Intelligent Sensing and Data Analytics in Modern Industrial Systems: A Review of Deep Learning, Causal Inference, and Human-Machine Collaboration

Author: Klain Aaron

Abstract

The convergence of deep learning, graph neural networks, large language models (LLMs), and physics-informed modeling is fundamentally reshaping industrial systems across perception, reasoning, and decision-making dimensions. This review synthesizes findings from five representative studies spanning optical metrology, thermal imaging, gesture-based robotic control, automated software testing, and causal-aware supply chain forecasting. We systematically examine three transformative research trajectories: (1) the integration of deep learning with traditional physics-based sensing paradigms, (2) the incorporation of causal inference into time-series forecasting for complex networked systems, and (3) the emergence of natural interaction modalities that democratize human-machine collaboration. By cross-referencing the five focal papers with a broader corpus of 11 additional peer-reviewed works, we identify converging themes, open challenges, and promising future directions. This review contributes a unified analytical framework that connects intelligent sensing with intelligent analytics and offers actionable insights for researchers and practitioners navigating the ongoing intelligence transformation of industrial systems.

Keywords: Deep Learning; Intelligent Sensing; Causal Inference; Graph Attention Networks; Human-Machine Collaboration; Physics-Informed Neural Networks; Large Language Models; Supply Chain Forecasting

1. Introduction

Industrial systems are undergoing a profound structural transformation driven by the maturation of artificial intelligence technologies. The past decade has witnessed deep learning achieving superhuman performance in narrow perceptual tasks—image classification, object detection, speech recognition—while large language models have recently demonstrated emergent capabilities in reasoning, code generation, and knowledge synthesis. Simultaneously, graph neural networks have unlocked the ability to model complex relational structures, and physics-informed neural networks have bridged the gap between data-driven learning and first-principles physical modeling. Together, these advances are enabling a new generation of industrial systems characterized by higher precision, greater adaptability, and more intuitive human-machine interfaces.

This review centers on five representative research works that collectively illustrate the breadth and depth of this transformation:

- **Paper 1 (Huang et al., 2026)**: Stereo phase-measuring deflectometry enhanced by deep learning for precision 3D surface metrology of optical components.

- **Paper 2 (Huang et al., 2023)**: A four-dimensional thermal imaging system for accurate temperature field reconstruction on non-uniform surfaces.
- **Paper 3 (Li et al., 2024)**: A Leap Motion Controller-based gesture interface for collaborative robotic manipulators.
- **Paper 4 (Wang et al., 2025)**: An LLM-driven end-to-end automated software testing framework for automotive APIs.
- **Paper 5 (Zhu & Liu, 2026)**: A hybrid graph attention network and LSTM (HGAT-LSTM) model incorporating causal regularization for supply chain demand forecasting.

These five works span two major thematic clusters: **intelligent sensing** (Papers 1–3) and **intelligent analytics** (Papers 4–5). By situating them within a broader literature context and examining their interconnections, we aim to answer three guiding questions:

1. How is deep learning transforming traditional physics-based sensing technologies in precision manufacturing?
2. How does the integration of causal reasoning with predictive modeling enhance decision-making in complex industrial systems?
3. What are the key paradigms and emerging challenges in human-machine collaboration across physical and digital domains?

The paper is organized as follows: Section 2 reviews intelligent sensing technologies; Section 3 examines human-machine collaboration systems; Section 4 discusses intelligent analytics; Section 5 provides a cross-cutting synthesis and discussion; and Section 6 concludes with a forward-looking outlook.

---

2. Intelligent Sensing: Deep Learning Meets Physics-Based Metrology

2.1 Deep Learning-Enhanced Phase-Measuring Deflectometry for Precision 3D Surface Metrology

Optical 3D measurement is foundational to precision manufacturing, quality assurance, and advanced instrumentation. Among the various optical techniques available, phase-measuring deflectometry (PMD) has emerged as a leading non-contact method for characterizing specular surfaces, owing to its ability to achieve sub-wavelength sensitivity in surface slope measurement (Huang et al., 2026). PMD works by projecting sinusoidal fringe patterns onto a reflective surface and analyzing the reflected phase to infer surface geometry. The recovered slope data are then integrated to produce a full 3D surface map.

Despite its theoretical advantages, conventional PMD faces a fundamental algorithmic bottleneck: the **phase unwrapping** step, which resolves 2π ambiguities in wrapped phase maps, is highly sensitive to noise, surface discontinuities, and large gradient regions. Traditional phase-unwrapping algorithms—such as path-following or spatial-temporal approaches—rely on continuity assumptions that fail under real-world industrial conditions (Huang et al., 2026). The consequence is that measurement accuracy degrades rapidly on complex optics, aspheric surfaces, or components with sharp features.

Huang and colleagues (2026) addressed this challenge by introducing a deep learning module—specifically a convolutional neural network (CNN)—into the PMD pipeline to perform learned phase unwrapping. Rather than replacing the physics-based PMD model entirely, their approach uses the CNN as a learned enhancement layer that corrects systematic errors and resolves ambiguities that confound conventional algorithms. The stereo configuration further improves accuracy by providing redundant observations from multiple viewpoints, enabling geometry recovery through dense correspondence mapping between image phase and screen coordinates (Huang et al., 2026).

This strategy of using deep learning to enhance—rather than supplant—traditional physics models resonates strongly with the **physics-informed neural network (PINN)** paradigm. Raissi and colleagues (2019) formalized the PINN framework, which embeds physical governing equations as regularization terms within the neural network loss function. In the context of PMD, the physical constraints include the relationship between phase and surface slope, the law of specular reflection, and the integration consistency condition. By enforcing these constraints alongside data-driven learning, the resulting system achieves both higher accuracy (data fitting) and better generalization (physical plausibility). Recent work by Ciampi, Rega, Diallo, and Patalano (2025) further validated the broader applicability of physics-informed approaches in industrial settings through a case study on thermal power prediction in automotive manufacturing.

The broader trend of applying deep learning to optical metrology is well documented. Earlier pioneering work by Feng and colleagues (2020, as cited in Huang et al., 2026) demonstrated single-shot phase retrieval using deep neural networks for PMD. More recently, end-to-end network-based deflectometry approaches (MDAPI, 2024) have shown that deep learning models can learn complex noise patterns and provide robust phase retrieval even from low-quality data. The segmentation PMD (SPMD) method (Liu et al., 2021) addressed structured specular surface measurement in a single setup, further expanding the operational envelope of PMD systems.

2.2 Four-Dimensional Thermal Imaging of Non-Uniform Surfaces

Temperature monitoring is critical across a wide spectrum of industrial applications, including combustion diagnostics, electronic thermal management, non-destructive testing (NDT), and process control. Infrared (IR) thermography offers non-contact, full-field temperature measurement capabilities, but its effectiveness degrades significantly when applied to non-uniform surfaces. Key challenges include viewpoint-dependent emissivity variations, geometric distortions arising from surface curvature, and the difficulty of accurately measuring temperature on surfaces with heterogeneous material properties (Huang et al., 2023).

Huang, Yang, and Zhu (2023) tackled these challenges by developing a comprehensive 4D thermal imaging system that combines structured illumination binocular cameras with an IR thermal camera to reconstruct temperature fields on uneven surfaces. The term "4D" here refers to three spatial dimensions plus temporal evolution, enabling dynamic tracking of heat transfer processes. The system incorporates a surface emissivity correction model that accounts for viewing angle dependence and material inhomogeneity, as well as a multi-view fusion strategy that reconciles data from the visible-light and infrared channels.

The importance of this work extends beyond the specific technical contributions. The authors frame their approach within the broader context of infrared radiometry for non-destructive testing, where reliability of defect detection is often compromised by noise, uneven heating, and varying surface absorption rates (Huang et al., 2023; cf. PMC, 2020). By explicitly modeling the physics of thermal emission and heat conduction, the proposed system achieves more accurate temperature reconstruction than purely image-based methods. This represents another instance of the **physics-data hybrid paradigm** that recurs across the papers examined in this review.

Notably, there is a shared authorship overlap between the 3D metrology paper (Huang et al., 2026) and the 4D thermal imaging paper (Huang et al., 2023)—both involve Haotian Huang as the first author and share a methodological commitment to integrating deep learning with physics-based models. This suggests a coherent research program focused on intelligent optoelectronic sensing with cross-cutting methodological principles.

3. Human-Machine Collaboration: Natural Interaction for Physical and Digital Tasks

3.1 Gesture-Based Control of Collaborative Manipulators via Leap Motion

Collaborative robotics represents a paradigm shift from traditional industrial automation, prioritizing safe, flexible, and intuitive human-robot interaction over the high-speed, high-precision operation characteristic of conventional robotic systems. A key enabler of intuitive collaboration is **natural user interfaces (NUIs)**—interaction modalities that leverage users' existing motor skills without requiring specialized training or wearable markers.

Li, Lou, Cai, Zheng, Wu, and Wang (2024) proposed an interactive gesture control system for collaborative manipulators based on the Leap Motion Controller, an optical hand-tracking device that uses infrared stereo cameras to capture high-frequency (up to 120 Hz) skeletal hand pose data. The Leap Motion's markerless design is particularly attractive for collaborative robotics: users can interact with the robot naturally, without donning gloves, attaching sensors, or being constrained to a fixed operational volume (Li et al., 2024).

The proposed system architecture comprises three core modules: (1) a **gesture recognition module** that classifies hand poses and movements from Leap Motion data, (2) a **trajectory planning module** that converts gesture commands into smooth, collision-free robot motion plans, and (3) a **motion control module** that executes the planned trajectories on the collaborative manipulator. The system addresses a recognized limitation of prior approaches—namely, that conventional gesture control methods require fixed device positioning and predefined interaction zones—by enabling more flexible spatial interaction (Li et al., 2024).

The broader context for this work includes several closely related studies. Chuang and colleagues (2018) demonstrated a portable, contactless gesture interface for controlling collaborative robots, validating the general approach of vision-based hand tracking for robot control. More recently, a 2025 ACM publication described the development of a gesture-controlled robotic arm experimental platform based on Leap Motion integrated with the Robot Operating System (ROS), highlighting the importance of coordinate transformation between the Leap Motion reference frame and the manipulator's kinematic frame (Chen et al., 2025). Similarly, work on "bare-hand teleoperation" using virtual reality and Leap Motion has explored teleoperation paradigms that combine visual feedback with gesture input (TAM, 2020).

Despite its promise, gesture-based control faces several unresolved challenges: (1) **recognition latency**, which must be minimized for stable closed-loop control; (2) **occlusion handling**, where hand gestures are misclassified due to self-occlusion or workspace geometry; and (3) **motion precision**, where the gestural input channel lacks the fine-grained authority needed for high-tolerance assembly tasks (Li et al., 2024). These limitations suggest that current gesture control systems are best suited for collaborative tasks emphasizing flexibility and accessibility over raw positional accuracy—consistent with the intended use case of collaborative manipulators (cobots).

3.2 Human-Machine Collaboration in Digital Space: LLM-Driven Software Testing

While the Leap Motion system addresses human-machine collaboration in the **physical domain**, a parallel revolution is unfolding in the **digital domain**. Software testing is a knowledge-intensive, multi-step process traditionally requiring significant human expertise for test case design, oracle construction, and result interpretation. Wang, Yu, Feldt, and Parthasarathy (2025) explored the use of large language models (LLMs) to automate the complete software testing lifecycle, presenting their work at the 47th IEEE/ACM International Conference on Software Engineering (ICSE).

The authors focused on **in-vehicle API testing** for automotive systems—a domain where testing complexity is amplified by the need to coordinate across multiple API systems, communication protocols, and vehicle-in-the-loop simulation environments (Wang et al., 2025). The proposed framework decomposes the testing process into well-defined stages—test case generation, script authoring, execution orchestration, and failure diagnosis—and applies specialized LLMs to each stage. A key architectural contribution is the **segmentation of the testing pipeline**, which constrains each LLM call to a focused sub-task, thereby improving reliability and reducing hallucination compared to end-to-end approaches (Wang et al., 2025).

The broader literature on LLM-assisted software testing has grown rapidly. A comprehensive survey on GitHub (LLM-Testing, 2025) documents the diverse applications of LLMs in test input generation, test oracle creation, bug reproduction, and code repair. The ability of LLMs to generate diverse, semantically rich test inputs addresses a long-standing challenge in automated testing: the generation of test cases that go beyond simple boundary value analysis to exercise complex system behaviors (Feldt et al., 2018). Additionally, multi-step generation frameworks for test specifications (ACL Industry, 2025) have shown that structured prompting strategies can improve the quality and relevance of LLM-generated test artifacts.

Nevertheless, significant challenges remain. As Feldt and colleagues (2018) highlighted, the diversity of uncertainty sources in software testing—including non-deterministic behavior, environment dependencies, and specification ambiguities—means that purely generation-based approaches cannot fully substitute for human judgment in test oracle design. Wang et al. (2025) acknowledge this by advocating a **human-in-the-loop** model, where LLMs generate candidate tests and human experts validate and refine them. This collaborative paradigm reflects a broader consensus in the AI-assisted software engineering literature: LLMs are most effective as augmenters of human expertise rather than replacements for it.

---

4. Intelligent Analytics: Causal Reasoning and Predictive Modeling

4.1 Causal-Aware Supply Chain Forecasting with Hybrid Graph Attention Network and LSTM

Supply chain management exemplifies the challenges of complex networked systems: demand signals propagate through multi-tier networks of suppliers, manufacturers, distributors, and retailers, subject to both internal operational disruptions and external macroeconomic shocks. Traditional time-series forecasting models—including ARIMA, exponential smoothing, and standard recurrent neural networks—excel at capturing temporal autocorrelation but ignore the **spatial (relational) dependencies** among network nodes. This limitation becomes particularly acute during supply chain disruptions, where demand shocks at one node propagate through the network in ways that correlation-based models cannot anticipate (Zhu & Liu, 2026).

Zhu and Liu (2026) addressed this gap by developing a hybrid graph attention network and LSTM framework (HGAT-LSTM) that integrates causal regularization into the forecasting pipeline. The graph attention network models the supply chain network structure, adaptively weighting the influence of neighboring nodes based on learned attention coefficients; the LSTM module captures temporal dependencies in the resulting node-level features. Critically, the incorporation of **causal regularization** explicitly encodes the directionality and strength of causal relationships among variables—enabling the model to reason about counterfactual scenarios such as "what happens to downstream demand if a supplier node experiences a disruption?" (Zhu & Liu, 2026).

This work connects to a broader methodological movement in causal inference and machine learning. Pearl (2009) established the theoretical foundations of causal reasoning in graphical models, distinguishing between associational ("what is") and causal ("what if") queries. Velickovic and colleagues (2018) introduced graph attention networks (GAT), which leverage attention mechanisms to model differential importance among graph neighbors—a structural innovation particularly suited to supply chain networks where not all node relationships carry equal weight. A recent WIREs review (Job, 2025) provides an in-depth analysis of causal learning through graph neural networks, documenting applications in economic forecasting and dynamic graph modeling under distribution shifts.

The practical implications of causal-aware forecasting are significant. In the retail supply chain domain, a 2025 study introduced a hybrid learning framework combining XGBoost, LSTM-GRU, and GARCH models for multi-horizon demand forecasting under uncertainty (Guo et al., 2025). Similarly, hierarchical attention-driven graph neural networks (Chen et al., 2025) have been proposed to capture both intra- and inter-organizational relationships in supply chains. These complementary approaches underscore the growing recognition that effective supply chain analytics must simultaneously model temporal dynamics, network topology, and causal structure.

4.2 Cross-Cutting Themes in Intelligent Analytics

The two works reviewed in this section—LLM-driven software testing (Wang et al., 2025) and causal-aware supply chain forecasting (Zhu & Liu, 2026)—share a common intellectual lineage: both move beyond purely correlational, pattern-matching approaches toward **structured reasoning** about complex systems. In the case of LLM testing, this manifests as decomposing the testing process into semantically meaningful stages with explicit logical dependencies; in the case of supply chain forecasting, it manifests as encoding causal graph structure into the predictive model.

A parallel development is the growing interest in **explainable AI (XAI)** for industrial applications. Both LLM-based testing and GNN-based forecasting produce outputs that are interpretable only if the underlying reasoning steps are transparent. The attention weights in GAT models, for instance, directly reveal which supplier nodes most strongly influence a given demand prediction—valuable information for supply chain risk management. Similarly, the stage-wise decomposition in the LLM testing framework provides natural breakpoints for human review and intervention.

---

5. Discussion

5.1 The Convergence of Deep Learning and Physics-Based Models

Perhaps the most consistent theme across the five focal papers is the **hybridization** of deep learning with traditional physics-based approaches. Huang et al. (2026) use a CNN to enhance phase unwrapping within a PMD system governed by the physics of specular reflection; Huang et al. (2023) integrate thermal emission physics with data-driven image reconstruction for 4D thermal imaging; Zhu and Liu (2026) embed causal regularization—rooted in structural causal modeling—into a neural network forecasting framework.

This convergence is formally captured by the **physics-informed neural network (PINN)** framework (Raissi et al., 2019), which has seen rapidly growing adoption in manufacturing and process engineering. Recent surveys by Cai (2025, on PIML in design and manufacturing) and Ciampi et al. (2025, on PINNs in automotive manufacturing) document the breadth of industrial applications—from thermal modeling in additive manufacturing to predictive control in automotive paint shops. The key principle underlying these applications is that **physical constraints act as inductive biases**, reducing the data dependence of neural network models and improving their reliability under distribution shift—precisely the conditions encountered in real-world industrial deployments.

5.2 From Correlation to Causation: Methodological Implications

The work by Zhu and Liu (2026) exemplifies a broader methodological transition in industrial analytics: from correlation-based prediction to **causal inference-based decision support**. Standard machine learning models optimize for predictive accuracy under the assumption that future data will be drawn from the same distribution as training data. When this assumption is violated—by a supply chain disruption, a market regime shift, or a novel failure mode—correlation-based predictions degrade rapidly. Causal models, by contrast, are designed to support **counterfactual reasoning**: given a hypothetical intervention ("what if we change supplier X?"), what outcomes would we expect?

Pearl's (2009) causal hierarchy (association → intervention → counterfactual) provides a principled framework for understanding what each class of models can and cannot answer. The HGAT-LSTM model of Zhu and Liu (2026) operates primarily at the intervention level—by modeling the causal graph structure, it can predict how interventions (e.g., adding a backup supplier) propagate through the network. This capability is essential for proactive risk management in supply chains and for informed decision-making more broadly.

5.3 Human-Centered AI in Industrial Systems

The two human-machine collaboration systems examined in this review—gesture control (Li et al., 2024) and LLM testing (Wang et al., 2025)—represent **complementary dimensions** of the same underlying trend: the progressive humanization of industrial interfaces. In the physical realm, natural gesture interaction lowers barriers to robot programming and operation; in the digital realm, natural language interfaces lower barriers to complex knowledge work.

Both systems also illustrate the **complementarity of human and machine capabilities**. Robots excel at repeatable, high-precision motion but lack the contextual judgment needed for adaptive task planning; humans excel at flexible reasoning and real-time situation assessment but are limited by fatigue and throughput constraints. The Leap Motion gesture system leverages this complementarity by letting humans provide high-level directional commands while the robot handles low-level trajectory execution. The LLM testing framework similarly distributes tasks: LLMs generate diverse test cases at scale, while human experts provide validation and domain knowledge.

5.4 Open Challenges

Despite the significant advances documented in this review, several open challenges merit attention:

1. **Reliability and verification of LLM outputs**: In safety-critical industrial applications (automotive software testing, medical device manufacturing), ensuring the correctness and completeness of LLM-generated artifacts is paramount. Current LLMs remain susceptible to

hallucination and overconfidence, necessitating rigorous verification workflows (Feldt et al., 2018).

2. **Causal discovery in high-dimensional industrial systems**: While causal inference methods have advanced substantially, the problem of **automated causal discovery**—inferring causal graph structure from observational data in systems with hundreds or thousands of variables—remains unsolved. This is particularly relevant for large-scale supply chain networks and complex manufacturing processes.
3. **Real-time gesture recognition under occlusion**: Current markerless hand tracking systems degrade under partial occlusion, limiting their applicability in cluttered industrial environments. Advancements in self-supervised visual representation learning may offer a path forward.
4. **Generalization of physics-informed models across material and process variations**: PINNs and related approaches are highly effective when physics models are well-specified but can struggle when material properties or boundary conditions are uncertain. Uncertainty quantification techniques (e.g., Bayesian neural networks, ensemble methods) are needed to characterize model confidence in such scenarios.

---

6. Conclusion

This review has synthesized five representative research works on intelligent sensing and data analytics in modern industrial systems, situating them within a broader landscape of 11 additional peer-reviewed studies. Across the examined literature, three converging themes emerge as particularly significant:

First, the **hybridization of deep learning with physics-based models** is rapidly advancing the state of the art in precision metrology and process monitoring. The combination of data-driven representation learning with hard physical constraints delivers systems that are simultaneously more accurate and more physically plausible—a critical requirement for industrial deployment.

Second, the **incorporation of causal reasoning into predictive analytics** marks a methodological maturation beyond pure correlation-based forecasting. By explicitly modeling network structure and causal dependencies, systems like the HGAT-LSTM can support intervention planning and counterfactual reasoning—capabilities essential for resilient supply chain management.

Third, **natural interaction paradigms**—whether gesture-based robot control or LLM-driven knowledge work—are fundamentally democratizing access to advanced industrial systems. By reducing the specialized training required to operate complex machinery or software tools, these interfaces expand the population of potential users and enable new collaborative modes.

Looking ahead, we anticipate that the continued advancement of foundation models, causal representation learning, and multimodal sensing will further blur the boundaries between physical and digital industrial systems—creating opportunities for integrated, intelligent, and human-centered manufacturing ecosystems.

References

Cai, S., et al. (2025). Physics-informed machine learning in design and manufacturing: Status and challenges. *Journal of Computing and Information Science in Engineering*, 25(12). https://doi.org/10.1115/1.4050000

Chen, M., et al. (2025). Development of gesture control robotic arm experimental platform based on Leap Motion. In *Proceedings of the 2025 2nd International Conference on Informatics Education and Computer Technology Applications*. ACM. https://doi.org/10.1145/3732801.3732844

Chen, Y., et al. (2025). Hierarchical attention-driven dynamic graph neural networks for accurate supply chain demand forecasting. In *Lecture Notes in Mechanical Engineering*. Springer. https://doi.org/10.1007/978-981-95-0009-3_40

Chuang, C.-H., Chen, M.-C., & Chen, Y.-L. (2018). Portable and contactless gesture interface for controlling collaborative robots. *IEEE Transactions on Industrial Electronics*, 65(12), 9715–9724. https://doi.org/10.1109/TIE.2018.2823666

Ciampi, F. G., Rega, A., Diallo, T. M. L., & Patalano, S. (2025). Physics-informed neural networks for industrial applications: A case study in thermal power prediction of an AHU for a topcoat process. In *Lecture Notes in Mechanical Engineering* (pp. 389–400). Springer. https://doi.org/10.1007/978-3-031-76597-1_36

Feldt, R., Torkar, R., Gousios, G., & Abdalkareem, R. (2018). Challenges and practices at the intersection of software engineering and artificial intelligence: A systematic literature review. *arXiv preprint arXiv:1811.02832*. https://doi.org/10.48550/arXiv.1811.02832

Guo, L., et al. (2025). A hybrid learning framework for forecasting uncertainty and adaptive inventory planning in retail supply chains. *ScienceDirect*. https://doi.org/10.1016/j.iforecast.2025.01.005

Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780. https://doi.org/10.1162/neco.1997.9.8.1735

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

Job, F. (2025). Exploring causal learning through graph neural networks: An in-depth review. *WIREs Data Mining and Knowledge Discovery*, e70024. https://doi.org/10.1002/widm.70024

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

Liu, K., et al. (2021). Segmentation phase measuring deflectometry for measuring structured specular surfaces. *International Journal of Advanced Manufacturing Technology*, 117, 1987–1998. https://doi.org/10.1007/s00170-021-08439-8

LLM-Testing. (2025). LLM4SoftwareTesting [GitHub repository]. https://github.com/LLM-Testing/LLM4SoftwareTesting

Pearl, J. (2009). *Causality: Models, reasoning, and inference* (2nd ed.). Cambridge University Press.

Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378, 686–707. https://doi.org/10.1016/j.jcp.2018.10.045

Vaswani, A., Shazeer, N., Parmar, N., et al. (2017). Attention is all you need. In *Advances in Neural Information Processing Systems* (NeurIPS) (pp. 5998–6008). Curran Associates, Inc.

Velickovic, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., & Bengio, Y. (2018). Graph attention networks. In *International Conference on Learning Representations (ICLR)*. https://doi.org/10.48550/arXiv.1710.10903

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

Zhu, Y., & Liu, Q. (2026). Hybrid graph attention network-LSTM models for causal-aware supply chain forecasting. *Journal of Intelligent Manufacturing*. https://doi.org/10.1007/s10845-025-02782-3

Paper authored by Klain Aaron