Generative Artificial Intelligence for Industrial Design and Manufacturing: Generative Design, Synthetic Data Generation, and Diffusion-Based Optimization

Author: Loius Nuudle

Abstract

The growing complexity of modern manufacturing—encompassing additive manufacturing, precision assembly, and high-mix low-volume production—demands design and optimization tools that can explore vast design spaces, generate high-quality training data for data-driven models, and accelerate the transition from conceptual design to validated product. **Generative artificial intelligence (Generative AI)**—comprising generative adversarial networks (GANs), variational autoencoders (VAEs), diffusion models, and large language models—has emerged as a transformative toolkit for industrial design and manufacturing, enabling the automatic generation of novel designs, the synthesis of training data for defect detection and quality inspection, and the optimization of manufacturing processes under complex physical constraints. This review provides a comprehensive and critical synthesis of the application of generative AI in industrial design and manufacturing. We examine three major application domains: **generative design** enabled by deep learning and reinforcement learning for topology optimization of additively manufactured structures; **synthetic data generation** using GANs and diffusion models to address data scarcity in industrial defect detection; and **diffusion-based 3D shape generation and microstructure design** for materials discovery and manufacturing process optimization. We further connect these generative AI methods to advances in industrial sensing—precision 3D optical metrology, four-dimensional thermal imaging, and collaborative robotic inspection—demonstrating their complementary roles in the intelligent manufacturing ecosystem. A central contribution of this review is the articulation of a unified **Generative-to-Physical (G2P) pipeline** framework that connects generative AI design exploration, physics-based validation, and manufacturing process control, charting a course toward AI-augmented creative engineering.

Keywords: Generative AI; Generative Design; Synthetic Data; GAN; Diffusion Models; Topology Optimization; Industrial Defect Detection; Additive Manufacturing; 3D Shape Generation; AI-Augmented Design

1. Introduction

The design of complex industrial products—from aerospace components and automotive structures to microelectronics and medical devices—is fundamentally a creative search problem conducted under constraints. Engineers must simultaneously satisfy functional requirements (structural integrity, thermal performance, fluid dynamics), manufacturing constraints (machinability, tolerances, process capabilities), and resource limitations (material costs, production time, energy consumption). Traditional computer-aided design (CAD) workflows begin with a human engineer's conceptual design, which is then iteratively refined through simulation and testing—a process that is time-consuming, expertise-dependent, and inherently limited by the human designer's ability to explore the vast space of possible solutions.

The past decade has witnessed a paradigm shift toward **generative design**—the automatic generation of designs by algorithms that explore design spaces subject to user-specified constraints and objectives. Initially, generative design was synonymous with topology optimization—a mathematical approach that optimizes material distribution within a design domain to minimize a structural compliance objective subject to constraints. However, classical topology optimization methods are computationally expensive for complex 3D problems, generate designs that are difficult to interpret and manufacture, and cannot incorporate aesthetic, stylistic, or functional preferences that are difficult to formalize mathematically.

The emergence of **generative AI**—deep learning models that learn to generate novel data samples from learned data distributions—has fundamentally transformed generative design. Rather than solving optimization problems from scratch, generative AI models learn design distributions from existing design datasets and generate novel designs that are simultaneously diverse (exploring the design space broadly) and high-quality (conforming to the learned distribution of feasible designs). This shift from optimization-based to learning-based generation represents a qualitative change in the nature of the design process: from iterative refinement of a single design candidate to sampling from a learned generative model that implicitly encodes the full space of feasible designs.

This review examines the application of generative AI to three critical domains in industrial design and manufacturing: **generative design** for topology optimization of additively manufactured structures; **synthetic data generation** for industrial defect detection and quality inspection; and **diffusion-based 3D shape and microstructure generation** for materials design and manufacturing process optimization. We demonstrate how these generative AI methods complement and integrate with advances in industrial sensing—including stereo phase-measuring deflectometry (SPMD), four-dimensional thermal imaging, and collaborative robotic inspection—to create a **Generative-to-Physical (G2P) pipeline** that spans design exploration, physics validation, and manufacturing process control.

Our contributions are: (1) a systematic review of generative AI methods for industrial design, covering GANs, VAEs, diffusion models, and their applications to defect detection, design optimization, and microstructure generation; (2) an analysis of the G2P pipeline concept, connecting generative AI to physical validation and manufacturing process control; (3) a synthesis of the complementary roles of industrial sensing and generative AI in the intelligent manufacturing ecosystem; and (4) an articulation of open challenges and future research directions.

The review is organized as follows: Section 2 reviews generative design and topology optimization; Section 3 examines synthetic data generation for industrial defect detection; Section 4 covers diffusion models for 3D shape and microstructure generation; Section 5 provides cross-cutting synthesis; and Section 6 concludes.

2. Generative Design: Deep Learning and Reinforcement Learning for Topology Optimization

2.1 From Classical Topology Optimization to Generative AI

Classical topology optimization—formalized through the Solid Isotropic Material with Penalization (SIMP) method and the Evolutionary Structural Optimization (ESO) approach—optimizes material distribution within a design domain by iteratively updating a density field according to gradient information computed from finite element analysis (FEA). While mathematically rigorous, classical methods suffer from several practical limitations: high computational cost for fine-resolution 3D problems; difficulty incorporating manufacturing constraints (support structures, minimum feature sizes, overhang angles) directly into the optimization; sensitivity to mesh discretization and initial conditions; and a tendency to generate complex, porous structures that are difficult to interpret, manufacture, or modify.

Generative AI addresses these limitations by learning design distributions from existing design datasets—computational or experimental—and generating novel designs through sampling. A generative model trained on a dataset of high-performing designs implicitly encodes the structural principles, aesthetic preferences, and manufacturing constraints that characterize good designs, and can generate new candidates that are diverse yet high-quality without re-solving the optimization from scratch.

2.2 Deep Learning for Generative Design of Additively Manufactured Structures

A 2024 *ResearchGate* study—*Integrating Deep Learning with Generative Design and Topology Optimization for Efficient Additive Manufacturing*—provided a comprehensive examination of how deep learning integrates with generative design and topology optimization in additive manufacturing (AM). Additive manufacturing is particularly amenable to generative design because it imposes fewer geometric constraints than subtractive manufacturing: complex lattice structures, organic shapes, and multi-scale architectures that are impossible to machine can be directly printed. However, the same geometric freedom makes design for AM challenging—human designers lack the intuitive understanding of printability and structural performance that experience has built up for conventional manufacturing processes.

Deep learning models trained on AM design datasets can learn the relationships between design parameters, process conditions, and resulting part properties, enabling the generation of AM-compatible designs that satisfy structural, thermal, and manufacturing constraints simultaneously. The authors demonstrated that the integration of deep learning with generative design reduces the number of design iterations required to achieve target performance by up to an order of magnitude compared to conventional iterative FEA-based refinement (ResearchGate, 2024).

2.3 Reinforcement Learning for Topology Optimization

A landmark 2025 study in *Materials Today Communications*—*Reinforcement Learning-Based Topology Optimization for Generative Designed Lightweight Structures*—demonstrated how **proximal policy optimization (PPO)**, a state-of-the-art deep reinforcement learning algorithm, can be applied to the topology optimization of lightweight mechanical structures. The key innovation is framing topology optimization as a sequential decision problem: at each step, the RL agent decides whether to add or remove material at each finite element location; the agent's policy network learns from reward signals (structural compliance, stress constraints, volume fraction) to generate material distributions that minimize structural compliance while satisfying manufacturing constraints.

This RL-based approach offers several advantages over classical gradient-based topology optimization. First, RL is naturally suited to problems with discrete decision spaces and constraint satisfaction requirements—precisely the characteristics of manufacturing-constrained topology optimization. Second, the learned policy network can generate designs in a single forward pass once trained, eliminating the iterative FEA loop that makes classical topology optimization computationally expensive. Third, RL policies can be pre-trained in simulation and fine-tuned on experimental data, enabling transfer from computational to physical design domains (Materials Today Communications, 2025).

2.4 Implicit Neural Representations for Design Optimization

An emerging trend in generative design is the use of **implicit neural representations (INRs)**—neural networks that represent 3D shapes as continuous functions mapping spatial coordinates to occupancy or signed-distance values—as design representations. Unlike explicit mesh or voxel representations, INRs are resolution-independent, enabling optimization at arbitrary scales and smooth gradient-based design exploration. A 2025 patent from Narnia Labs described mapping generated designs into a **low-dimensional parametric space** where topology optimization can be applied analytically, unifying generative exploration and structural optimization within a single differentiable framework (PatSnap, 2025). This differentiable design paradigm—where the entire design, simulation, and optimization pipeline is differentiable—promises to dramatically accelerate the generative design process by enabling end-to-end gradient-based learning.

2.5 Industrial Sensing for Design Validation

The generative design pipeline does not end with the generation of candidate designs: physical validation is essential to verify that generated designs meet structural, thermal, and manufacturing requirements. Industrial sensing technologies provide the data foundation for this validation.

Huang and colleagues' **stereo phase-measuring deflectometry (SPMD)** system (2026)—originally developed for precision 3D surface metrology of optical components—demonstrates how optical measurement technologies serve as critical validation tools in the generative design pipeline. By achieving high-precision, non-contact surface measurement, SPMD enables the rapid verification of additively manufactured parts against the generative design specification, detecting deviations in surface form, waviness, and roughness that may compromise structural or optical performance (Huang et al., 2026). The deep learning-enhanced phase unwrapping component of the SPMD system further enables automated defect detection in the measurement data, creating a closed loop between generative design and physical validation.

---

3. Synthetic Data Generation for Industrial Defect Detection

3.1 The Data Scarcity Problem

Industrial defect detection—the automated identification of defects such as scratches, cracks, voids, and dimensional deviations in manufactured products—is a cornerstone of quality assurance. Deep learning-based defect detection systems have demonstrated superhuman accuracy when trained on large, annotated datasets of defective and non-defective samples. However, the fundamental challenge in industrial defect detection is **data scarcity**: defective samples are inherently rare, expensive to collect, and may represent only a subset of the possible defect types. This data scarcity problem is particularly acute for **novel or rare defect types**—the very cases where detection matters most but labeled data is most scarce.

3.2 Generative Adversarial Networks for Synthetic Defect Generation

Generative adversarial networks (GANs)—comprising a generator network that synthesizes samples and a discriminator network that distinguishes real from synthetic samples—have emerged as a powerful tool for addressing data scarcity in industrial defect detection. By learning to model the distribution of defective samples from limited training data, GANs can generate synthetic defect images that augment real training datasets, improving the robustness and coverage of defect detection models.

A 2024 *Scientific Reports* study—*DG2GAN: Improving Defect Recognition Performance with Generated Defect Image Samples*—demonstrated how DG2GAN, a domain-adaptive GAN architecture, significantly improves defect recognition performance by generating high-quality synthetic defect samples that bridge the gap between limited real training data and the diverse space of real-world defect appearances. The authors showed that DG2GAN-augmented training sets enable defect detection models to generalize better to novel defect types not seen during training—a critical capability for real-world deployment where defect types evolve as manufacturing processes, materials, and products change (Scientific Reports, 2024).

A comprehensive review in the *Journal of Computational Design and Engineering* (2024) examined imbalanced fault diagnosis technology based on GANs in industrial applications, documenting a range of GAN architectures—including **CycleGAN**, **ISU-GAN**, and **semi-supervised CAE-SGAN**—applied to defect data augmentation in rolling element bearings, hot-rolled steel sheets, and infrared thermography inspection. The review identified three key requirements for effective GAN-based defect augmentation: **defect fidelity** (synthetic defects must be realistic enough to improve detection models), **defect diversity** (the generated defects must span the space of real defect variations), and **domain adaptation** (synthetic defects must match the characteristics of the target inspection domain, including sensor modality, lighting, and material appearance) (JCDE, 2024).

3.3 Diffusion Models for Synthetic Defect Generation

A more recent development is the application of **diffusion probabilistic models** (diffusion models) to synthetic defect generation. Unlike GANs, which learn a generator through adversarial training, diffusion models learn to synthesize data by denoising random noise through an iterative refinement process. Diffusion models have demonstrated superior synthesis quality and diversity compared to GANs in many domains—including image synthesis, molecular design, and audio generation—and are less prone to mode collapse (the generator's tendency to produce a limited variety of samples).

A 2024 *ScienceDirect* study applied **latent diffusion models** (Stable Diffusion) to generate synthetic defect images for improving the robustness of visual defect segmentation networks in steel surface inspection. The authors showed that latent diffusion models trained on normal (defect-free) surface images can be adapted to generate defective images through inpainting—masking a region of a normal image and prompting the model to fill the masked region with defect-like textures. The resulting synthetic defect images significantly improved the segmentation network's accuracy on real defective samples, particularly for rare defect types that were absent or underrepresented in the real training set (ScienceDirect, 2024).

3.4 Time-Series Anomaly Detection with GANs

Beyond image-based defect detection, GANs have been applied to **time-series anomaly detection** in industrial systems. A 2024 *ScienceDirect* study developed DEGAN—an **unsupervised anomaly detection framework using GAN discriminators and density estimation**—for infrastructure systems monitoring. The DEGAN framework trains a GAN on normal operational data (sensor readings, process variables); at inference, samples that the discriminator identifies as significantly different from the learned normal distribution are flagged as anomalies. The density estimation component provides an interpretable anomaly score—the estimated probability density of the input under the normal data distribution—enabling the ranking of anomalies by severity and the identification of anomalous sub-sequences within multivariate time series (ScienceDirect, 2024).

3.5 Industrial Sensing for Defect Data Acquisition

The quality of synthetic defect generation depends critically on the quality and diversity of the real defect data used for training. Industrial sensing technologies provide the measurement systems that acquire this foundational data.

Huang and colleagues' **four-dimensional (4D) thermal imaging system** (2023)—which combines structured illumination binocular cameras with infrared thermography to reconstruct temperature fields on non-uniform surfaces—provides a unique modality for thermal defect detection. Thermal imaging reveals subsurface defects (delaminations, voids, inclusions) that are invisible to visible-light cameras, by detecting the anomalous thermal patterns these defects produce under加热 or cooling. The multi-view fusion and emissivity correction capabilities developed in this work are directly applicable to **thermal defect dataset acquisition** for training synthetic defect generators (Huang et al., 2023).

Li and colleagues' **Leap Motion Controller-based gesture control system** (2024) for collaborative robotic manipulators exemplifies another pathway for defect data acquisition: **robot-guided inspection**. Collaborative robots equipped with high-resolution vision sensors, guided by intuitive gesture-based programming, can autonomously navigate inspection trajectories over complex product surfaces, acquiring dense, high-quality image data for defect dataset construction. The gesture-based programming paradigm makes robotic inspection programming accessible to non-expert operators, enabling the rapid deployment of automated inspection systems across diverse product types (Li et al., 2024).

---

4. Diffusion Models for 3D Shape Generation and Microstructure Design

4.1 Diffusion Models: Principles and Industrial Potential

Diffusion models have emerged as the dominant paradigm in generative AI for 2D image synthesis, achieving state-of-the-art results in image generation, image-to-image translation, and conditional synthesis tasks. Their core principle—learning to reverse a gradual noising process that transforms data into random noise—provides several advantages over GANs: more stable training (no adversarial equilibrium), higher sample diversity (less mode collapse), and better scalability to high-resolution data. These advantages have catalyzed rapid adoption in scientific and engineering domains beyond computer vision.

The extension of diffusion models to **3D data**—point clouds, meshes, voxel grids, and implicit representations—has become one of the most active research frontiers in generative AI for manufacturing. 3D generative models enable the automatic creation of novel part geometries, the inverse design of structures with specified properties, and the generative exploration of material microstructures for targeted macroscopic behavior.

4.2 Diffusion Models for 3D Shape Generation

A 2025 *Tsinghua Computational Visual Media* survey—*Diffusion Models for 3D Generation: A Survey*—comprehensively reviewed the rapidly evolving landscape of 3D diffusion models, documenting architectures including **point cloud diffusion** (operating directly on sets of 3D coordinates), **voxel diffusion** (operating on 3D grids), **mesh diffusion** (operating on triangle mesh representations), and **hybrid diffusion** (combining multiple 3D representations). The survey identified the key challenge in 3D diffusion modeling: 3D data is more complex, irregular, and memory-intensive than 2D images, requiring specialized architectures and sampling strategies.

A 2025 *Cell Reports Physical Science* study introduced **MPaDiffusion**—a **multi-modal, property-aware diffusion model for 3D reconstruction and on-demand design** of material microstructures. MPaDiffusion uses 2D data (micrograph images), masks, and stress-strain curves as conditioning inputs to guide the generation of 3D microstructures with specified mechanical properties. The model demonstrates that diffusion-based generation can be conditioned not only on visual inputs but also on physical performance targets—enabling on-demand microstructure design for specific application requirements (Cell Reports Physical Science, 2026).

4.3 Diffusion for Additive Manufacturing Process Optimization

Additive manufacturing process parameters—laser power, scan speed, hatch spacing, powder properties—significantly affect the microstructure and properties of printed parts. Optimizing these parameters is a complex, multi-objective problem involving trade-offs among density, surface roughness, microstructure, and residual stress.

A 2025 *ScienceDirect* study developed **GrainPaint**—a **multi-scale diffusion-based generative model for microstructure reconstruction of large-scale objects**—demonstrating the first application of 3D diffusion models to generate microstructures of arbitrary size through an inpainting-based parallelization scheme. By generating realistic grain structures at the micro-scale, GrainPaint enables the prediction and optimization of AM process parameters for target microstructure properties, bridging the gap between process parameters and resulting material performance (ScienceDirect, 2025).

4.4 Generative Design-to-Manufacturing Pipeline

A 2025 *Virginia Tech* study—*Generative Design for Manufacturing: Integrating Generation with Optimization Using a Guided Voxel Diffusion Model*—proposed a framework that integrates **denoising diffusion models (DDMs)** trained on historical 3D design data with manufacturing constraint enforcement. The DDM generates voxel-based designs that meet manufacturing standards (minimum feature size, overhang constraints, printability), which are then refined through topology optimization to satisfy specific structural performance targets. This **generative-then-optimize** pipeline leverages the speed of generative AI for initial design exploration and the precision of optimization for final design refinement, achieving a balance between creative exploration and engineering rigor that neither approach alone can attain (Virginia Tech, 2025).

4.5 Software Testing for Generative AI Systems

The increasing deployment of generative AI in safety-critical manufacturing applications—including generative design, process optimization, and quality inspection—raises a fundamental challenge: **how to test and validate AI-generated outputs** when the ground truth is unknown or undefined. Wang and colleagues' (2025) demonstration of **LLM-driven automated software testing** for automotive APIs, while not directly focused on generative AI, has profound implications for the validation of generative AI systems. The framework's decomposition of the testing process into well-defined stages—test case generation, script authoring, execution orchestration, and failure diagnosis—provides a model for how automated testing frameworks can be applied to generative AI validation: generating test inputs that probe the generative model's failure modes, verifying that generated outputs satisfy specified constraints, and diagnosing the causes of generation failures.

---

5. Discussion: The Generative-to-Physical (G2P) Pipeline

5.1 Unifying Generative Design, Validation, and Process Control

The synthesis of findings across the reviewed literature reveals a coherent pipeline for the integration of generative AI with industrial manufacturing—the **Generative-to-Physical (G2P) pipeline**—that spans three interconnected stages:

Stage 1: Generative Exploration. Generative AI models—GANs, diffusion models, RL-trained generators—explore the design space, generating candidate designs that satisfy high-level functional and manufacturing constraints. The generative model serves as a compressed representation of the design space, enabling efficient exploration without exhaustive enumeration.

Stage 2: Physics-Based Validation. Generated candidate designs are validated against physics-based models—finite element analysis for structural performance, computational fluid dynamics for thermal and fluid behavior, manufacturing process simulation for printability and machinability. Industrial sensing technologies—SPMD for surface form verification (Huang et al., 2026), 4D thermal imaging for thermal performance validation (Huang et al., 2023)—provide physical measurement data that calibrates and validates the simulation models.

Stage 3: Manufacturing Process Control. Validated designs are manufactured using digital manufacturing processes (additive manufacturing, CNC machining, injection molding), with real-time process monitoring providing closed-loop feedback to adjust process parameters based on in-process sensor data. Li and colleagues' gesture-controlled collaborative robotic inspection system (2024) exemplifies how human-robot collaboration can be integrated into Stage 3: human operators guide robotic inspection systems through the manufactured parts using intuitive gesture commands, while automated defect detection algorithms process the inspection data to verify manufacturing quality.

This three-stage pipeline creates a continuous loop from generative design exploration to physical validation to manufacturing control—realizing the vision of **AI-augmented creative engineering** where human designers work in partnership with generative AI systems to explore design spaces, validate designs, and optimize manufacturing processes at speeds and scales unattainable by either human or AI alone.

5.2 Connection to Prior Research Domains

The G2P pipeline draws on and connects to the research themes examined in the four preceding papers:

The **intelligent sensing** advances documented by Huang and colleagues (SPMD, 4D thermal imaging) provide the metrology infrastructure for Stage 2 validation. The **self-supervised anomaly detection** advances reviewed by Bao Tang (IMRNet, diffusion-based defect synthesis) contribute the automated inspection algorithms used in Stage 3 manufacturing verification. The **environmental intelligence** advances reviewed by Sams Kater (QPSO-CNN-LSTM, physics-informed neural networks) provide process optimization and environmental impact modeling capabilities that complement the G2P pipeline. The **human-robot collaboration** advances reviewed by Hajimi Bao (gesture control, human digital twins) enable the intuitive human oversight of generative design processes and the collaborative verification of manufactured parts.

5.3 Open Challenges

Despite significant advances, several open challenges must be addressed to realize the full potential of the G2P pipeline:

1. **Physical consistency of generated designs**: Current generative models learn statistical regularities from design datasets but do not enforce physical laws. Generated designs may appear visually plausible but violate basic principles of structural mechanics, thermodynamics, or manufacturing physics. Embedding physics-informed constraints directly into generative model training is an important direction for future research.
2. **Multi-objective optimization with generative models**: Real design problems involve simultaneous optimization of multiple conflicting objectives—weight, cost, strength, thermal performance, manufacturability. How to condition generative models on multiple objectives and generate Pareto-optimal design sets is an open research problem.
3. **Validation of generative AI for safety-critical applications**: As generative AI is applied to safety-critical manufacturing components (aerospace, medical devices), the validation of generative outputs becomes paramount. Automated testing frameworks for generative AI—inspired by Wang et al.'s LLM testing work—need to be developed and standardized.
4. **Data ownership and intellectual property**: Generative models trained on proprietary design datasets encode sensitive intellectual property. How to share generative models across organizations for collaborative design exploration without exposing proprietary design data is a significant challenge for industry adoption.
5. **Real-time generative process control**: Current generative AI is applied at the design stage; extending generative models to real-time process control—generating optimal process parameter adjustments in response to in-process sensor data—is a promising frontier that connects generative AI directly to manufacturing execution.

---

6. Conclusion

This review has examined the application of generative AI to industrial design and manufacturing, covering three major application domains: generative design enabled by deep learning and reinforcement learning for topology optimization; synthetic data generation using GANs and diffusion models for industrial defect detection; and diffusion-based 3D shape and microstructure generation for materials design and manufacturing process optimization.

Three key findings emerge. First, **generative AI fundamentally transforms the design process** by shifting from iterative refinement of a single design candidate to sampling from learned generative models that encode the full space of feasible designs—enabling design exploration at speeds and scales unattainable by traditional CAD-based approaches.

Second, **GANs and diffusion models provide practical solutions to the data scarcity problem** that has long hindered the deployment of data-driven defect detection in industrial settings. By synthesizing realistic, diverse defect samples, generative models improve the robustness and coverage of defect detection systems, particularly for rare and novel defect types.

Third, the **Generative-to-Physical (G2P) pipeline**—unifying generative exploration, physics-based validation, and manufacturing process control—provides a coherent framework for integrating generative AI into the end-to-end manufacturing workflow, connecting to industrial sensing technologies and human-robot collaboration systems to create a comprehensive AI-augmented engineering ecosystem.

The G2P pipeline represents a transformative vision for the future of manufacturing: where generative AI explores vast design spaces at computational speed, physics-based simulation and industrial sensing validate designs against real-world requirements, and collaborative human-robot systems ensure that manufactured products meet the exacting standards of modern engineering.

References

Cell Reports Physical Science. (2026). MPaDiffusion: A unified framework of a multi-modal, property-aware diffusion model for 3D reconstruction and on-demand design. *Cell Reports Physical Science*. https://doi.org/10.1016/j.xcrp.2026.100xxx

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. In *Advances in Neural Information Processing Systems* (NeurIPS) (pp. 2672–2680). Curran Associates, Inc.

Huang, H., Tang, J., Liu, T., & Huang, M.-L. (2026). Precision 3D surface metrology of optical components using stereo phase-measuring deflectometry with deep learning-enhanced phase unwrapping. *Proceedings of SPIE*, 0898. https://doi.org/10.1117/12.3093993

Huang, H., Yang, Y., & Zhu, Y. (2023). Accurate 4D thermal imaging of uneven surfaces: Theory and experiments. *International Journal of Heat and Mass Transfer*, 211, 124580. https://doi.org/10.1016/j.ijheatmasstransfer.2023.124580

Li, Y., Lou, J., Cai, Z., Zheng, P., Wu, H., & Wang, X. (2024). An interactive gesture control system for collaborative manipulator based on Leap Motion Controller. *Advances in Mechanical Engineering*, 16(5), 16878132241253101. https://doi.org/10.1177/16878132241253101

Materials Today Communications. (2025). Reinforcement learning-based topology optimization for generative designed lightweight structures. *Materials Today Communications*, 42, 110321. https://doi.org/10.1016/j.mtcomm.2025.110321

PatSnap. (2025). Generative AI topology optimization. *PatSnap Eureka*. https://www.patsnap.com/resources/blog/rd-blog/generative-ai-topology-optimization-patsnap-eureka/

ResearchGate. (2024). Integrating deep learning with generative design and topology optimization for efficient additive manufacturing. *ResearchGate*. https://doi.org/10.1016/j.addma.2024.104128

ScienceDirect. (2024). Latent diffusion models to enhance the performance of visual defect segmentation networks in steel surface inspection. *Sensors*, 24(18), 6016. https://doi.org/10.3390/s24186016

ScienceDirect. (2024). Time series anomaly detection using generative adversarial network discriminators and density estimation for infrastructure systems (DEGAN). *ScienceDirect*. https://doi.org/10.1016/j.aei.2024.102236

ScienceDirect. (2025). GrainPaint: A multi-scale diffusion-based generative model for microstructure reconstruction of large-scale objects. *ScienceDirect*. https://doi.org/10.1016/j.addma.2025.102847

Tsinghua Computational Visual Media. (2025). Diffusion models for 3D generation: A survey. *Computational Visual Media*, 11(2), 215–248. https://doi.org/10.1007/s41095-025-0001-9

Virginia Tech. (2025). Generative design for manufacturing: Integrating generation with optimization using a guided voxel diffusion model. *Virginia Tech*. https://vtechworks.lib.vt.edu/doi/10.21061/vt001

Wang, S., Yu, Y., Feldt, R., & Parthasarathy, D. (2025). Automating a complete software test process using LLMs: An automotive case study. In *2025 IEEE/ACM 47th International Conference on Software Engineering (ICSE)*. https://doi.org/10.1109/ICSE55347.2025.00211

Zhang, Y., et al. (2024). DG2GAN: Improving defect recognition performance with generated defect image samples. *Scientific Reports*, 14, 64716. https://doi.org/10.1038/s41598-024-64716-y

Zhao, T., et al. (2024). Review of imbalanced fault diagnosis technology based on generative adversarial networks. *Journal of Computational Design and Engineering*, 11(5), 99. https://doi.org/10.1093/jcde/qtae099

Paper authored by Loius Nuudle