

Physical structure and graph learning for manufacturing defect detection

Vladimir Shapovalov
The Polytechnic University of Turin

Abstract

Manufacturing inspection combines measurements taken on irregular geometries with decisions about conformity, rework, and further examination. Graph neural networks can represent these measurements, but geometric connectivity, physical consistency, and defect evidence are different forms of information. This review develops a framework for separating them. It examines graph construction, anomaly reconstruction, explicit physical residuals, and the consequences of smoothing for localized faults. Analytical examples quantify signal attenuation, the effect of defect prevalence on alarm precision, and the distinction between node and part decisions. A proposed evaluation protocol links these calculations to manufacturing inspection through controlled topology changes, grouped validation, sensor uncertainty, and independently confirmed defect labels.

1. From measurements to an inspection decision

An inspection system observes a manufactured object through a measurement procedure. A thickness map, solder-paste measurement, infrared image, and displacement field do not measure the same property, even when they share spatial coordinates. The first modeling task is therefore to specify the measurand, units, spatial support, acquisition conditions, and decision being supported. A low reconstruction error may indicate predictable measurements without demonstrating conformity to a dimensional tolerance or absence of an internal defect. The broad anomaly-detection literature describes multiple objectives and assumptions that must be selected explicitly for this purpose [1].

The patent publication WO2024205617A1 describes graphs of measurement locations, distance-derived relationships, and neural models for reconstruction or supervised classification [2]. This is a relevant architecture reference. Its comparison of reconstruction on nominal manufacturing observations supports a particular reconstruction claim; reconstruction improvement alone does not determine sensitivity to every defect mechanism. The present analysis treats the disclosed arrangement as one candidate pipeline whose decision performance requires separate validation.

Graph convolutional networks [3], neighborhood aggregation for inductive graph learning [4], graph attention [5], and analyses of graph representational power [6] provide distinct architectural foundations. Their original benchmarks do not establish manufacturing reliability. For an engineering comparison, the learning architecture must be connected to physical

measurement semantics and the cost of incorrect decisions. This connection is where a physically structured model can be useful and where unsupported generalizations should be avoided.

2. Geometry, physical similarity, and governing equations

Let node i carry a measurement vector x_i at location r_i . A graph G contains nodes, edges, and edge weights w_{ij} . A Euclidean neighborhood expresses spatial proximity. A surface-geodesic neighborhood expresses proximity along a part surface. A process graph might connect locations deposited consecutively or components sharing a thermal path. These graphs encode different hypotheses. Two locations close across an air gap need not exchange heat like two locations connected through metal; a nearest-neighbor graph can silently confuse these cases.

The physics-informed learning literature spans data constraints, constitutive structure, and explicit differential equations [7, 8]. This breadth does not make every geometric graph a PDE-constrained model. A useful distinction is between physical organization of inputs, a prior about correlations, and a residual derived from a balance law. Learned mesh simulation illustrates a graph representation associated with physical evolution [9]. Its prediction task and supervision remain different from labeling defective parts.

Frequency-compensated fluid PINNs use physical signal structure to motivate learned frequency and phase features [10]. Rate- and temperature-dependent plasticity PINNs require outputs consistent with the mechanical formulation [11]. Neither paper validates an inspection classifier. Their relevance is methodological: the representation should correspond to the quantity being predicted. A model that predicts displacement may require a constitutive map before its output can support a stress-based inspection criterion; a model that predicts temperatures needs a separate connection to a defect mechanism.

For graphs built from mixed measurements, normalize each channel using training-set statistics or an independently specified physical scale. Keep measurement uncertainty separate from feature normalization. A channel with small numerical magnitude can be highly informative, and a large-variance channel can be mostly sensor noise. Correlated channels should not be counted as independent evidence merely because they occupy different columns of a measurement matrix.

3. What smoothing does to a localized defect

For symmetric nonnegative weights, define the unnormalized graph Laplacian $L = D - W$, where D contains weighted degrees. A common smoothness measure has the exact identity below. The factor one-half assumes that the double sum includes both edge directions.

$$x^T L x = (1/2) \sum_{i,j} w_{ij} (x_i - x_j)^2.$$

Penalizing this expression favors similar neighboring values. That is reasonable for a smooth nominal thickness field, but a crack, missing deposit, or sharp dimensional discontinuity may produce precisely the difference being penalized. Geometric smoothness is therefore not

automatically defect preservation. Spectral bias in neural approximation [12] and Fourier-feature representations [13] provide a related warning: fitting dominant smooth patterns and resolving a small high-frequency deviation are separate requirements.

Consider two nodes whose nominal values are zero and whose defect signal is $x = (a, o)$. A smoothing step P with diagonal entries $1 - \alpha$ and off-diagonal entries α gives the following exact result for zero to one-half mixing strength.

$$Px = ((1 - \alpha)a, \alpha a).$$

$$(Px)_1 - (Px)_2 = (1 - 2\alpha)a.$$

At $\alpha = 1/2$, the localized contrast disappears although the mean signal remains. If a downstream detector operates on local contrast, it can miss the defect after averaging. Conversely, a detector based on reconstruction residual may amplify evidence if the nominal model reconstructs a smooth baseline. These are opposite effects of where smoothing enters the pipeline. The algebraic example isolates the effect of smoothing on localized contrast.

Repeated propagation can enlarge the spatial footprint of a fault and alter localization metrics even when a part-level alarm remains correct. A study should therefore report both localization and part decisions and compare the raw-measurement residual with the latent-space score. Periodic activations [14] or graph attention cannot be assumed to solve this problem; their value depends on the learned map, training examples, and retained measurement resolution.

4. Reconstruction, classification, and physical residuals

An autoencoder minimizes a discrepancy between x and its reconstruction. A classifier estimates a label-dependent score. A physics model predicts an admissible field or its evolution. These objectives can share an encoder, but they answer different questions. Variational graph autoencoders [15] are an architectural reference whose original graph-link reconstruction should not be confused with reconstructing thickness or temperature channels. Industrial anomaly methods based on nominal feature memories, such as PatchCore, provide another relevant family of baselines [16].

Suppose a defect itself is easy to reconstruct because the model has sufficient capacity or contaminated nominal training data. Its reconstruction error can be small. Conversely, an acceptable but rare operating condition can be difficult to reconstruct and receive a large score. An anomaly score therefore measures departure from the learned representation of nominal behavior, not a direct posterior probability of physical failure. Supervised labels reduce this ambiguity only for the label definition and population that generated them.

An explicit physical residual may add useful evidence if its inputs and observation operator are justified. For transient temperature, a heat-balance residual requires temporal information, conductivity, sources, and boundary conditions. A surface snapshot alone generally does not supply these quantities. Work reconstructing thermal fields from partial measurements provides a relevant process-modeling example [17]; temperature and melt-pool modeling provides another

[18]. Neither allows every measured thermal irregularity to be labeled porosity without an additional validated relation.

Thermodynamics-based constitutive networks [19], solid-mechanics PINNs [20], and stress-oriented mechanics formulations [21] illustrate stronger physical structure than geometric proximity. Introducing such structure into inspection can improve interpretability but also introduces model discrepancy. A nominal constitutive law may fail exactly where damage changes the mechanism. Penalizing all departures from that law too strongly can suppress the evidence the inspection system needs to retain.

5. A hybrid objective and its failure modes

A proposed hybrid model can combine nominal reconstruction, labeled classification, and physically interpretable residuals. Let m_i identify observed channels and let Σ_i describe their measurement covariance. One possible objective is the following, with terms evaluated only where their assumptions are meaningful.

$$J = \lambda_{\text{rec}} \Sigma_i e_i^T \Sigma_i^{-1} e_i + \lambda_{\text{cls}} J_{\text{classification}} + \lambda_{\text{phys}} J_{\text{physics}}.$$

Here e_i contains reconstruction errors only for channels selected by m_i , and Σ_i is restricted to that observed subspace. Zeroing a missing channel is not equivalent to observing zero. Classification loss applies to labeled samples. Physical residuals may apply to latent fields rather than raw measurements. These distinctions should be specified before selecting scalar weights. Gradient imbalance and optimization difficulties documented for PINNs [22, 23, 24] explain why minimizing a composite objective does not ensure that each engineering requirement is satisfied.

Adaptive loss balancing [25] and self-adaptive weighting [26] are candidate training approaches, not substitutes for independent metrics. A reduced reconstruction term can coexist with a worse missed-defect rate. Report each constituent loss and evaluate classification at a threshold chosen without test labels. Sampling strategies [27, 28] likewise change which locations receive training emphasis. If rare defect labels are oversampled, correct the interpretation of calibrated probabilities and retain the true deployment prevalence for decision analysis.

For a physical residual evaluated on a graph, distinguish a graph difference from a consistent spatial derivative approximation. Anisotropic spacing, varying sampling density, and boundary truncation affect the operator. Weak or variational residuals [29, 30] may offer useful alternatives when pointwise derivatives are poorly supported, but they still require an integration rule and meaningful test functions. More residual evaluations cannot repair an incorrectly specified measurement operator.

6. Alarm precision depends on prevalence

Let π be the probability that an inspected part is defective, t the true-positive rate, and f the false-positive rate at a fixed threshold. Bayes' rule gives the fraction of alarms expected to correspond to defects when these rates describe the deployment population.

$$P(\text{defect} \mid \text{alarm}) = \pi t / [\pi t + (1 - \pi)f].$$

For an illustrative prevalence of 0.01, sensitivity 0.95, and false-positive rate 0.05, alarm precision is about 0.161. Among 10,000 parts, the expected counts are 95 true alarms and 495 false alarms. The counts follow from the assumed prevalence and operating rates. Keeping sensitivity fixed while reducing the false-positive rate to 0.005 raises precision to about 0.657. Reporting sensitivity alone conceals the inspection workload created by false alarms.

Node and part thresholds introduce a second issue. If a part is flagged whenever any of n independent nominal node tests alarms with probability q , the part-level false-alarm probability is $1 - (1 - q)^n$. With $q = 0.001$ and $n = 1,000$, this is approximately 0.632. Independence is unlikely on a physical graph, so this value is an illustrative limiting calculation. It demonstrates why changing sensor density can change a maximum-score alarm even when per-node performance appears stable.

Calibration research [31] motivates checking probability reliability on held-out parts. Deep ensembles [32] and Bayesian PINNs [33] provide different uncertainty mechanisms; neither by itself ensures reliable defect probabilities under a new production process. Separate uncertainty about a field estimate from uncertainty about the mapping from that field to a label. The latter can dominate when confirmed defects are rare or annotations disagree.

7. An engineering evaluation that isolates physical information

The proposed benchmark uses whole manufactured parts as independent units. Each part has measurement locations, acquisition metadata, nominal operating conditions, and an independently obtained conformity or defect label. Group parts by production batch before splitting. Repeated scans, adjacent image crops, and nodes from the same part remain together. Labels used to tune the operating threshold belong to validation data, and the final test set is evaluated once under the frozen procedure.

Comparison	Controlled change	Engineering interpretation
Independent baseline versus geometric graph	Add only neighborhood information	Value of spatial relationships
Geometric versus process-informed graph	Change justified edge semantics	Value of the proposed physical pathway
Graph model versus graph plus residual	Add an explicit balance or admissibility term	Incremental value of governing constraints

Comparison	Controlled change	Engineering interpretation
Nominal versus shifted acquisition	Change density, calibration, or missingness	Robustness of the complete inspection procedure

Hold channel preprocessing, training cases, tuning budget, and access to labels constant across variants. Compare computational cost from measurement ingestion through the final alarm, including graph construction. A graph-free baseline should receive the same coordinates if the scientific question is whether relational processing helps. Otherwise, an apparent graph advantage may merely reflect access to geometry withheld from the baseline.

PINN benchmark designs [34] and PDEBench [35] offer useful examples of organizing model comparisons, although their PDE tasks do not define defect-detection acceptance criteria. Computational comparisons with finite elements [36] further motivate reporting the task and budget on which a speed or accuracy claim rests. Inspection baselines should include a simple tolerance rule and a conventional statistical detector, alongside the more flexible learned models.

8. Measurement resolution and independently confirmed labels

The inspection instrument limits what a learned model can identify. Suppose a thickness instrument reports the average over a footprint of width h , and a narrow rectangular depression of depth d occupies width a within that footprint. In a one-dimensional idealization with uniform sensitivity and $a \ll h$, the measured depression is da/h . A defect occupying one-tenth of the footprint therefore contributes one-tenth of its true depth to that measurement. This deterministic attenuation occurs before graph construction or learning. A more expressive encoder cannot recover the depth uniquely without additional observations or a justified prior on defect width.

Spatial registration also matters. Moving the measurement footprint changes the overlap with the same defect, potentially changing the score without any change in the part. A proposed robustness study should retain the acquisition kernel or footprint specification and vary registration within the instrument's documented uncertainty. Report performance against defect size relative to measurement support, rather than using only a global average over all defects. This allows an engineer to distinguish an algorithmic limitation from insufficient physical resolution.

Ground truth should describe how a defect was confirmed, its spatial uncertainty, and whether the inspection method can observe it. Destructive sectioning, an independently calibrated instrument, or a subsequent functional test may provide different labels. A functional failure need not localize its initiating defect, and a visible discontinuity need not violate the functional acceptance criterion. These label types should have separate endpoints rather than being silently merged into one binary variable.

If several assessors label the same part, retain disagreement rather than replacing it with a seemingly exact consensus without documentation. Evaluate borderline cases separately and report how changing the acceptance threshold affects the conclusion. A model can appear poorly calibrated when the recorded label is itself inconsistent, yet label uncertainty must not be used to excuse systematic missed defects. The appropriate response is to quantify both sources of uncertainty and obtain stronger independent evidence for the decision-critical cases.

Finally, an inspection pipeline can abstain and request an additional measurement when its acquisition conditions fall outside the validated range. Evaluate this as a complete policy: report the fraction of parts referred, the additional inspection cost, and the residual error rate among automatically decided parts. Excluding referred cases from the denominator without reporting their frequency would overstate practical performance. This accounting connects uncertainty assessment to a concrete engineering action.

9. Perturbation tests and interpretation

Perturb graph edges while preserving approximate node degree to test sensitivity to topology, then separately perturb sensor locations within their calibration uncertainty. These interventions answer different questions: the first challenges the relational hypothesis, while the second challenges measurement robustness. Removing sensors tests information loss, whereas resampling collocation points tests optimization. Sensor-placement methods [37] help formalize the former and should not be treated as interchangeable with residual sampling.

Test localized defects of different spatial extent, acceptable gradients near boundaries, global offsets, and process-condition changes. A global sensor drift should be distinguishable from a local defect only when the observations or contextual measurements support that distinction. Multifidelity learning [38] can incorporate simulated scenarios, but synthetic labels require a documented transfer assumption. A simulation that lacks the real defect mechanism cannot establish the detector's sensitivity to that mechanism.

The strongest conclusion from such a study would identify which physical information improves a stated inspection decision and under which conditions. It would report missed defects, false alarms, localization error, uncertainty, and total inspection cost, rather than treating a lower training loss as sufficient evidence. Graph learning is valuable when its relationships correspond to the part, process, and measurement system. Its engineering credibility comes from demonstrating that correspondence and testing the decision that follows from it.

References

- [1] Guansong Pang, Chunhua Shen, Longbing Cao, Anton van den Hengel. Deep Learning for Anomaly Detection: A Review. arXiv:2007.02500 (2020). arXiv.
- [2] Pratik Kakkar, Bruno Paes Leao, Julian Busch, Shinjan Ghosh. Applying physics to a neural network model for detection of manufacturing defects. WO2024205617A1. PCT/US2023/030943. Published 3 October 2024 (2024). Source.

- [3] Thomas N. Kipf, Max Welling. Semi-Supervised Classification with Graph Convolutional Networks. arXiv:1609.02907 (2016). arXiv.
- [4] William L. Hamilton, Rex Ying, Jure Leskovec. Inductive Representation Learning on Large Graphs. arXiv:1706.02216 (2017). arXiv.
- [5] Petar Veličković, Guillem Cucurull, Arantxa Casanova, et al.. Graph Attention Networks. arXiv:1710.10903 (2017). arXiv.
- [6] Keyulu Xu, Weihua Hu, Jure Leskovec, Stefanie Jegelka. How Powerful are Graph Neural Networks?. arXiv:1810.00826 (2018). arXiv.
- [7] George Em Karniadakis, Ioannis G. Kevrekidis, Lu Lu, et al.. Physics-informed machine learning. *Nature Reviews Physics* 3, 422-440 (2021). DOI.
- [8] Maziar Raissi, Paris Perdikaris, George Em Karniadakis. 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 (2019). DOI.
- [9] Tobias Pfaff, Meire Fortunato, Alvaro Sanchez-Gonzalez, Peter W. Battaglia. Learning Mesh-Based Simulation with Graph Networks. arXiv:2010.03409 (2020). arXiv.
- [10] Tongtao Zhang, Biswadip Dey, Pratik Kakkar, Arindam Dasgupta, Amit Chakraborty. Frequency-compensated PINNs for Fluid-dynamic Design Problems. ML4Eng Workshop, NeurIPS 2020. arXiv:2011.01456 (2020). arXiv.
- [11] Rajat Arora, Pratik Kakkar, Amit Chakraborty, Biswadip Dey. Physics-informed neural networks for modeling rate- and temperature-dependent plasticity. Machine Learning and the Physical Sciences Workshop, NeurIPS 2022. arXiv:2201.08363v3 (2022). arXiv.
- [12] Nasim Rahaman, Aristide Baratin, Devansh Arpit, et al.. On the Spectral Bias of Neural Networks. arXiv:1806.08734 (2018). arXiv.
- [13] Matthew Tancik, Pratul P. Srinivasan, Ben Mildenhall, et al.. Fourier Features Let Networks Learn High Frequency Functions in Low Dimensional Domains. arXiv:2006.10739 (2020). arXiv.
- [14] Vincent Sitzmann, Julien N. P. Martel, Alexander W. Bergman, et al.. Implicit Neural Representations with Periodic Activation Functions. arXiv:2006.09661 (2020). arXiv.
- [15] Thomas N. Kipf, Max Welling. Variational Graph Auto-Encoders. arXiv:1611.07308 (2016). arXiv.
- [16] Karsten Roth, Latha Pemula, Joaquin Zepeda, et al.. Towards Total Recall in Industrial Anomaly Detection. arXiv:2106.08265 (2021). arXiv.
- [17] Shuheng Liao, Tianju Xue, Jihoon Jeong, et al.. Hybrid thermal modeling of additive manufacturing processes using physics-informed neural networks for temperature prediction and parameter identification. arXiv:2206.07756 (2022). arXiv.
- [18] Qiming Zhu, Zeliang Liu, Jinhui Yan. Machine learning for metal additive manufacturing: Predicting temperature and melt pool fluid dynamics using physics-informed neural networks. arXiv:2008.13547 (2020). arXiv.
- [19] Filippo Masi, Ioannis Stefanou, Paolo Vannucci, Victor Maffi-Berthier. Thermodynamics-based Artificial Neural Networks for constitutive modeling. arXiv:2005.12183 (2020). arXiv.
- [20] Ehsan Haghighat, Maziar Raissi, Adrian Moure, et al.. A deep learning framework for solution and discovery in solid mechanics. arXiv:2003.02751 (2020). arXiv.
- [21] R. Sharma, Y. B. Guo. Thermal-Mechanical Physics Informed Deep Learning For Fast Prediction of Thermal Stress Evolution in Laser Metal Deposition. arXiv:2412.18786 (2024). arXiv.
- [22] Sifan Wang, Yujun Teng, Paris Perdikaris. Understanding and mitigating gradient pathologies in physics-informed neural networks. arXiv:2001.04536 (2020). arXiv.
- [23] Sifan Wang, Xinling Yu, Paris Perdikaris. When and why PINNs fail to train: A neural tangent kernel perspective. arXiv:2007.14527 (2020). arXiv.

- [24] Aditi S. Krishnapriyan, Amir Gholami, Shandian Zhe, et al.. Characterizing possible failure modes in physics-informed neural networks. arXiv:2109.01050 (2021). arXiv.
- [25] Rafael Bischof, Michael Kraus. Multi-Objective Loss Balancing for Physics-Informed Deep Learning. arXiv:2110.09813 (2021). arXiv.
- [26] Levi McClenny, Ulisses Braga-Neto. Self-Adaptive Physics-Informed Neural Networks using a Soft Attention Mechanism. arXiv:2009.04544 (2020). arXiv.
- [27] Chenxi Wu, Min Zhu, Qinyang Tan, et al.. A comprehensive study of non-adaptive and residual-based adaptive sampling for physics-informed neural networks. arXiv:2207.10289 (2022). arXiv.
- [28] Mohammad Amin Nabian, Rini Jasmine Gladstone, Hadi Meidani. Efficient training of physics-informed neural networks via importance sampling. arXiv:2104.12325 (2021). arXiv.
- [29] E. Kharazmi, Z. Zhang, G. E. Karniadakis. Variational Physics-Informed Neural Networks For Solving Partial Differential Equations. arXiv:1912.00873 (2019). arXiv.
- [30] Ehsan Kharazmi, Zhongqiang Zhang, George Em Karniadakis. hp-VPINNs: Variational Physics-Informed Neural Networks With Domain Decomposition. arXiv:2003.05385 (2020). arXiv.
- [31] Chuan Guo, Geoff Pleiss, Yu Sun, Kilian Q. Weinberger. On Calibration of Modern Neural Networks. arXiv:1706.04599 (2017). arXiv.
- [32] Balaji Lakshminarayanan, Alexander Pritzel, Charles Blundell. Simple and Scalable Predictive Uncertainty Estimation using Deep Ensembles. arXiv:1612.01474 (2016). arXiv.
- [33] Liu Yang, Xuhui Meng, George Em Karniadakis. B-PINNs: Bayesian Physics-Informed Neural Networks for Forward and Inverse PDE Problems with Noisy Data. arXiv:2003.06097 (2020). arXiv.
- [34] Zhongkai Hao, Jiachen Yao, Chang Su, et al.. PINNacle: A Comprehensive Benchmark of Physics-Informed Neural Networks for Solving PDEs. arXiv:2306.08827 (2023). arXiv.
- [35] Makoto Takamoto, Timothy Praditia, Raphael Leiteritz, et al.. PDEBENCH: An Extensive Benchmark for Scientific Machine Learning. arXiv:2210.07182 (2022). arXiv.
- [36] Tamara G. Grossmann, Urszula Julia Komorowska, Jonas Latz, Carola-Bibiane Schönlieb. Can Physics-Informed Neural Networks beat the Finite Element Method?. arXiv:2302.04107 (2023). arXiv.
- [37] Krithika Manohar, Bingni W. Brunton, J. Nathan Kutz, Steven L. Brunton. Data-Driven Sparse Sensor Placement for Reconstruction. arXiv:1701.07569 (2017). arXiv.
- [38] Xuhui Meng, George Em Karniadakis. A composite neural network that learns from multi-fidelity data: Application to function approximation and inverse PDE problems. arXiv:1903.00104 (2019). arXiv.