidence characterization distinguishes reported evaluation from industrial adoption.

4.1. Corpus Profile and Temporal Development

The strict corpus comprises 252 studies: 201 articles, 28 proceedings papers, 17 reviews, four early-access articles, one article/proceedings paper, and one early-access review. Figure 2 combines annual publication output with the number of unique normalized author keywords, allowing growth in volume to be considered together with vocabulary diversification.




Figure 2. Annual publication output and unique normalized author-keyword counts. The asterisk (*) indicates incomplete annual data as of the search date.

Figure 2 shows modest growth between 2022 and 2024, followed by a pronounced increase in 2025. The parallel increase in unique keywords suggests thematic diversification, but it should not be interpreted as evidence that the field has reached industrial maturity. The 2026 observation is partial and cannot be compared directly with complete calendar years.

Table 2 provides the exact annual values behind Figure 2, including each year's share of the 252-study corpus and the year-over-year change for complete years.

Table 2. Annual output, relative frequency, and vocabulary breadth.

Year	Publications	Share (%)	YoY (%)	Keyword Instances	Unique Keywords
2022	21	8.3	—	96	76
2023	26	10.3	23.8	126	101
2024	29	11.5	11.5	159	125
2025	98	38.9	237.9	493	336
2026 *	78	31.0	—	369	269

Note: YoY = year-over-year change. The asterisk (*) indicates incomplete annual data as of the search date.

Table 2 confirms that 2025 accounts for 38.9% of the corpus and records a 237.9% increase over 2024. This jump is descriptively important, but the short observation window and database-indexing effects preclude a strong statistical trend claim.

4.2. Keyword Structure and Normalization Sensitivity

Among the 252 retained studies, 225 records contained author keywords. The normalization procedure retained 1243 keyword instances and produced 769 unique normalized labels. Table 3 presents the 20 most frequent labels used to construct the keyword co-occurrence network.

Table 3. Twenty most frequent normalized author keywords in the strict corpus.

Rank	Normalized Keyword	Frequency
1	large language models	70
2	machine learning	45
3	generative ai	31
4	generative adversarial networks	29
5	generative design	26
6	artificial intelligence	21
7	diffusion models	16
8	topology optimization	12
9	variational autoencoders	11
10	deep learning	11
11	inverse design	10
12	electronic design automation	9
13	knowledge graph	9
14	additive manufacturing	9
15	multi-objective optimization	9
16	design automation	8
17	design optimization	8
18	data-driven design	8
19	smart manufacturing	7
20	computer-aided design	6

Note: Frequencies count records containing the normalized keyword; synonymous surface forms were consolidated according to the supplementary thesaurus.

Large language models (LLMs) are the most frequent normalized keyword, followed by machine learning, GAI, generative adversarial networks (GANs), and generative design. The list contains model families, broad technological concepts, engineering methods, and application areas. This mixture shows that the vocabulary of industrial GAI research combines computational approaches with design, manufacturing, optimization, and automation tasks. However, keyword frequency alone does not establish that these technologies are technically integrated.

Normalization affected the resulting ranking. The normalized and minimally cleaned raw top-20 lists share 16 terms, producing a Jaccard similarity of 0.667. Thus, four terms differ between the lists because abbreviations, singular and plural forms, spelling variants, and synonymous labels were consolidated during normalization. The raw-versus-normalized comparison is, therefore, reported as a sensitivity assessment rather than assuming that normalization is a neutral preprocessing operation.

Binary document-level co-occurrence was calculated for the 20 keywords in Table 3. Figure 3 displays the 42 strongest of the 94 nonzero edges. Node size represents document frequency, and edge width represents the number of retained records in which the connected keywords co-occur. Node positions were calculated using a weighted Fruchterman–Reingold spring-layout algorithm with a fixed random seed. Stronger edges

exert greater attractive force in this layout. All nodes use the same color, and no thematic categories or manually assigned coordinates determine their positions. The orientation and exact spatial distances are outputs of the visualization algorithm and should not be interpreted as independently measured similarities.




Figure 3. Weighted co-occurrence network of the 20 most frequent normalized author keywords. Node size represents document frequency, and edge width represents co-occurrence weight. Notes: Node positions were calculated using a weighted spring-layout algorithm with a fixed random seed. All degree and weighted-degree values discussed below refer to the 42-edge network displayed in Figure 3. Corresponding statistics for the complete 94-edge network are reported separately in the repository workbook.

The displayed network forms one connected component, but its keywords differ substantially in connectivity. Machine learning has the largest number of displayed connections, with 12 incident edges and a weighted degree of 63, where weighted degree is the sum of its incident co-occurrence weights. LLMs have nine displayed connections and a weighted degree of 52. This distinction is informative: LLMs have the highest document frequency, whereas machine learning connects with the widest range of model families, design approaches, and optimization concepts in the thresholded network.

The strongest LLM relationships are with artificial intelligence, machine learning, electronic design automation, design automation, and GAI. LLMs also co-occur with knowledge graphs, smart manufacturing, generative design, and computer-aided design among the displayed edges. This pattern associates LLM research with both general AI terminology and specialized engineering or industrial contexts. It does not, however, demonstrate that LLMs were operationally integrated with all these technologies.

Machine learning has strong or repeated connections with GANs, generative design, topology optimization, data-driven design, design optimization, multi-objective optimization, artificial intelligence, GAI, LLMs, deep learning, variational autoencoders, and computer-aided design. Its broad connectivity indicates that machine learning terminology is used across multiple forms of industrial generation, prediction, optimization, and automation.

The model families show different local co-occurrence patterns. GANs are most strongly associated with machine learning and are also connected with deep learning and topology optimization. Diffusion models co-occur with inverse design and topology optimization, linking them to constrained engineering generation problems. Variational autoencoders have one connection within the 42-edge thresholded network, namely, with machine learning. GAI connects with LLMs, machine learning, artificial intelligence, additive manufacturing, multi-objective optimization, and smart manufacturing.

Keywords with fewer displayed connections should not be interpreted as unimportant or isolated. Their apparent peripheral positions may reflect lower document frequencies or the restriction of Figure 3 to the 42 strongest edges. For example, electronic design automation has a strong connection to LLMs, while knowledge graphs connect to LLMs and smart manufacturing. The complete 94-edge network is, therefore, provided in the supplementary workbook.

Keyword co-occurrence represents thematic association at the publication level. It does not demonstrate technical interoperability, causal relationships, successful implementation, or operational industrial adoption. Figure 3 should, consequently, be interpreted as a map of the vocabulary relationships reported in the retained literature.

Because displaying all edge values would reduce readability, Table 4 reports the 15 strongest associations numerically.

Table 4. Fifteen strongest normalized-keyword co-occurrence edges.

Keyword 1	Keyword 2	Weight
large language models	artificial intelligence	11
machine learning	generative adversarial networks	10
large language models	machine learning	8
large language models	electronic design automation	7
large language models	design automation	6
large language models	generative ai	6
machine learning	generative design	6
machine learning	artificial intelligence	5
machine learning	data-driven design	5
machine learning	design optimization	5
machine learning	generative ai	5
machine learning	multi-objective optimization	5
machine learning	topology optimization	5
generative ai	additive manufacturing	4
generative ai	artificial intelligence	4

Note: Weight is the number of retained records in which both normalized author keywords occur.

Table 4 quantifies the relationships that visually dominate the network. LLMs have their strongest association with artificial intelligence, whereas machine learning has its strongest association with GANs. The remaining leading edges connect general AI concepts and model families with design, optimization, manufacturing, and automation terminology. These relationships demonstrate cross-domain thematic overlap in the literature without implying technical integration or deployment. The complete table of 94 nonzero edges is supplied in the supplementary workbook.

4.3. Dominant Themes and Cross-Domain Overlap

Table 5 presents a mutually exclusive classification based on the primary research objective of each retained study. Each record is assigned to one dominant category, providing an interpretable overview of the principal application areas. Table 6 subsequently

uses non-exclusive secondary tags to capture thematic relationships that extend across these categories.

Table 5. Dominant thematic categories by publication year.

No.	Dominant Research Topic	2022	2023	2024	2025	2026 *	Total
1	Mechanical/component generative design, topology optimization, and additive manufacturing	2	2	2	9	9	24
2	Materials discovery, molecular/alloy/catalyst design, and property prediction	2	8	8	23	20	61
3	LLM/RAG/knowledge graph/agent systems for engineering design, CAD, and process planning	0	0	2	18	11	31
4	Lattice, metamaterial, architected, biomimetic, and biomedical structure design	4	0	2	6	4	16
5	Civil, architectural, building, construction, and infrastructural applications	0	2	3	7	2	14
6	Product styling, customization, human–AI co-design, and user-driven design	0	2	4	10	14	30
7	GAN/VAE/diffusion/synthetic data approaches for process modeling, monitoring, reliability, and quality	9	5	3	4	3	24
8	Electronics, semiconductor, EDA, photonics, and electrical device design	2	4	5	15	7	33
9	Factory/smart-manufacturing systems, robotics, digital twins, production systems, and supply chains	1	1	0	4	6	12
10	Cross-domain industrial GAI and enabling systems	1	2	0	2	2	7

Note: Each study is counted once according to its dominant research objective. The asterisk (*) indicates incomplete annual data as of the search date.

Table 6. Secondary-tag overlap matrix.

Secondary_Tag	Engineering Generation	Materials	LLM/Knowledge Systems	Manufacturing/ Process	Electronics/ EDA	Human–AI/ Product	Robotics/ Digital Twins
Engineering generation	166	77	73	148	30	38	24
Materials	77	98	37	94	12	10	4
LLM/knowledge systems	73	37	109	93	27	43	26
Manufacturing/process	148	94	93	222	38	59	36
Electronics/EDA	30	12	27	38	49	11	17
Human–AI/product	38	10	43	59	11	70	19
Robotics/digital twins	24	4	26	36	17	19	41

Note: Diagonal values are individual tag frequencies, and off-diagonal values are joint assignments. Multiple secondary tags may be assigned to each study; therefore, diagonal frequencies should not be summed as mutually exclusive totals.

Materials discovery is the largest dominant category, accounting for 61 studies. It is followed by electronics, semiconductor, and EDA applications with 33 studies; LLM, RAG, knowledge graph, and agent systems with 31; and product styling and human–AI co-design with 30. Mechanical and component generation and GAN-, VAE-, diffusion-, and synthetic-data-based process applications each contain 24 studies.

The annual distribution indicates differences in the development of the principal themes. GAN-, VAE-, diffusion-, and synthetic-data-based process applications are represented from the beginning of the observation period, with nine studies in 2022. In contrast, the dedicated LLM and knowledge systems category first appears in 2024 and expands substantially in 2025. Materials discovery remains prominent throughout the period, while product styling and human–AI co-design also increase in the later years. Because the 2026

search period is incomplete, its counts should not be compared directly with complete annual totals.

Representative studies illustrate the breadth of the dominant categories. Product and mechanical applications include LLM-assisted engineering [27], product appearance design [28], controllable concept generation and analogy-based ideation [29,30], and human–LLM collaboration [14]. Learned generative models support structural component generation [31], metamaterial design [32], topology optimization [33–35], and lattice design [36]. EDA research includes domain reviews [13,37], interactive copilots [38], and autonomous agents for circuit and layout workflows [39]. Process-planning research includes multi-agent LLM-based route generation [12] and the integration of LLMs, smaller language models, and knowledge graphs [40]. Materials and smart-manufacturing studies extend the field toward foundation-model-supported production [6,41] and learned materials generation and discovery [42,43].

The dominant classification necessarily simplifies studies that address several domains simultaneously. Table 6, therefore, reports non-exclusive secondary-tag assignments. Diagonal values show the number of records receiving each tag, whereas off-diagonal values indicate the number receiving both corresponding tags.

Table 6 complements the dominant classification by reporting joint secondary-tag assignments. The diagonal gives the number of records receiving each tag, and the off-diagonal cells show how often two tags occur together.

Manufacturing/process is the most frequent secondary tag, appearing in 222 records, followed by engineering generation in 166, LLM/knowledge systems in 109, and materials in 98. This distribution shows that manufacturing and process considerations occur across many studies whose primary objectives are classified elsewhere.

The largest joint assignment count connects manufacturing/process with engineering generation, involving 148 records. Manufacturing/process also overlaps extensively with materials in 94 records and LLM/knowledge systems in 93. Engineering generation co-occurs with materials in 77 records and with LLM/knowledge systems in 73. Human–AI/product applications overlap with manufacturing/process in 59 records and with LLM/knowledge systems in 43, while robotics/digital-twin studies most frequently intersect with manufacturing/process.

These values represent raw joint assignment counts and are influenced by the overall prevalence of the individual tags. They should, therefore, not be interpreted as normalized measures of association strength. Nevertheless, the overlap matrix demonstrates that the dominant categories are analytical summaries rather than isolated research domains. The retained literature frequently combines artifact generation, materials knowledge, process constraints, electronics, human interaction, and knowledge-intensive assistance within the same studies.

4.4. Abstract-Reported Evidence and Adoption Maturity

Table 7 summarizes the evidence characterization derived from titles and abstracts. The table separates methodological form, validation context, reproducibility resource reporting, and apparent adoption stage so that high publication activity is not interpreted as equivalent to operational readiness.

Table 7 shows that empirical or computational evaluations constitute 72.2% of the corpus, while 19.8% are conceptual or methodological, and 7.9% are reviews. Simulation, benchmark, or case evaluation is reported for 49.2%; industrial cases or operational contexts for 12.3%; and laboratory/prototype validation for 6.3%. Nevertheless, 84.5% are classified as research-stage, 14.7% as pilot/case, and only 0.8% as indicating operational evidence. Code or data resources are mentioned in only 2.0% of abstracts.

Table 7. Abstract-reported evidence characterization of the 252 included studies.

Dimension	Category	N	Percent
Study_Type	Empirical/computational evaluation	182	72.2
Study_Type	Conceptual/methodological	50	19.8
Study_Type	Review/synthesis	20	7.9
Validation_Context	Simulation/benchmark/case evaluation reported	124	49.2
Validation_Context	No empirical validation identified in abstract	81	32.1
Validation_Context	Industrial case or operational context reported	31	12.3
Validation_Context	Laboratory/prototype validation reported	16	6.3
Reproducibility_Resource	Not reported in abstract	247	98.0
Reproducibility_Resource	Code/data resource reported	5	2.0
Industrial_Adoption	Research-stage evidence	213	84.5
Industrial_Adoption	Pilot/case context indicated	37	14.7
Industrial_Adoption	Operational/industrial evidence indicated	2	0.8

Note: Categories describe signals reported in titles and abstracts. They are not a formal full-text quality or risk-of-bias assessment.

The obtained results answer the research questions at complementary levels. Publication and keyword analyses describe field development; the network and thematic tables reveal two connected research trajectories; the overlap matrix demonstrates cross-domain integration; and the evidence characterization shows that deployment maturity remains substantially lower than bibliometric visibility. These findings provide the basis for the interpretive and stakeholder-oriented discussion that follows.

5. Discussion

The results reveal a rapidly expanding but uneven field. This section moves beyond description by interpreting why the technological emphasis has changed, how the application domains can be connected as an industrial system, what evidential limits remain, and which actions are appropriate for different stakeholders.

5.1. From Engineering Generation to Integrated Industrial Intelligence

The corpus does not show a simple replacement of generative engineering by LLM-based industrial intelligence. Instead, it shows convergence. The earlier engineering trajectory concentrated on constrained generation of shapes, structures, materials, and process representations. The later LLM trajectory adds conversational access, knowledge retrieval, task decomposition, agent coordination, and tool use. Foundation model availability after 2022, wider digitization of engineering documentation, improved access to computational infrastructure, and demand for accessible human–tool interfaces help explain why LLM, RAG, knowledge graph, and agent studies expanded sharply in 2025–2026 [4,6,7,9,12,15,18].

This transition also reflects a change in the object being generated. Earlier systems primarily generated an artifact or candidate design. Newer systems increasingly generate an intermediate representation, explanation, plan, command, query, or coordinated sequence of actions. Such outputs can connect unstructured requirements with engineering tools and enterprise knowledge, but their value depends on grounding and verification. Fluent interaction alone cannot guarantee that an output satisfies physics, manufacturability, regulation, or company-specific practice.

The strongest cross-topic overlaps indicate a potential workflow chain: materials models propose candidates; generative engineering explores geometry; simulation and manufacturability checks constrain alternatives; EDA or process-planning tools translate designs into domain artifacts; digital twins and production systems provide operational context; and LLM/RAG agents connect documents, tools, and human decisions. This life-

cycle interpretation explains why industrial value lies less in a single model than in verified integration across data, knowledge, physics, software, and accountable human work.

5.2. Challenges, Evidence, and Systems Integration

The maturity profile tempers claims based on publication growth. A high keyword frequency or dense network position measures research attention, not comparative effectiveness. The small share of operational evidence reflects persistent obstacles: proprietary and heterogeneous data, hallucination and unverifiable output, weak physical constraint handling, limited benchmark comparability, model drift, intellectual property exposure, cybersecurity, computational cost, and integration with CAD/CAE, product lifecycle management (PLM), manufacturing execution (MES), enterprise resource planning (ERP), and quality systems [2–6,9–13].

Research-stage systems may perform well on curated benchmarks yet fail when faced with changing product variants, incomplete documentation, safety constraints, or legacy interfaces. Industrial evaluation should report not only model accuracy but also constraint satisfaction, provenance, failure modes, human override, latency, integration effort, lifecycle monitoring, and economic value. These measures are necessary to connect technical performance with systems practice.

Data governance is particularly important because industrial GAI often combines intellectual property, customer requirements, engineering drawings, maintenance records, and operational data. Organizations must define which information can be used for prompting, retrieval, fine tuning, logging, and external service calls. Generated artifacts also require ownership, version control, and traceability so that later design or operational decisions can be reconstructed.

5.3. Stakeholder-Oriented Recommendations

Table 8 translates the review findings into differentiated recommendations. This distinction is important because methodological, organizational, technical, and regulatory responsibilities cannot be assigned to one stakeholder group alone.

Table 8. Stakeholder-specific implications derived from this review.

Stakeholder	Priority Actions
Researchers	Shared benchmarks; non-generative baselines; external validity; code/data and failure reporting
Managers	Bounded use cases; data/integration readiness; value metrics; workforce preparation; ownership
Engineers	Approved sources; physical/manufacturability checks; provenance; testing; human approval
Policymakers/standards bodies	Accountability, IP, safety, cybersecurity, auditability, interoperability, and SME support

Note: The recommendations synthesize recurring needs in the reviewed literature and the evidence maturity results; they are not tested causal effects.

Table 8 indicates that researchers should strengthen reproducibility and comparative evaluation, managers should focus on organizational readiness and measurable use cases, engineers should establish verification and approval controls, and policymakers should support common rules for accountability and interoperability. Their actions are complementary: a technically strong model cannot be deployed responsibly without suitable data governance, workflow ownership, and standards.

A practical maturity sequence begins with low-autonomy, reversible applications such as document retrieval, drafting, design ideation, and synthetic test-data support. Pilot deployment should add domain grounding, access control, logging, benchmark and user evaluation, and explicit approval gates. Higher-impact applications involving geometry release, process planning, equipment control, quality disposition, or safety decisions require independent verification, configuration management, monitoring, and accountable human authorization.

5.4. Limitations

This review is limited by its three selected databases, English-language restriction, 2022–2026 window, search-string dependence, and the incomplete 2026 observation. Screening and thematic synthesis rely mainly on indexed titles, abstracts, and keywords; full-text outcomes were not extracted systematically for all 252 studies. The evidence characterization is rule-based and abstract-reported, not a validated risk-of-bias instrument. No formal inter-rater reliability coefficient was calculated. Keyword normalization includes judgment, although the thesaurus, raw comparison, and sensitivity results are disclosed. Citation-based impact was not used because recent papers have unequal exposure time. Finally, rapid model and product change makes implementation conclusions time-sensitive.

These limitations define the appropriate interpretation of this review. This study provides a reproducible map of indexed research and reported maturity signals, not a definitive comparison of model effectiveness. Future reviews can extend the database coverage, conduct independent double coding, extract full-text validation outcomes, and update the corpus as industrial deployments become better documented.

Overall, the discussion connects the bibliometric structure with technological convergence and organizational readiness. Industrial GAI appears most promising when generative capabilities are embedded in governed workflows that combine domain knowledge, computational verification, interoperable tools, and human accountability. This systems interpretation informs the final conclusions.

6. Conclusions

This review set out to clarify how learned generative models are being studied across manufacturing and industrial systems, how the application domains connect, and how far the published literature indicates movement toward operational adoption. The final 252-study corpus provides a stricter basis for answering these questions than a search set containing classical generative design records.

Two increasingly connected trajectories organize the literature. Engineering generation concerns artifacts, structures, materials, and electronics; industrial intelligence concerns LLMs, RAG, knowledge graphs, agents, and human–AI collaboration. Their overlap suggests that future value will come from integrated workflows in which models generate candidates, knowledge systems provide context, simulation and engineering rules verify feasibility, and humans retain responsibility for consequential decisions.

The evidence remains less mature than publication growth alone suggests. Most records are research-stage, and only a small fraction indicate operational industrial use in their abstracts. Organizations should, therefore, progress from bounded and reversible applications toward higher-impact use only when data governance, domain grounding, manufacturability, traceability, cybersecurity, intellectual property protection, interoperability, monitoring, workforce competence, and accountable oversight are in place.

Future research should publish reproducible benchmarks, compare GAI-enabled approaches with appropriate conventional baselines, report failures and lifecycle costs, and validate integrated workflows in real industrial settings. Longitudinal studies are also needed to determine whether promising pilots produce sustained improvements in quality, productivity, sustainability, and engineering decision making. In this way, industrial GAI can develop from a collection of visible research topics into trustworthy systems practice.

Supplementary Materials: The following supporting information can be downloaded at <https://www.mdpi.com/article/10.3390/electronics15153485/s1>, (Section S1: Database-Specific Search Strategies; Section S2: Keyword Normalization and Co-Occurrence Mapping; Section S3: Reproducibility Tables for Keyword Frequencies and Co-Occurrence). The supplementary pack-

age contains the PRISMA 2020 checklist and a methodological supplement with the final search strings, screening audit, exclusion reasons, coding protocol, secondary-tag overlap matrix, keyword normalization thesaurus and sensitivity results, co-occurrence data, and abstract-reported evidence characterization. Corresponding data files are available through Mendeley Data at <https://doi.org/10.17632/nrssr637h3>.

Author Contributions: Conceptualization, G.I.; methodology, G.I. and Y.I.; data curation, G.I. and Y.I.; formal analysis, G.I.; writing—original draft preparation, G.I.; writing—review and editing, G.I. and Y.I.; visualization, G.I. and Y.I.; supervision, Y.I. All authors have read and agreed to the published version of the manuscript.

Funding: This research was partially supported by the project “Digitalisation of the Economy in a Big Data Environment” (BG16RFPR002-1.014-0013), funded by the European Regional Development Fund through the program “Research, Innovation and Digitalisation for Smart Transformation” 2021–2027.

Data Availability Statement: The screening audit, retained and excluded record lists, normalized keyword data, co-occurrence results, thematic coding, evidence characterization, and analysis code are available in Mendeley Data, Version 2: <https://doi.org/10.17632/nrssr637h3.2>.

Acknowledgments: The authors thank the Academic Editor and the three anonymous reviewers for their constructive comments.

Conflicts of Interest: Author Yuliy Iliev was employed by Teletek Electronics. The remaining author declares that this research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Abbreviations

The following abbreviations are used in this manuscript:

Abbreviation	Meaning
ACM	Association for Computing Machinery
AI	Artificial intelligence
CAD	Computer-aided design
CAE	Computer-aided engineering
EDA	Electronic design automation
ERP	Enterprise resource planning
GAI	Generative artificial intelligence
GAN	Generative adversarial network
LLM	Large language model
MES	Manufacturing execution system
PLM	Product lifecycle management
PRISMA	Preferred Reporting Items for Systematic Reviews and Meta-Analyses
RAG	Retrieval-augmented generation
SME	Small and medium-sized enterprise
VAE	Variational autoencoder

References

1. Doanh, D.C.; Dufek, Z.; Ejdy, J.; Ginevičius, R.; Korzynski, P.; Mazurek, G.; Paliszkievicz, J.; Wach, K.; Ziemba, E. Generative AI in the manufacturing process: Theoretical considerations. *Eng. Manag. Prod. Serv.* **2023**, *15*, 76–89. [\[CrossRef\]](#)
2. Shahin, M.; Hosseinzadeh, A.; Chen, F.F. Generative artificial intelligence in manufacturing: Applications, case studies, and future directions for next-generation intelligent production systems. *Int. J. Adv. Manuf. Technol.* **2025**, *141*, 1159–1265. [\[CrossRef\]](#)
3. Leng, J.; Su, X.; Liu, Z.; Zhou, L.; Chen, C.; Guo, X.; Wang, Y.; Wang, R.; Zhang, C.; Liu, Q.; et al. Diffusion model-driven smart design and manufacturing: Prospects and challenges. *J. Manuf. Syst.* **2025**, *82*, 561–577. [\[CrossRef\]](#)
4. Yu, G.; Wang, Y.; Wang, Z.; Chen, L.; Zheng, Y.; Liu, Y. LLMs in industrial domains: A systematic review of adaptation techniques and applications from the product lifecycle perspective. *Adv. Eng. Inform.* **2026**, *74*, 104655. [\[CrossRef\]](#)

5. Peckham, O.; Raines, J.; Bulsink, E.; Goudswaard, M.; Gopsill, J.; Barton, D.; Nassehi, A.; Hicks, B. Artificial intelligence in generative design: A structured review of trends and opportunities in techniques and applications. *Designs* **2025**, *9*, 79. [\[CrossRef\]](#)
6. Chen, C.; Li, L.; Wang, T.; Wang, Z.; Leng, J.; Shao, H.; Cheng, L. Leveraging large language models for smart manufacturing: Reviews, enablers, challenges, and opportunities. *J. Manuf. Syst.* **2026**, *85*, 593–612. [\[CrossRef\]](#)
7. Li, Z.; Liu, Z.; Sa, G.; Sun, J.; Hou, M.; Tan, J.; Sun, L.; Wei, J. Knowledge-enhanced large language models for ideation to implementation: A new paradigm in product design. *Appl. Soft Comput.* **2025**, *176*, 113147. [\[CrossRef\]](#)
8. Wang, Z.; Yan, X.; Melkote, S.; Rosen, D. McGAN: Generating manufacturable designs by embedding manufacturing rules into conditional generative adversarial network. *Adv. Eng. Inform.* **2025**, *64*, 103074. [\[CrossRef\]](#)
9. Li, S.; Corney, J. MechRAG: A multimodal large language model for mechanical engineering. *Commun. Eng.* **2025**, *4*, 187. [\[CrossRef\]](#) [\[PubMed\]](#)
10. Mata, O.; Ponce, P.; Perez, C.; Ramirez, M.; Anthony, B.; Russel, B.; Apte, P.; MacCleery, B.; Molina, A. Digital twin designs with generative AI: Crafting a comprehensive framework for manufacturing systems. *J. Intell. Manuf.* **2026**, *37*, 1049–1072. [\[CrossRef\]](#)
11. Daareyni, A.; Martikkala, A.; Mokhtarian, H.; Ituarte, I.F. Generative AI meets CAD: Enhancing engineering design to manufacturing processes with large language models. *Int. J. Adv. Manuf. Technol.* **2025**, 1–10. [\[CrossRef\]](#)
12. Xie, Y.; Liu, J.; Wang, R.; Wang, Z.; Yu, K.; Song, Z. Rapid generation method of process routes based on multi-agent collaboration with LLMs. *Adv. Eng. Inform.* **2025**, *68*, 103733. [\[CrossRef\]](#)
13. Pan, J.Y.; Zhou, G.L.; Chang, C.C.; Jacobson, I.; Hu, J.; Chen, Y.R. A survey of research in large language models for electronic design automation. *ACM Trans. Des. Autom. Electron. Syst.* **2025**, *30*, 1–21. [\[CrossRef\]](#)
14. Wang, X.; Jiang, Z.; Xiong, Y.; Liu, A. Human–LLM collaboration in generative design for customization. *J. Manuf. Syst.* **2025**, *80*, 425–435. [\[CrossRef\]](#)
15. Oyekan, J.; Turner, C.; Bax, M.; Graf, E. From ontologies to knowledge-augmented large language models for automation: Decision-making guidance for achieving human–robot collaboration in Industry 5.0. *Comput. Ind.* **2025**, *171*, 104329. [\[CrossRef\]](#)
16. Sikandar, H.; Khan, N.; Falahat, M.; Qureshi, M.I. Generative AI for sustainable product design: A technology convergence framework integrating multi-objective optimisation and smart manufacturing. *IET Collab. Intell. Manuf.* **2026**, *8*, e70051. [\[CrossRef\]](#)
17. Akhtar, P.; Ghouri, A.M.; Ashraf, A.; Lim, J.J.; Khan, N.R.; Ma, S. Smart product platforming powered by AI and generative AI: Personalization for the circular economy. *Int. J. Prod. Econ.* **2024**, *273*, 109283. [\[CrossRef\]](#)
18. Liang, X.; Wang, Z.; Liu, J. A survey of large language model-augmented knowledge graphs for advanced complex product design. *J. Manuf. Syst.* **2025**, *80*, 883–901. [\[CrossRef\]](#)
19. Page, M.J.; McKenzie, J.E.; Bossuyt, P.M.; Boutron, I.; Hoffmann, T.C.; Mulrow, C.D.; Shamseer, L.; Tetzlaff, J.M.; Akl, E.A.; Brennan, S.E.; et al. The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ* **2021**, *372*, n71. [\[CrossRef\]](#) [\[PubMed\]](#)
20. Liu, F.; Shao, Y.; Chen, M. Generative design strategies for additive manufacturing of lattice structures: A review. *Mater. Des.* **2026**, *262*, 115431. [\[CrossRef\]](#)
21. Jabbar, R.; Jabbar, R.; Kamoun, S. Recent progress in generative adversarial networks applied to inversely designing inorganic materials: A brief review. *Comput. Mater. Sci.* **2022**, *213*, 111612. [\[CrossRef\]](#)
22. Badini, S.; Regondi, S.; Pugliese, R. Enhancing mechanical and bioinspired materials through generative AI approaches. *Next Mater.* **2025**, *6*, 100275. [\[CrossRef\]](#)
23. Choi, J.; Nam, G.; Choi, J.; Jung, Y. A perspective on foundation models in chemistry. *JACS Au* **2025**, *5*, 1499–1518. [\[CrossRef\]](#) [\[PubMed\]](#)
24. Song, Z.; Cai, Y.; Yu, Y.; Xiang, P.; Li, M.; Liu, Z.; Dai, M. New opportunities for the integration of artificial intelligence with materials science: From large language models to embodied large models. *Mater. Genome Eng. Adv.* **2026**, *4*, e70042. [\[CrossRef\]](#)
25. Li, C.; Xian, Y.; Zhou, Y.; Ding, X.; Sun, J.; Xue, D. Generative design for alloys: Harnessing generative models for faster discovery. *Adv. Mater.* **2026**, *38*, e20478. [\[CrossRef\]](#) [\[PubMed\]](#)
26. Tang, J.; Cao, L.; Xu, G.; Li, M. From high-throughput screening to generative design: Artificial intelligence-driven dielectric materials discovery. *ACS Appl. Mater. Interfaces* **2026**, *18*, 29529–29563. [\[CrossRef\]](#) [\[PubMed\]](#)
27. Mustapha, K.B. A survey of emerging applications of large language models for problems in mechanics, product design, and manufacturing. *Adv. Eng. Inform.* **2025**, *64*, 103066. [\[CrossRef\]](#)
28. Liu, Y.; Wang, S.; Li, S.; Chen, G. From deep learning to generative AI: Research status and prospects of product appearance design in the artificial intelligence era. *J. Eng. Des.* **2026**, *37*, 1776–1817. [\[CrossRef\]](#)
29. Chen, L.Q.; Zuo, H.Y.; Cai, Z.B.; Yin, Y.; Zhang, Y.; Sun, L.Y.; Childs, P.; Wang, B.H. Toward controllable generative design: A conceptual design generation approach leveraging the function–behavior–structure ontology and large language models. *J. Mech. Des.* **2024**, *146*, 121401. [\[CrossRef\]](#)
30. Wang, B.; Zhao, X.; Zuo, H.; Song, Y.; Han, J.; Childs, P.; Chen, L. From analogy to innovation: A creative conceptual design approach leveraging large language models. *Adv. Eng. Inform.* **2025**, *67*, 103427. [\[CrossRef\]](#)

31. Li, J.; Du, W.; Xue, H.; Yang, M. Intelligent generative design of novel aluminum joints connecting wood members using generative adversarial networks. *Structures* **2026**, *87*, 111755. [[CrossRef](#)]
32. Wang, Z.; Xu, H. Manufacturability-aware deep generative design of 3D metamaterial units for additive manufacturing. *Struct. Multidiscip. Optim.* **2024**, *67*, 22. [[CrossRef](#)]
33. Wang, Z.C.; Melkote, S.; Rosen, D.W. Generative design by embedding topology optimization into conditional generative adversarial network. *J. Mech. Des.* **2023**, *145*, 111702. [[CrossRef](#)]
34. Parrott, C.M.; Abueidda, D.W.; James, K.A. Multidisciplinary topology optimization using generative adversarial networks for physics-based design enhancement. *J. Mech. Des.* **2023**, *145*, 061704. [[CrossRef](#)]
35. Zhang, W.; Zhao, G.; Su, L. Research on multi-stage topology optimization method based on latent diffusion model. *Adv. Eng. Inform.* **2025**, *63*, 102966. [[CrossRef](#)]
36. Zhang, J.; Chen, S.; Martin, R.J.; Liu, B.; Zhang, R.; Xiao, D. Conditional diffusion models for the inverse design of lattice structures. *Struct. Multidiscip. Optim.* **2025**, *68*, 58. [[CrossRef](#)]
37. Noori Zadeh, D.; Elamien, M.B. Generative AI for analog integrated circuit design: Methodologies and applications. *IEEE Access* **2025**, *13*, 58043–58059. [[CrossRef](#)]
38. Liu, B.; Zhang, H.; Gao, X.; Kong, Z.; Tang, X.; Lin, Y.; Wang, R.; Huang, R. LayoutCopilot: An LLM-powered multiagent collaborative framework for interactive analog layout design. *IEEE Trans. Comput. Aided Des. Integr. Circuits Syst.* **2025**, *44*, 3126–3139. [[CrossRef](#)]
39. Wu, H.Y.; He, Z.L.; Zhang, X.Y.; Yao, X.F.; Zheng, S.; Zheng, H.S.; Yu, B. ChatEDA: A large language model-powered autonomous agent for EDA. *IEEE Trans. Comput. Aided Des. Integr. Circuits Syst.* **2024**, *43*, 3184–3197. [[CrossRef](#)]
40. Ma, Y.; Zheng, S.; Yang, Z.; Zheng, P.; Leng, J.; Hong, J. Aircraft assembly process planning based on a knowledge graph constructed by integrating LLMs and SLMs. *J. Manuf. Syst.* **2026**, *84*, 1–19. [[CrossRef](#)]
41. Zheng, H.; Geng, Z.Y.; Xu, X. Foundation model-based generative AI for smart manufacturing: A paradigm shift. *Eng. Manag.* **2026**, *13*, 265–291. [[CrossRef](#)]
42. Liu, Y.; Yang, Z.W.; Yu, Z.Y.; Liu, Z.T.; Liu, D.H.; Lin, H.L.; Li, M.Q.; Ma, S.C.; Avdeev, M.; Shi, S.Q. Generative artificial intelligence and its applications in materials science: Current situation and future perspectives. *J. Mater.* **2023**, *9*, 798–816. [[CrossRef](#)]
43. Alverson, M.; Baird, S.G.; Murdock, R.; Ho, S.H.; Johnson, J.; Sparks, T.D. Generative adversarial networks and diffusion models in material discovery. *Digit. Discov.* **2024**, *3*, 62–80. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.