Computer Vision-Driven In-Situ Defect Detection and Geometrical Quality Assessment for Automated Composite Formwork Fabrication

Azra van der Linden
Department of Computer Science, University of Central Florida, Orlando, FL, USA.
azravanderlinden@ucf.edu

Kerrie A. Delap
Department of Electrical Engineering and Computer Science, University of Kansas, Lawrence, KS, USA.
kerrie.a.delap@ku.edu

Ema Recio
Department of Computer Science, University of New Hampshire, Durham, NH, USA.
emarecio89@unh.edu

Abstract

Automated fabrication of composite formwork creates a quality-control problem that cannot be solved reliably by post-process inspection alone. Defects such as voids, gaps, overlaps, wrinkling, fibre bridging, edge fraying, and local geometric deviation can be introduced layer by layer, while the material and tooling state remain accessible for correction. This manuscript develops a computer-vision-driven methodological framework for in-situ defect detection and geometrical quality assessment during automated composite formwork fabrication. The central contribution is an integrated inspection logic that combines calibrated multi-view structured-light imaging, learned defect segmentation, uncertainty estimation, scan-to-computer-aided-design registration, and a tolerance-aware acceptance rule. Two-dimensional defect evidence is not treated as a standalone classification output; it is fused with three-dimensional signed surface deviation and with calibrated predictive confidence. The framework defines an imaging model, a class-weighted segmentation objective, an entropy-based uncertainty measure, a signed-distance geometric deviation field, a defect-severity score, a calibration audit, and a final accept-or-review decision rule. Analytical assessment shows that the main failure modes are not only missed defects but also false rejection under specular highlights, registration bias at thin free edges, and spatial misalignment between defect masks and geometric deviation. The proposed decision logic makes those trade-offs explicit, supports targeted rework before subsequent layers are deposited, and provides a traceable basis for process control in large-scale composite formwork production.

Keywords

composite formwork; in-situ defect detection; geometric quality assessment; structured-light vision; uncertainty calibration; automated fabrication

1. Introduction

Automated fabrication of large composite formwork requires a quality-control strategy that observes the workpiece while the process can still be corrected. Composite formwork combines structural requirements, freeform geometry, and layer-wise material deposition. A defect introduced at an intermediate layer can be buried by subsequent material, making external inspection ineffective after consolidation. The inspection problem is therefore both spatial and temporal: defects must be detected at the layer scale, localized with respect to the nominal geometry, and converted into an action before the next deposition step.

Computer vision offers a natural sensing route because cameras provide dense spatial coverage without contacting the workpiece. Deep learning has become a practical foundation for industrial visual inspection, particularly where defect appearance varies with lighting, material texture, and process state [1]. Convolutional neural networks can learn hierarchical features from raw images [2], fully convolutional architectures support dense prediction [3], encoder-decoder designs recover fine spatial detail [4], and residual learning enables deeper feature extractors [5]. Unsupervised anomaly-detection benchmarks further show that industrial inspection often involves rare or previously unseen defects rather than a fixed closed set of classes [6]. These developments are relevant to composite formwork, but they do not by themselves answer the manufacturing question: whether a visually detected anomaly violates a geometric or structural acceptance condition at the current layer.

The knowledge gap addressed here is the absence of a coherent decision framework that links defect appearance to geometric quality and to an actionable accept-or-review decision for large-scale composite formwork. Existing inspection logic is often separated into two streams. Visual defect detection produces masks or classifications in image space, while dimensional metrology produces point clouds and surface deviations in three-dimensional space. When these streams are not aligned, a defect mask may be visually convincing but geometrically misplaced; conversely, a geometric deviation may be measured accurately but not attributed to a specific material anomaly. For formwork, that separation is costly because rework decisions depend on both defect type and location relative to the nominal surface.

The central contribution of this manuscript is an uncertainty-aware inspection framework that fuses two-dimensional defect evidence with three-dimensional geometric deviation before an acceptance decision is made. The framework is constructed around four technical elements. First, a calibrated multi-view structured-light imaging model provides metric surface samples and image features under controlled illumination. Second, a learned segmentation module estimates defect probability maps with class weighting and predictive entropy. Third, scan-to-design registration produces a signed deviation field on the nominal formwork surface. Fourth, a tolerance-aware severity and decision rule separates acceptable process variation from defects requiring review. The analysis is developed as a design specification and an analytical assessment of failure modes, not as a report of a particular production trial. This framing is appropriate because acceptable thresholds depend on product tolerances, material system, layer height, and the cost of false acceptance versus false rejection.

2. Manufacturing Context and Problem Formulation

Additive and automated manufacturing processes place inspection close to the deposition event, where process parameters and material state can still be adjusted [7]. In composite formwork, the deposited material is anisotropic and often fibre-reinforced. Tape or tow

placement introduces directional material properties, laps, gaps, and compaction requirements. The geometric quality of the formwork depends on the accumulated layer stack, tooling alignment, and the stability of the deposited polymer or composite tape. Defects can be broadly grouped into material discontinuities, placement errors, consolidation defects, and geometric deviations. Voids, resin-rich or resin-poor regions, and incomplete compaction affect local stiffness and surface finish. Gaps, overlaps, and fibre bridging alter the effective thickness and load path. Wrinkles and out-of-plane fibre waviness can create local resin pockets and reduce structural efficiency. Edge fraying and tow termination errors affect dimensional stability and may propagate into subsequent layers.

The inspection task is not equivalent to generic surface-defect classification. A dark region in an image may be a void, a shadow caused by surface slope, a resin-rich area, or a specular artifact. A local elevation in a point cloud may be an intended ply build-up, a lap, or an unintended wrinkle. Therefore, the framework must represent both appearance and geometry and must preserve spatial correspondence between the two. The problem can be stated as follows: for each fabrication layer, determine whether the observed material state satisfies the nominal design and process constraints, and if not, identify the defect class, spatial extent, geometric severity, and confidence sufficiently for a rework or continuation decision.

Let the nominal formwork surface be S and the current-layer observation consist of calibrated images and a metric surface sample P. The system output is not a single class label but a set of defect regions Ω_k with class labels c_k, geometric deviation statistics, and calibrated confidence. A region may be acceptable if its deviation is within process tolerance or if it lies in a non-structural zone. A region must be reviewed if its severity exceeds a defect-specific threshold or if uncertainty prevents reliable attribution. This dual condition is important because a small but deep void can be more consequential than a larger shallow depression, while a geometric deviation at a free edge can be more critical than the same deviation in a non-mating surface.

3. Methodological Framework

3.1 Calibrated imaging and metric surface sampling

The imaging subsystem combines multiple calibrated cameras with structured-light projection. Camera calibration recovers intrinsic parameters and lens distortion, while multi-view calibration recovers the rigid transformations among cameras and the projector [8]. Structured-light patterns provide dense surface correspondences, and phase-shifted fringe projection provides sub-pixel phase information for surface profilometry [9], [10]. The basic projection relation for a calibrated camera is:

Here x is the homogeneous image point, X is the three-dimensional point expressed in the current fabrication-layer frame, K is the intrinsic camera matrix, and R and t are the rotation and translation from the layer frame to the camera frame. The symbol $\sim$ denotes equality up to scale. The relation is used to back-project image evidence into the layer frame and to forward-project nominal surface samples into each camera for occlusion reasoning. Calibration residuals are not treated as negligible; they define a spatial uncertainty envelope that must be carried into geometric deviation assessment.

The illumination design must address two material-specific problems. Composite tapes can be specular, and the fibre direction can produce oriented highlights that resemble linear defects. The surface can also be dark or translucent, reducing contrast for fringe projection. A

practical configuration therefore uses polarization filtering, multi-frequency fringe patterns, and short exposure bracketing. The purpose is not to produce a uniformly bright image but to retain phase validity and surface coverage. Pixels with invalid phase or saturated intensity are masked and excluded from metric fitting rather than silently interpolated.

3.2 Defect segmentation and spatial evidence

Defect segmentation is formulated as a dense binary or multi-class prediction problem over the image domain. Learned segmentation is preferred over thresholding because defect appearance depends on illumination angle, fibre orientation, and local surface slope. The segmentation network should use an encoder-decoder structure with skip connections to preserve thin boundaries such as gaps and frayed edges [3], [4]. The encoder can use residual blocks to extract stable features from reflective and low-contrast regions [5]. The output is a defect probability map p_i for each pixel i, with optional class-specific channels for void, gap, overlap, wrinkle, and edge defect.

A class-weighted objective is necessary because defect pixels are rare relative to background pixels and because some defect classes are more consequential than others. A suitable segmentation objective combines weighted cross-entropy with a confidence penalty:

$$L_seg = -\sum_{i=1}^{N} w_i y_i \log p_i + \lambda \sum_{i=1}^{N} p_i(1-p_i)$$

In this expression, N is the number of valid pixels, y_i is the binary or one-hot defect label, p_i is the predicted probability for the corresponding defect class, w_i is a class-weighting factor, and λ controls the strength of the confidence penalty. The second term discourages overconfident predictions in ambiguous regions. It does not replace a calibrated uncertainty model; it regularizes the segmentation output so that ambiguous reflections and shadows are not forced into high-confidence defect masks.

For multi-class defects, the binary formulation is applied per class and the class with the highest calibrated evidence is retained. The segmentation output is then filtered by geometric consistency: a predicted defect must have a metric footprint in the reconstructed surface, or it must be supported by multiple camera views, before it can enter the severity calculation. This step reduces false positives caused by single-view specular highlights.

3.3 Registration and signed geometric deviation

The reconstructed surface is registered to the nominal computer-aided-design model using a rigid or locally flexible registration. Iterative closest point provides a standard alignment between the measured point set and the nominal surface [11], and efficient variants improve correspondence selection and convergence behavior [12]. For formwork, rigid registration alone may be insufficient because the deposited layer can deform or because the nominal model is a design surface rather than the as-built tooling state. A locally flexible registration or a per-patch rigid alignment can be used, but the transformation must be regularized to prevent the registration from absorbing real defects.

Once correspondence is established, geometric quality is represented by a signed deviation field. For a measured point p_j with closest nominal point q_j and nominal unit normal n(q_j), the signed deviation is:

$$d_j = (p_j - q_j*)\cdot n(q_j*)$$

The point q_j is the closest point on the nominal surface S to p_j, and the sign indicates whether the measured surface lies outside or inside the nominal material envelope. Positive and negative deviations have different process meanings: outward deviation may indicate an overlap, wrinkle, or excess material, while inward deviation may indicate a gap, void, or incomplete deposition. The closest-point operator is defined as q_j = argmin_{q∈S} ||p_j - q||. Registration and deviation computation must be restricted to valid surface samples, excluding invalid phase pixels, occlusions, and edge regions where the normal estimate is unstable.

The signed deviation field is more informative than a single global dimensional error because it preserves the spatial pattern of the layer. A global root-mean-square error can hide a localized wrinkle or a deep void, whereas the deviation field allows defect regions to be compared against local tolerance and local design intent. The field is also used to define the geometric term in the severity score.

3.4 Uncertainty, severity, and decision logic

Predictive uncertainty is essential because the cost of a false accept in structural formwork is high. Bayesian deep-learning approximations and entropy-based measures provide a way to represent model uncertainty in dense prediction [13]. For a binary defect probability p_i, predictive entropy is:

$$H_i = -p_i \log p_i - (1 - p_i)\log(1 - p_i)$$

High entropy indicates that the segmentation model is uncertain about the pixel, often because of glare, low contrast, or an appearance not represented in the training distribution. Entropy alone is not a calibrated probability; it must be audited against observed accuracy. Neural-network calibration can be assessed using reliability diagrams and expected calibration error [14]. For a set of N predictions partitioned into M confidence bins B_m, the calibration audit is:

$$ECE = \sum_{m=1}^{M} (|B_m|/N)\, |acc(B_m) - conf(B_m)|$$

Here |B_m| is the number of predictions in bin m, acc(B_m) is the observed accuracy in that bin, and conf(B_m) is the mean predicted confidence. The audit is used to adjust decision thresholds before the inspection logic is accepted for a given material and lighting configuration. This is particularly important for composite surfaces, where specular artifacts can produce confident but incorrect masks.

The defect severity score combines spatial extent, geometric deviation, and uncertainty for each connected defect region Ω_k:

$$S_k = \alpha A_k + \beta \max_{j \in \Omega_k} |d_j| + \gamma H_k$$

A_k is the normalized area of the defect region, max|d_j| is the maximum absolute signed deviation within the region, and H_k is the mean entropy over the region. The coefficients α, β, and γ are process-dependent weights that express the relative importance of area, geometric severity, and evidence uncertainty. They should be set from design tolerances, structural requirements, and rework cost, not from visual preference alone. The maximum deviation is used rather than the mean because a small deep void can dominate the local mechanical consequence.

A final acceptance rule converts the continuous evidence into a manufacturing action. Let A = 1 denote acceptance and A = 0 denote review or rework. A conservative rule is:

$$A = 1\left[\max_j |d_j| \le T_g \wedge \max_k S_k \le T_d\right]$$

T_g is the geometric tolerance for the current layer, and T_d is the defect-severity threshold. The symbol ∧ denotes logical conjunction, and 1[·] is the indicator function. The rule can be relaxed for non-structural zones and tightened for mating surfaces, inserts, or load-transfer regions. The important design choice is that no defect region bypasses both geometric and severity checks. A visually small defect that produces a large local deviation is still flagged, and a geometrically smooth region with high uncertainty is sent to review rather than accepted automatically.

Large-scale additive formwork with rapid fibre-reinforced polymer tape lamination illustrates why layer-wise inspection is necessary [16]. In such a process, the nominal geometry is not a static external mold; it is produced by the deposition itself. Registration to a static design model must therefore be accompanied by a layer-state model that records which nominal features have already been fabricated. The required reference is relevant here because it establishes the fabrication context of automated large-scale formwork and tape lamination, not because it supplies a defect-detection algorithm. The inspection framework must therefore operate on a layer-specific nominal surface and on a process state that changes after each deposition pass.

4. Analytical Assessment of Failure Modes and Trade-offs

The first failure mode is appearance ambiguity caused by specular reflection and fibre orientation. A deep-learning segmentation model can learn to associate bright streaks with wrinkles or gaps, but the same streak can be produced by a change in surface slope. The geometric deviation field provides a partial disambiguation: a wrinkle should produce a coherent out-of-plane deviation, whereas a specular highlight should not. The severity score handles this by giving weight to max|d_j|. However, if the surface reconstruction is invalid in the highlight region, the defect should be classified as uncertain rather than accepted. This is the correct behavior for formwork because a missed wrinkle can be buried by subsequent layers.

The second failure mode is registration bias at free edges and thin features. Iterative closest point tends to align well on broad surfaces but can slide along edges or collapse thin features. A locally flexible registration can reduce this bias but may also absorb a real edge defect. The trade-off is controlled by regularization: too much flexibility hides defects, while too little flexibility produces large false deviations at legitimate tooling transitions. The signed distance formulation makes the problem visible because edge bias appears as a systematic deviation pattern, not as a random segmentation error.

The third failure mode is spatial mismatch between defect masks and geometric deviation. A defect mask produced in image space must be projected into the surface frame before its severity is computed. If the projection uses an incorrect depth or an outdated calibration, the mask can be shifted by several pixels and overlap the wrong deviation values. The imaging equation and calibration residual envelope are therefore not preliminary details; they determine whether the fused severity score is meaningful. Multi-view consistency is the main safeguard. A defect region that is not supported by at least two views, or that does not project onto a valid surface patch, should be marked as uncertain.

The fourth failure mode is threshold mis-specification. A single global threshold for all defects will either over-inspect benign surface texture or under-inspect structurally relevant gaps. Defect-specific thresholds are preferable. Voids and gaps should be judged by depth and area relative to the layer thickness. Wrinkles should be judged by out-of-plane amplitude and

spatial wavelength. Edge defects should be judged by their distance from the mating surface. The acceptance rule can be implemented with a threshold vector rather than a scalar, while the severity score provides a common ranking for review.

The fifth failure mode is distribution shift across material batches, lighting conditions, and tooling states. Calibration metrics such as expected calibration error are useful because they reveal when the model's confidence no longer matches its accuracy. A high-confidence false positive is more dangerous than a low-confidence false positive because it can trigger unnecessary rework and delay. The decision logic should therefore include a review zone: regions with moderate severity or high entropy are neither accepted nor rejected automatically but are flagged for human inspection or additional sensing. This three-way outcome is more realistic than a binary accept/reject output in a production setting where rework is expensive but feasible.

Offline non-destructive testing remains a necessary reference for composite quality because some internal defects, such as delamination and porosity, are not directly visible at the surface [17]. Automated prepreg layup has long recognized placement defects and process-induced anomalies [18]. Delamination can dominate failure in fibre-reinforced composites [19], and out-of-plane wrinkling is a recognized geometric and structural defect in composite forming [20]. These works inform the defect taxonomy and the severity logic, but the in-situ framework is not a replacement for volumetric inspection. Instead, it provides early evidence that can prevent the most visible and correctable defects from being buried.

5. Technical Discussion and Implementation Conditions

The proposed framework is most defensible when the inspection system is calibrated to the actual material and lighting conditions before thresholds are fixed. A composite tape that is highly specular under one angle may be diffuse under another; a single illumination geometry is unlikely to satisfy all defect classes. Polarization and multi-frequency structured light address this problem, but they increase acquisition time and data volume. The design trade-off is between spatial coverage and temporal latency. For large formwork, a scanning head that moves with the deposition machine can provide dense coverage, but the motion stage must be synchronized with the camera exposure and fringe projection. Otherwise, motion blur invalidates the phase estimate and the geometric deviation field becomes unreliable.

The computational architecture should separate deterministic calibration and registration from learned perception. Calibration, projection, and signed-distance computation are well-defined geometric operations and should be auditable. Segmentation and uncertainty estimation are learned components and should be monitored with calibration audits and out-of-distribution checks. This separation makes it possible to diagnose a false rejection as either a geometric problem, a calibration problem, or a model-confidence problem. It also allows the acceptance rule to be recomputed when tolerances change without retraining the segmentation model.

The framework has clear boundary conditions. It is limited to surface-visible and near-surface defects; it cannot certify internal consolidation or delamination. It assumes that the nominal design surface and layer-state model are available and correct. It also assumes that the deposition machine can pause or mark a region for rework without compromising the process. Where these assumptions fail, the inspection output remains useful as a screening

record but not as a final structural quality certificate. These are natural constraints of in-situ optical inspection, not defects of the decision logic.

6. Conclusion

This manuscript has developed a computer-vision-driven framework for in-situ defect detection and geometrical quality assessment in automated composite formwork fabrication. The central argument is that defect appearance and geometric deviation must be fused before a manufacturing decision is made. Calibrated multi-view structured-light imaging provides metric surface samples, learned segmentation provides defect evidence, registration provides signed deviation, and calibrated uncertainty provides a guard against overconfident decisions. The severity score and acceptance rule convert those quantities into an actionable outcome that can support rework before defects are buried.

The analytical assessment identifies the principal failure modes as specular ambiguity, edge registration bias, spatial mismatch between image masks and surface deviations, threshold mis-specification, and distribution shift. The proposed decision logic addresses these failure modes by requiring multi-view support, excluding invalid phase pixels, carrying calibration uncertainty into the deviation field, using defect-specific severity weights, and including a review zone for high-entropy cases. The framework does not replace volumetric non-destructive testing, but it provides a traceable and timely layer-wise quality record. Its practical value lies in making inspection decisions explicit and auditable, which is a prerequisite for closed-loop control in large-scale composite formwork production.

References

1. Wang, J., Ma, Y., Zhang, L., Gao, R. X., & Wu, D. (2018). Deep learning for smart manufacturing: Methods and applications. Journal of Manufacturing Systems, 48, 144–156.

2. Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet classification with deep convolutional neural networks. Advances in Neural Information Processing Systems, 25, 1097–1105.

3. Long, J., Shelhamer, E., & Darrell, T. (2015). Fully convolutional networks for semantic segmentation. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (pp. 3431–3440).

4. Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolutional networks for biomedical image segmentation. In Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015 (pp. 234–241). Springer.

5. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (pp. 770–778).

6. Bergmann, P., Fauser, M., Sattlegger, D., & Steger, C. (2019). MVTec AD—A comprehensive real-world dataset for unsupervised anomaly detection. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 9592–9600).

7. Gao, W., Zhang, Y., Ramanujan, D., Ramani, K., Chen, Y., Williams, C. B., Wang, C. C. L., Shin, Y. C., Zhang, S., & Zavattieri, P. D. (2015). The status, challenges, and future of additive manufacturing in engineering. Computer-Aided Design, 69, 65–89.

8. Zhang, Z. (2000). A flexible new technique for camera calibration. IEEE Transactions on Pattern Analysis and Machine Intelligence, 22(11), 1330–1334.

9. Salvi, J., Fernandez, S., Pribanic, T., & Llado, X. (2010). A state of the art in structured light patterns for surface profilometry. Pattern Recognition, 43(8), 2666–2680.

10. Gorthi, S. S., & Rastogi, P. (2010). Fringe projection techniques: Whither we are? Optics and Lasers in Engineering, 48(2), 133–140.

11. Besl, P. J., & McKay, N. D. (1992). A method for registration of 3-D shapes. IEEE Transactions on Pattern Analysis and Machine Intelligence, 14(2), 239–256.

12. Rusinkiewicz, S., & Levoy, M. (2001). Efficient variants of the ICP algorithm. In Proceedings of the Third International Conference on 3-D Digital Imaging and Modeling (pp. 145–152). IEEE.

13. Kendall, A., & Gal, Y. (2017). What uncertainties do we need in Bayesian deep learning for computer vision? Advances in Neural Information Processing Systems, 30, 5574–5584.

14. Guo, C., Pleiss, G., Sun, Y., & Weinberger, K. Q. (2017). On calibration of modern neural networks. In Proceedings of the 34th International Conference on Machine Learning (pp. 1321–1330). PMLR.

15. Savio, E., De Chiffre, L., Schmitt, R., & Müller, U. (2007). Metrology of freeform shaped parts. CIRP Annals, 56(2), 810–835.

16. Xin, Z. Y., Zhu, G. Q., Gattas, J. M., & Luo, D. (2025). Automated large-scale additive manufacturing of structural formwork with rapid fibre-reinforced polymer tape lamination. Automation in Construction, 171, 105978.

17. Wang, B., Zhong, S., Lee, T.-L., Fancey, K. S., & Mi, J. (2020). Non-destructive testing and evaluation of composite materials/structures: A state-of-the-art review. Advances in Mechanical Engineering, 12(4), 1687814020913761.

18. Lukaszewicz, D. H.-J. A., Ward, C., & Potter, K. D. (2012). The engineering aspects of automated prepreg layup: History, present and future. Composites Part B: Engineering, 43(3), 997–1009.

19. Wisnom, M. R. (2012). The role of delamination in failure of fibre-reinforced composites. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, 370(1965), 1850–1870.

20. Boisse, P., Colmars, J., Hamila, N., Naouar, N., & Steer, Q. (2018). Bending and wrinkling of composite fiber preforms and prepregs. A review and new developments in the draping simulations. Composites Part B: Engineering, 141, 234–249.