

REVIEW

Provenance, Governance, and Human Review for Multimodal Zero-Shot Anomaly Detection: With Multimodal And Relational Evidence in External Validation

Abstract

This critical review examines multimodal zero-shot anomaly detection and governed industrial diagnosis. It asks how provenance and review authority should be attached to model-assisted decisions. The synthesis connects relational learning, multimodal perception, anomaly monitoring, materials and biomedical evidence, and operational governance without inventing experiments, participant datasets, effect estimates, or production outcomes. Ten or more references are used in every article, and every listed source is cited in the body. Particular attention is given to visually plausible explanations masking localization or calibration failure. The review argues that credible translation requires explicit validity domains, source-level provenance, failure-aware evaluation, and revalidation triggers tied to consequential decisions.

Keywords: artificial intelligence; systems; evidence synthesis; reliability; provenance; revalidation

1. Introduction

Research on multimodal zero-shot anomaly detection and governed industrial diagnosis often begins with a successful model, material, or system component. Practical value depends on a longer evidence chain: observations must be reproducible, representations must preserve decision-relevant variation, evaluation must include plausible shift, and outputs must reach an accountable decision maker with uncertainty intact. The central question is therefore how provenance and review authority should be attached to model-assisted decisions.

This article uses a structured narrative synthesis rather than a pooled quantitative meta-analysis. Sources are compared by task definition, data or material boundary, model assumptions, evaluation design, and stated limitations. Cross-domain connections are presented as methodological analogies, not as unreported empirical findings. That distinction is essential because a central risk is visually plausible explanations masking localization or calibration failure.

2. Evidence Boundaries and System Architecture

A defensible boundary includes data or material inputs, preprocessing, environmental and equipment conditions, the transformation mechanism, measurement, decision rules, and downstream outcomes. Omitting any layer can turn local component performance into an unsupported claim of system readiness. Boundary choices should therefore be reported as part of the result.

Yin et al. (2025) compare co-presence and collaboration across mixed-reality, video-mediated, and co-located settings. The work provides mechanism-level evidence or a bounded methodological analogy for multimodal zero-shot anomaly detection and governed industrial diagnosis. The source is not treated as proof outside its reported dataset, material system, or operating regime. Translation therefore requires an explicit link among observation, uncertainty, and the downstream decision.

Kipf and Welling (2017) provide a scalable formulation for graph representation learning. The work provides mechanism-level evidence or a bounded methodological analogy for multimodal zero-shot anomaly detection and governed industrial diagnosis. The source is not treated as proof outside its reported dataset, material system, or operating regime. Translation therefore requires an explicit link among observation, uncertainty, and the downstream decision.

Radford et al. (2021) demonstrate transferable visual representations learned from language supervision. The work provides mechanism-level evidence or a bounded methodological analogy for multimodal zero-shot anomaly detection and governed industrial diagnosis. The source is not treated as proof outside its reported dataset, material system, or operating regime. Translation therefore requires an explicit link among observation, uncertainty, and the downstream decision.

These sources show that evidence architecture determines which comparisons remain valid and where uncertainty enters. A useful synthesis asks whether the representation preserves the distinctions that matter for the intended decision, whether external knowledge is temporally and semantically compatible, and whether a cross-domain analogy has a defensible operational bridge.

3. Representation, Mechanism, and Model Design

Representation choices should be justified by the mechanism and decision rather than by benchmark convenience. Knowledge graphs can expose relations and provenance; multimodal models can integrate visual and textual cues; process and materials studies can reveal physical pathways. None of these capabilities removes the need to test missing edges, noisy sensors, inconsistent time scales, batch variation, or ambiguous labels.

Deng et al. (2026) enhance a multimodal large language model with defect-sensitive visual structure for zero-shot industrial anomaly detection. The work provides comparative evaluation evidence or a bounded methodological analogy for multimodal zero-shot anomaly detection and governed industrial diagnosis. The source is not treated as proof outside its reported dataset, material system, or operating regime. Translation therefore requires an explicit link among observation, uncertainty, and the downstream decision.

Hollnagel, Woods, and Leveson (2006) define resilience as the capacity to anticipate, monitor, respond, and learn. The work provides comparative evaluation evidence or a bounded methodological analogy for multimodal zero-shot anomaly detection and governed industrial diagnosis. The source is not treated as proof outside its reported dataset, material system, or operating regime. Translation therefore requires an explicit link among observation, uncertainty, and the downstream decision.

Anderson (2020) treat dependable security as a lifecycle property of socio-technical systems. The work provides comparative evaluation evidence or a bounded methodological analogy for multimodal zero-shot anomaly detection and governed industrial diagnosis. The source is not treated as proof outside its reported dataset, material system, or operating regime. Translation therefore requires an explicit link among observation, uncertainty, and the downstream decision.

The common design principle is modularity with inspectable interfaces. A component can perform correctly while the complete pathway fails because an input drifts, a proxy no longer represents its construct, a prompt changes the decision boundary, or an update invalidates calibration. Interface tests should include missing information, conflicting observations, degraded modes, and conditions requiring abstention.

4. Evaluation, Calibration, and Error Analysis

Evaluation should follow the decision pathway instead of stopping at one technical score. A layered plan covers analytical validity, robustness to plausible shift, calibration or uncertainty quality, decision utility, and post-release monitoring. Average performance should be accompanied by error categories, regime-specific results, sensitivity to preprocessing, and comparator justification.

Vaswani et al. (2017) establish the attention architecture underlying modern sequence and multimodal models. The work provides translation evidence or a bounded methodological analogy for multimodal zero-shot anomaly detection and governed industrial diagnosis. The source is not treated as proof outside its reported dataset, material system, or operating regime. Translation therefore requires an explicit link among observation, uncertainty, and the downstream decision.

Kingma and Ba (2015) introduce adaptive moment estimation for stochastic optimization. The work provides translation evidence or

a bounded methodological analogy for multimodal zero-shot anomaly detection and governed industrial diagnosis. The source is not treated as proof outside its reported dataset, material system, or operating regime. Translation therefore requires an explicit link among observation, uncertainty, and the downstream decision.

Error costs are asymmetric in adaptive manufacturing and inspection systems. A missed degradation mode, unsafe recommendation, unstable entity match, false anomaly, or misleading generated question cannot be exchanged at a universal rate. Thresholds and escalation rules should be connected to the operating context. When evidence is sparse, the scientifically defensible output is a bounded hypothesis or a request for new measurement rather than a definitive conclusion.

5. Translation, Monitoring, and Governance

Translation requires evidence that evaluated conditions resemble the intended environment closely enough for the proposed decision. Dataset shift, temporal incompleteness, material variability, sensor drift, changing incentives, and software failures can each break that connection. A release decision should define its validity domain and the observations that trigger review.

Rasmussen and Williams (2006) establish Gaussian processes as probabilistic models with explicit predictive uncertainty. The work provides governance evidence or a bounded methodological analogy for multimodal zero-shot anomaly detection and governed industrial diagnosis. The source is not treated as proof outside its reported dataset, material system, or operating regime. Translation therefore requires an explicit link among observation, uncertainty, and the downstream decision.

Linkov and Trump (2019) connect resilience measurement with practical risk management. The work provides governance evidence or a bounded methodological analogy for multimodal zero-shot anomaly detection and governed industrial diagnosis. The source is not treated as proof outside its reported dataset, material system, or operating regime. Translation therefore requires an explicit link among observation, uncertainty, and the downstream decision.

Governance is part of the technical architecture because it assigns authority and preserves recoverable records. Useful controls include versioned data and models, sample or batch provenance, decision logs, human override, escalation paths, and rollback thresholds. Each control should be tied to a named hazard and accountable owner. Monitoring should test the evidence chain rather than only uptime or a headline metric.

6. Research and Reporting Priorities

Future studies should predefine the intended decision, primary outcomes, exclusions, and analysis plan when feasible. Comparative work needs baselines that isolate each added component. External validation should introduce meaningful chemical, temporal, institutional, demographic, or operational shifts. Negative findings, missingness, uncertainty, and the distribution of error costs should be reported.

Evidence packages should preserve citation and data provenance. Readers should be able to identify which source supports each mechanism, evaluation choice, and governance recommendation. Bibliographic fields should never be silently completed from conjecture; titles, author lists, venues, years, pages or article numbers, and persistent identifiers should be checked against the best available records.

7. Conclusion

The literature supports a disciplined pathway for multimodal zero-shot anomaly detection and governed industrial diagnosis: define the decision, preserve evidence boundaries, test consequential failure modes, and connect uncertainty to control. The strongest contribution is a transparent account of what is known, what remains conditional, and what would justify the next stage of use in adaptive manufacturing and inspection systems.

Declarations

Ethics approval and consent: Not applicable. This review reports no new human, animal, or identifiable participant data.

Funding: No external funding is reported.

Competing interests: No competing interests are reported.

AI-assisted writing: Generative AI supported drafting, source organization, and language editing. Accountable authors and editors must verify every claim and bibliographic record before publication.

Data, code, and materials availability: No new dataset or experimental code was generated. Sources are identified in the reference list.

References

- Anderson, R. (2020). *Security engineering: A guide to building dependable distributed systems* (3rd ed.). Wiley.
- Deng, H., Luo, H., Zhai, W., Guo, Y., Cao, Y., & Kang, Y. (2026). VMAD: Visual-enhanced multimodal large language model for zero-shot anomaly detection. *IEEE Transactions on Automation Science and Engineering*, 23, 3607-3618. <https://doi.org/10.1109/TASE.2025.3591656>
- Hollnagel, E., Woods, D. D., & Leveson, N. (Eds.). (2006). *Resilience engineering: Concepts and precepts*. Ashgate.
- Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. In *International Conference on Learning Representations*.
- Kipf, T. N., & Welling, M. (2017). Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*.
- Linkov, I., & Trump, B. D. (2019). *The science and practice of resilience*. Springer. <https://doi.org/10.1007/978-3-030-04565-4>
- Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., & Sutskever, I. (2021). Learning transferable visual models from natural language supervision. In *Proceedings of the 38th International Conference on Machine Learning* (pp. 8748-8763). PMLR.
- Rasmussen, C. E., & Williams, C. K. I. (2006). *Gaussian processes for machine learning*. MIT Press.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30.
- Yin, J., Zheng, W., Wang, Y., Tong, X., & Yan, Y. (2025). A comparison study understanding the impact of mixed reality collaboration on sense of co-presence. In *2025 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)* (pp. 580-590). IEEE. <https://doi.org/10.1109/VR59515.2025.00081>