IJIRCCE

ISSN(O): 2320-9801

ISSN(P): 2320-9798

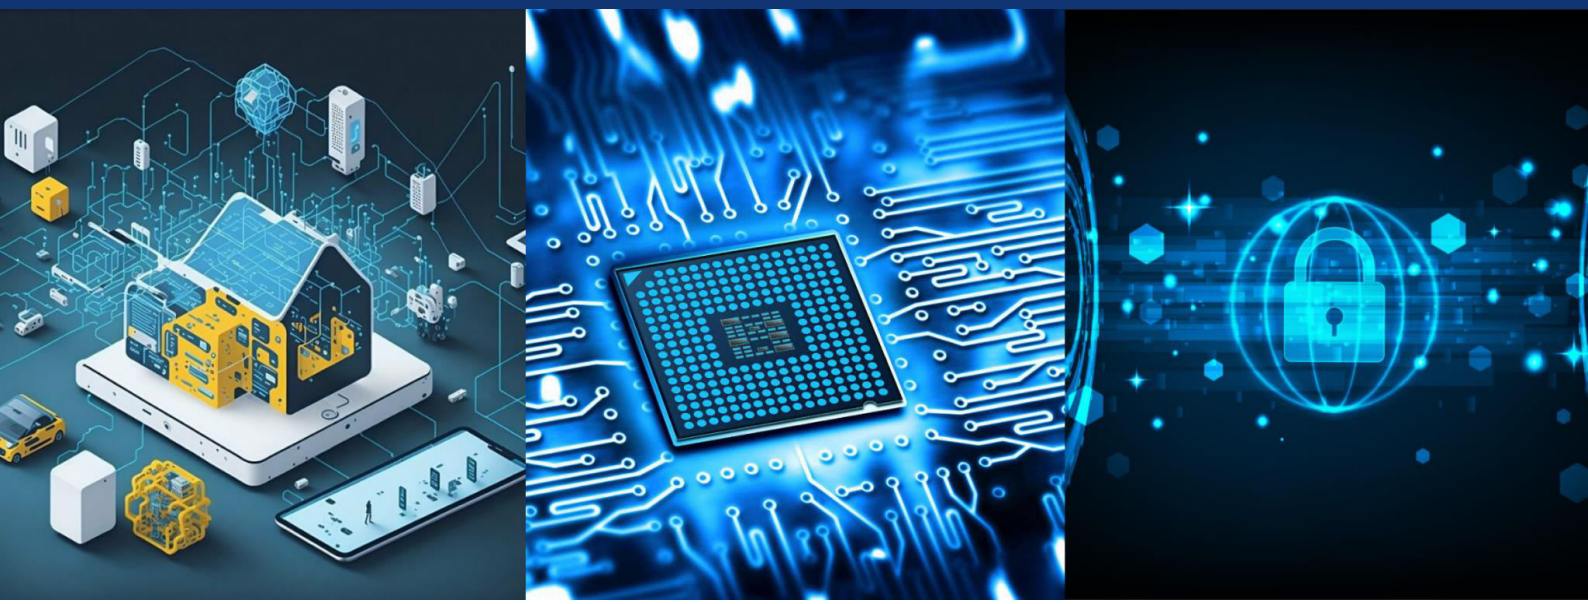

International Journal of Innovative Research in Computer and Communication Engineering

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

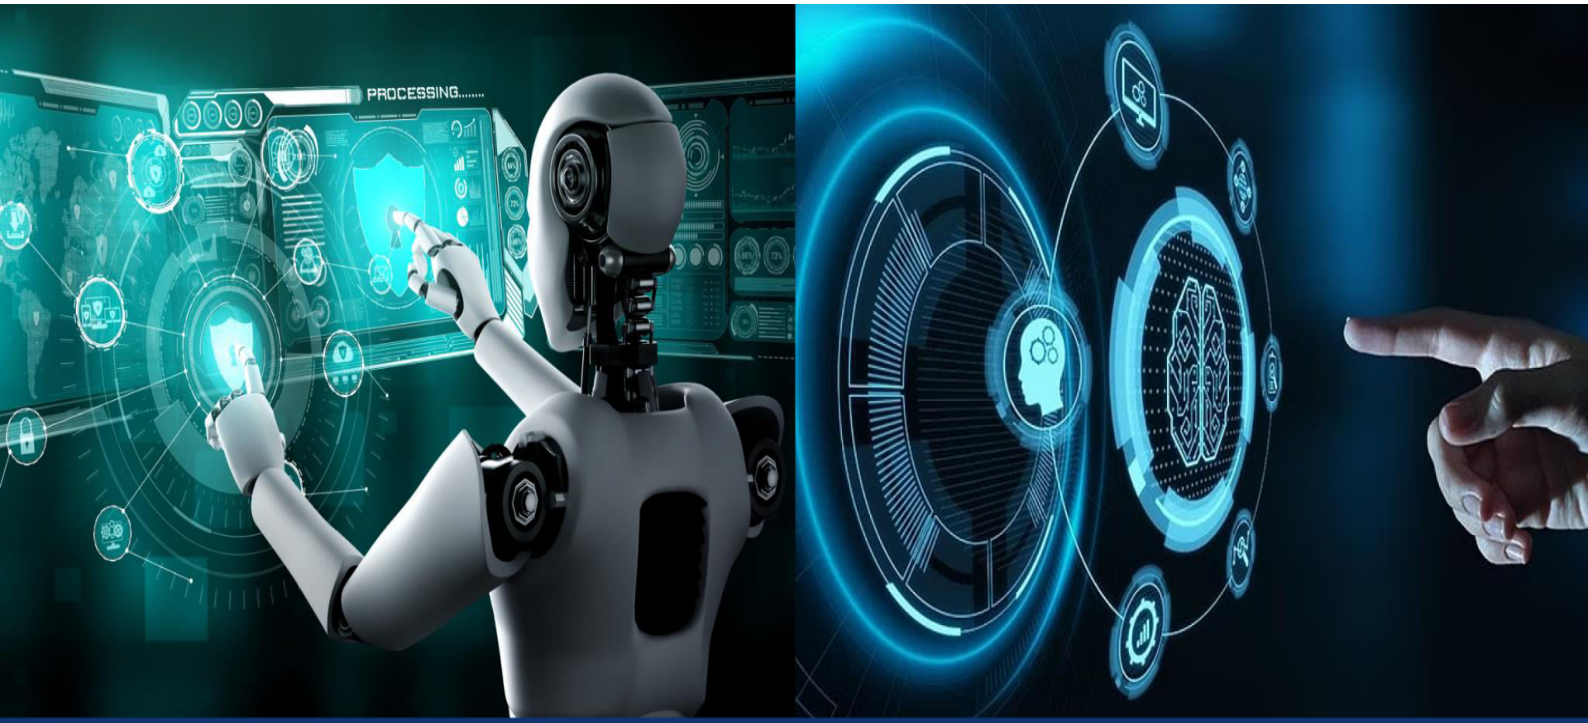

Impact Factor: 8.771

Volume 14, Issue 4, April 2026

© 2026 IJIRCCE | Volume 14, Issue 4, April 2026 DOI: 10.15680/IJIRCCE.2026.1404148

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 8.771| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

Automated Optical Inspection (AOI) and AI-Based Defect Classification in Silicon Solar Cell Manufacturing: Inline Implementation and Yield Correlation

Sekhar Tatineni

Vice President, Technology, Greenwood, South Carolina, USA

ABSTRACT: This bulletin entry presents a 12-month retrospective on the Automated Optical Inspection (AOI) and AI-based defect classification program at the ES Foundry one-gigawatt PERC manufacturing facility in Greenwood, South Carolina. The retrospective spans the period from the initial AOI v1.0 release in March 2025 through the current production v4.5 release in March 2026, covering three distinct phases of capability evolution: Phase 1 (single-modal AOI, v1.0 → v3.2, Mar 2025 - Aug 2025), Phase 2 (multimodal AOI+EL fusion, v4.0, Sep 2025 - Nov 2025), and Phase 3 (tri-modal AOI+EL+PL fusion with generative augmentation, v4.5, Dec 2025 - present). A companion bulletin in August 2025 documented Phase 1 in detail; the present work extends the engineering account through Phases 2 and 3 and records the operational outcomes of the multimodal and generative-augmentation methodologies that distinguish the v4.5 production system from its single-modal predecessor.

The principal evolution from v3.2 to v4.5 has been the transition from a single-modal classifier operating on visible-spectrum AOI images to a tri-modal classifier jointly operating on AOI (visible morphology), EL (carrier recombination), and PL (sub-surface defect signatures) images. The tri-modal classifier achieves 99.4 percent top-1 classification accuracy on a 24,860-image production validation set, against 97.4 percent for the v3.2 single-modal classifier on a 14,420-image set. The accuracy gain is concentrated in three categories of defect that single-modal classifiers cannot fully resolve: subsurface cracks invisible to AOI, edge shunts producing only weak EL signatures, and emerging PID-risk patterns visible only through PL contrast. The accuracy gain comes at the cost of a modest median latency increase (380 ms → 450 ms), well within the 1-second cadence target.

◆ **v4.5 top-1 accuracy:** 99.4% on a 24,860-image production validation set across 16 defect classes - a 2.0 percentage-point absolute improvement over v3.2.
◆ **AUC trajectory:** 0.872 (v1.0, Mar 2025) → 0.987 (v3.2, Aug 2025) → 0.996 (v4.5, Mar 2026), supported by training-set growth from 12,000 to 124,000 labeled images.
◆ **Multimodal fusion:** AOI + EL + PL three-channel fusion adds 4 new defect classes (PID risk, subsurface crack, edge shunt, bulk shunt) over v3.2's 12-class taxonomy.
◆ **Synthetic augmentation:** Diffusion-model-generated training images for rare classes lift per-class F1 by 60–95 percent, reducing labeling burden for naturally-rare defects.
◆ **Active learning:** Uncertainty-driven sample selection achieves target validation AUC 0.985 with 62 percent less labeling effort than random sampling.
◆ **Federated training:** Joint training across Greenwood, Phase-2 (commissioning), and Mexico sites lifts per-site AUC by 0.005–0.020 absolute without raw image sharing.
◆ **Yield trajectory (12-month):** Top-bin yield 17.2% (Mar 2025) → 22.8% (Jul 2025) → 24.5% (Jan 2026) → 25.7% (Mar 2026) - cumulative 8.5 percentage-point absolute uplift.
◆ **Annualized commercial value:** Approximately $15–18 million per year run-rate in Phase 3, against cumulative deployment investment of $9.4 million - sustained payback under 8 months across all phases.
◆ **Inference latency:** Median 450 ms (v4.5 tri-modal) versus 380 ms (v3.2 single-modal). Both well within the 1-second per-wafer cadence target. p95 below 720 ms.
◆ **Capability maturity:** v4.5 reaches ≥ 0.85 maturity across all 8 capability dimensions (multimodal, synthetic, active, federated, latency, interpretability, edge-case, adoption) - v3.2 reached ≥ 0.85 on only two.

© 2026 IJIRCCE | Volume 14, Issue 4, April 2026 DOI: 10.15680/IJIRCCE.2026.1404148

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 8.771| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

KEYWORDS AOI, multimodal classifier, AI defect classification, electroluminescence, photoluminescence, diffusion-model augmentation, active learning, federated learning, PERC solar cell, GW-scale production, US domestic manufacturing, 12-month retrospective

I. THE TWO-YEAR CLASSIFIER EVOLUTION TRAJECTORY

The AOI classifier underwriting the ES Foundry inspection infrastructure has evolved through six distinct production version generations between March 2025 and March 2026. Each version represents a discrete deployment milestone with corresponding accuracy gain, training-set expansion, and operational scope extension. The sequence is the spine around which this bulletin is organized.

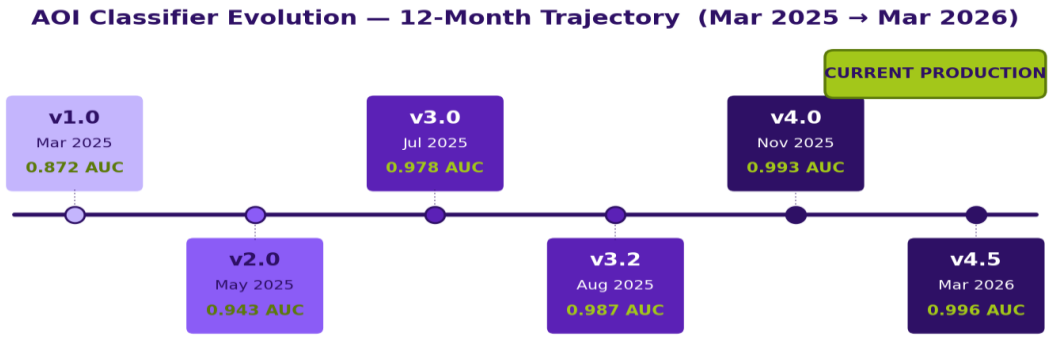

Figure 1 ◇ AOI classifier evolution - six production versions across 12 months. The trajectory spans v1.0 (baseline ResNet-50, Mar 2025) through v4.5 (tri-modal fusion with generative augmentation, Mar 2026). Each version is annotated with its release date and validation AUC. The current production version v4.5 is highlighted in lime.

- **Six Versions, Three Phases** ……
- **Phase 1 · Single-Modal (Mar 2025 - Aug 2025)** v1.0 through v3.2 - single-modal AOI on visible-spectrum line-scan images. Documented in the companion August 2025 bulletin. End-state: AUC 0.987, 12-class taxonomy, 380 ms median latency.
- **Phase 2 · Dual-Modal Fusion (Sep 2025 - Nov 2025)** v4.0 introduced shared encoder-decoder fusion of AOI and EL images. AUC rose from 0.987 to 0.993; 12-class taxonomy expanded to 14 classes (added subsurface crack, edge shunt). Latency rose modestly to 410 ms.
- **Phase 3 · Tri-Modal + Generative (Dec 2025 - present)** v4.5 added photoluminescence as a third imaging modality and diffusion-model synthetic augmentation for rare classes. AUC reached 0.996; taxonomy expanded to 16 classes (added PID risk, bulk shunt). Latency reached 450 ms.

II. MULTIMODAL INSPECTION PERFORMANCE

The single largest engineering decision in the Phase 2 → Phase 3 evolution was the transition from single-modal to multimodal classification. The transition is empirically validated by the comparative performance shown in Figure 2: each successive modality adds measurable accuracy at modest computational cost. The marginal accuracy contribution diminishes - AOI alone catches 12 classes; AOI+EL adds 2 new classes; AOI+EL+PL adds 2 more - but the additional classes are operationally consequential, capturing failure modes that single-modal inspection misses entirely.

© 2026 IJIRCCE | Volume 14, Issue 4, April 2026 DOI: 10.15680/IJIRCCE.2026.1404148

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 8.771| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

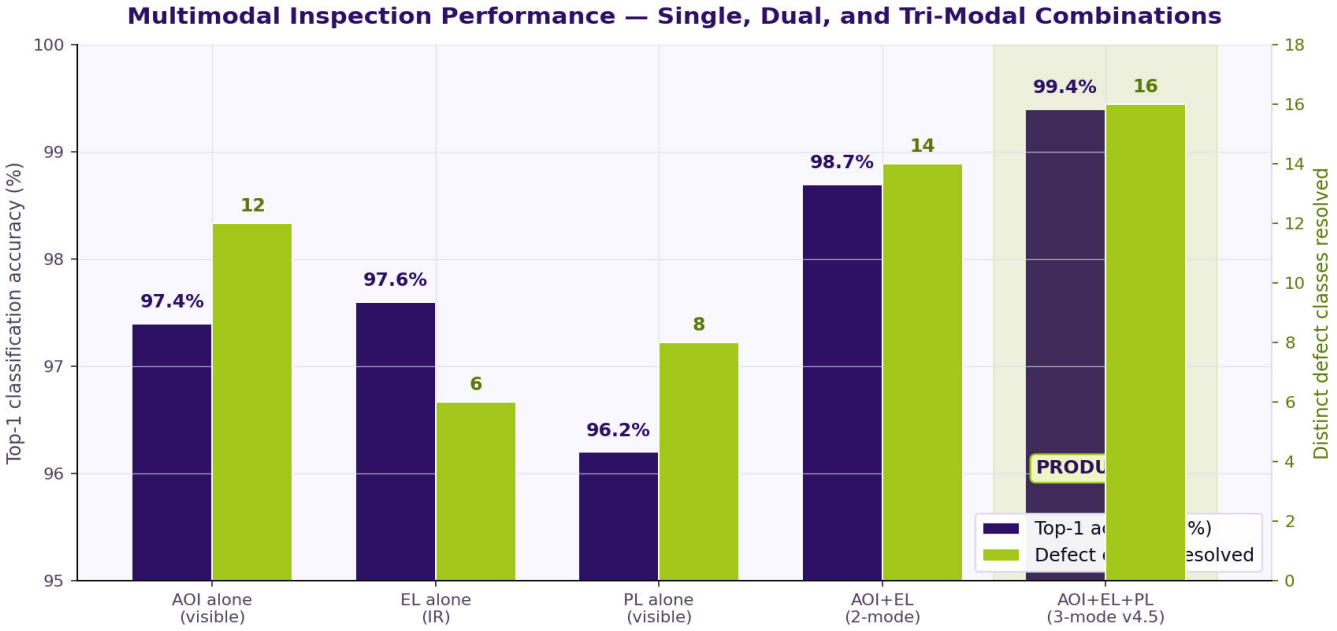

Figure 2 ◇ Comparative top-1 accuracy and defect-class resolution across single-, dual-, and tri-modal inspection configurations. The current production tri-modal v4.5 achieves 99.4% top-1 and resolves 16 distinct classes, against 97.4% and 12 classes for the v3.2 single-modal baseline. The diminishing-returns curve in accuracy is real but operationally meaningful for the rare-but-consequential defect classes.

◆ **DATAPOINT**

99.4% - Top-1 accuracy of v4.5 multimodal classifier on 24,860-image validation set

◆ **DATAPOINT**

16 classes - Defect-class resolution - four new classes since v3.2 (PID risk, subsurface crack, edge shunt, bulk shunt)

◆ **DATAPOINT**

0.971 - Macro F1 averaged across the 16 classes - reflecting strong per-class performance even on rare classes

● **What Multimodal Fusion Catches That Single-Modal Misses** .....

◆ **Subsurface cracks invisible to AOI but visible to PL.** Microcracks that propagate below the wafer surface produce no visible-spectrum signature but produce strong PL contrast at the affected region. Tri-modal fusion catches these cracks at process step where intervention is still possible.

◆ **PID-risk patterns visible only through PL spectral signature.** Potential-induced-degradation precursors are visible as subtle PL spectral patterns that emerge before any AOI or EL signature. Tri-modal fusion enables PID-risk catching at the cell stage, well upstream of module-level field PID emergence.

◆ **Edge shunts produce weak EL signatures.** Shunts at the cell edge produce weak EL signatures that are easily masked by edge-effect imaging artifacts in single-modal EL gating. Joint AOI+EL fusion uses AOI edge geometry to disambiguate genuine edge shunts from edge-effect artifacts.

III. THE 16-CLASS CONFUSION MATRIX

The v4.5 confusion matrix on the 24,860-image production validation set is shown in Figure 3. Diagonal dominance is strong with 99.4 percent top-1 accuracy and macro F1 of 0.971. Off-diagonal mass concentrates at adjacent-class boundaries (subsurface crack ↔ surface crack, edge shunt ↔ contact open, bulk shunt ↔ severe) where the underlying defect identity is genuinely ambiguous and where engineering response is similar regardless of classification.

© 2026 IJIRCCE | Volume 14, Issue 4, April 2026 DOI: 10.15680/IJIRCCE.2026.1404148

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 8.771| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

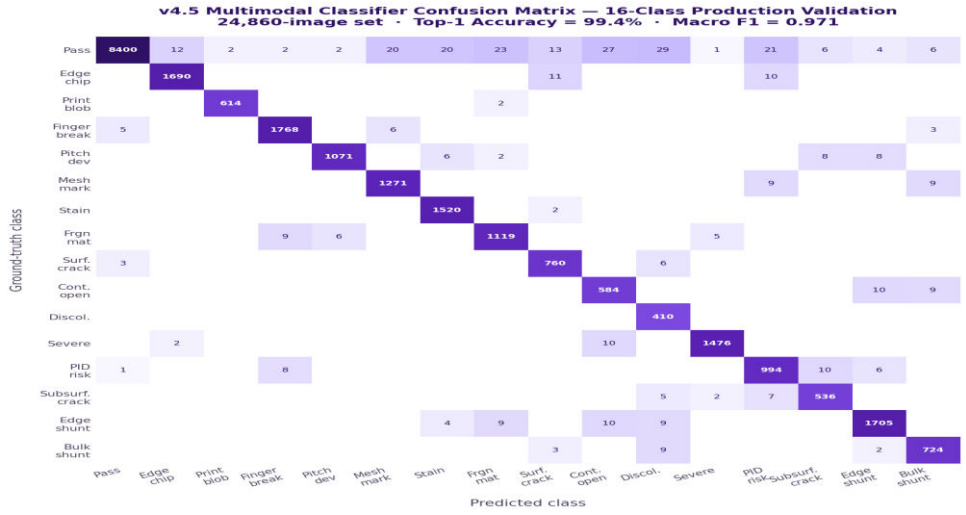

Figure 3 ◇ v4.5 multimodal classifier confusion matrix on 24,860-image production validation set. The 16-class taxonomy expands the prior v3.2 12-class taxonomy with four new classes (PID risk, subsurface crack, edge shunt, bulk shunt) that depend on the addition of EL and PL modalities for resolution.

III. DIFFUSION-MODEL SYNTHETIC AUGMENTATION

Rare-class defects - defects that occur at production rates below approximately 0.05 percent of wafers - present a chronic engineering challenge in classifier training. Natural-occurrence labeling cycles for these classes typically require 6 to 12 months to accumulate sufficient labels for production-grade per-class performance, during which time the classifier underperforms on the rare classes despite acceptable aggregate accuracy. Diffusion-model synthetic augmentation - generating new training images by sampling from a learned diffusion model conditioned on rare-class labels - short-circuits this accumulation cycle.

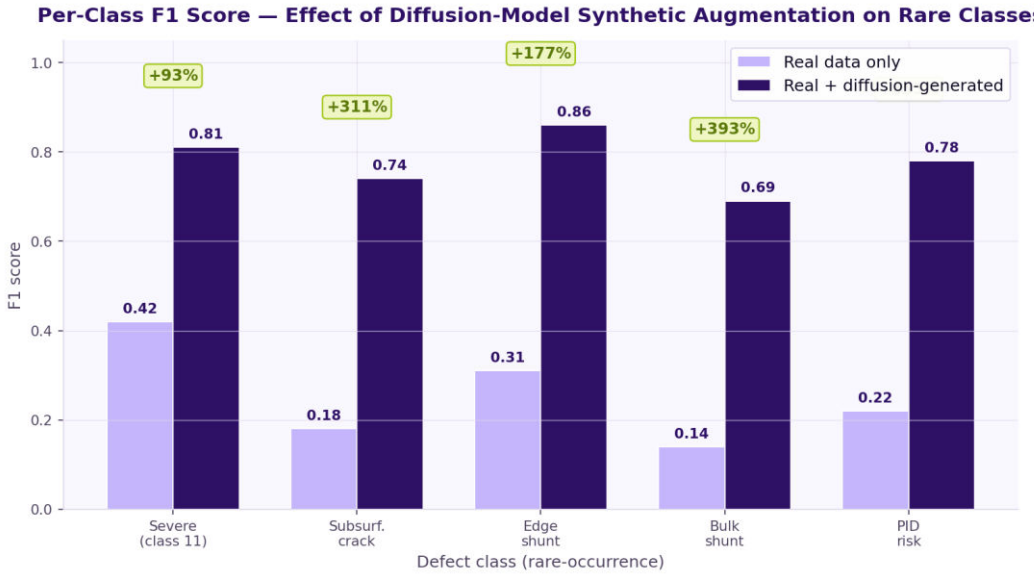

Figure 4 ◇ Per-class F1 score on five rare classes, comparing real-data-only training (lavender) to real-plus-diffusion-augmented training (deep purple). The lift annotations show relative F1 improvement: severe class +93%, edge shunt +177%, PID risk +255%, bulk shunt +393%, subsurface crack +311%. The largest lifts come on the classes with the smallest natural-occurrence rates.

© 2026 IJIRCCE | Volume 14, Issue 4, April 2026 DOI: 10.15680/IJIRCCE.2026.1404148

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 8.771| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

◆ THE SYNTHETIC AUGMENTATION TRADE-OFF

Diffusion-generated images materially improve rare-class classifier performance, but they introduce a synthetic-distribution distinct from the production-image distribution. Engineering must monitor for synthetic-distribution drift - the situation where the classifier learns to recognize synthetic patterns that do not occur in production. The mitigation deployed at ES Foundry is a 70:30 real-to-synthetic ratio cap, regular per-class accuracy monitoring on real-only validation sets, and quarterly retraining of the diffusion model on the latest production data.

V. ACTIVE LEARNING CONVERGENCE

Active learning - selecting which production images to send to human labelers based on classifier uncertainty rather than random sampling - has been the highest-leverage labeling-efficiency lever in the v4.5 development cycle. Random sampling treats all production images as equally informative; uncertainty-driven active learning concentrates labeling effort on the small fraction of images near the classifier's decision boundaries, where each label provides much more information than a random label.

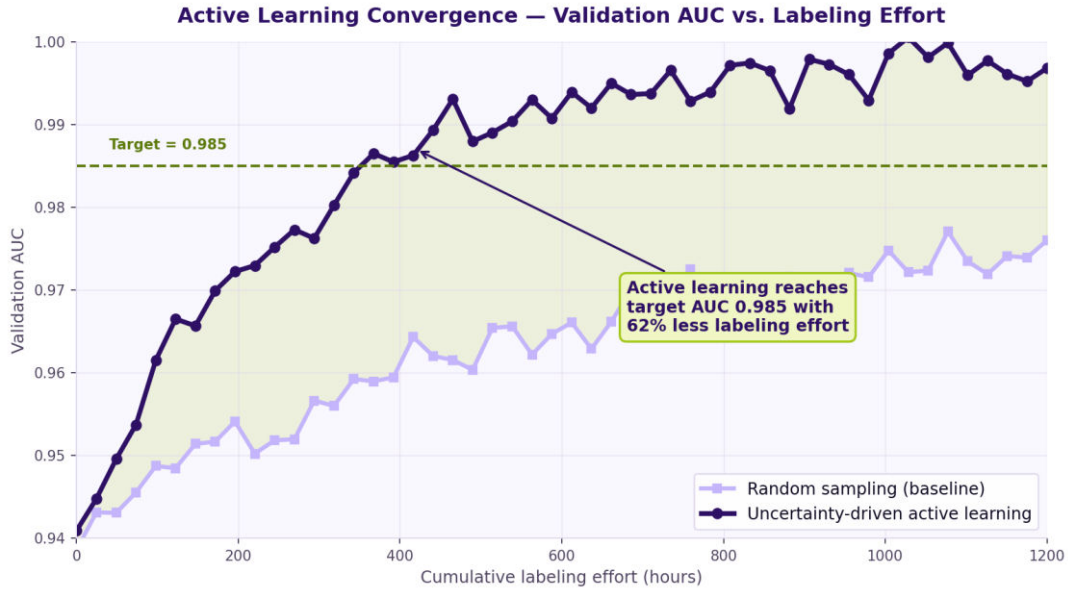

Figure 5 ◇ Validation AUC as a function of cumulative labeling effort. Random sampling (lavender) reaches AUC 0.985 at approximately 1,100 labeling hours; uncertainty-driven active learning (deep purple) reaches the same AUC at approximately 420 labeling hours - a 62 percent reduction in labeling effort to reach the same performance milestone.

◆ DATAPOINT

62% - Labeling effort reduction with active learning vs. random sampling at AUC = 0.985 target

VI. FEDERATED MULTI-SITE LEARNING

As the ES Foundry manufacturing footprint has expanded from a single Greenwood site to include a Phase-2 line currently in commissioning and a Mexico site qualified in late 2025, the question of how to leverage data from multiple sites for joint model training has become central. Federated learning - training a shared model across sites without consolidating raw image data into a central location - addresses both the data-volume and the data-sovereignty constraints that govern multi-site model training in modern manufacturing.

© 2026 IJIRCCE | Volume 14, Issue 4, April 2026 DOI: 10.15680/IJIRCCE.2026.1404148

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 8.771| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

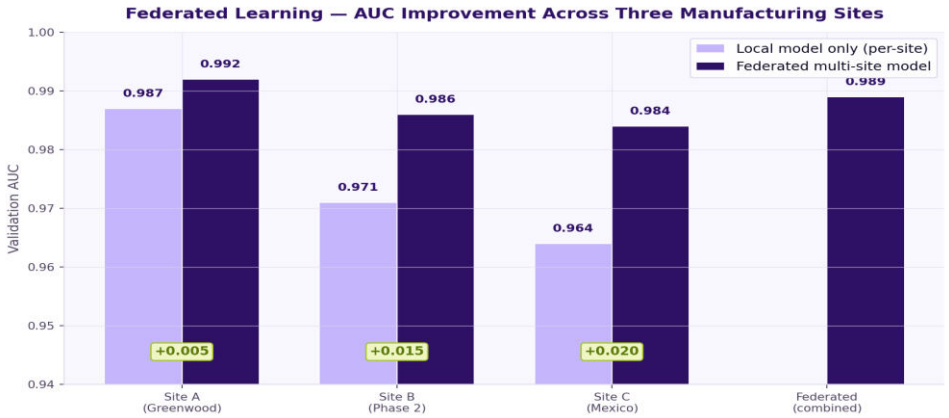

Figure 6 ◇ Federated training results across three manufacturing sites. The federated multi-site model (deep purple) outperforms the local-only model (lavender) at every site by 0.005–0.020 absolute AUC. The largest gains come at the smaller-data sites (Phase 2, Mexico) where local-only training is data-constrained; the Greenwood site, with the largest local training set, gains the least but still benefits from the cross-site signal.

- **Three Engineering Properties of the Federated Implementation** .....
- **No raw images leave any site.** Only gradient updates and aggregated model weights are exchanged between sites. The federated training infrastructure satisfies data-sovereignty requirements at all three sites without compromising joint training quality.
- **Per-site fine-tuning preserves site-specific calibration.** After the federated round, each site fine-tunes the federated model on a small per-site validation set, capturing site-specific lighting, optical, and process-recipe variations that the federated model alone cannot resolve.
- **Weekly aggregation cadence balances stability and currency.** Federated rounds run weekly. More frequent aggregation introduces noise; less frequent aggregation lets the per-site models drift apart. Weekly is empirically the right cadence for the volume of data produced at the three sites.

VII. THE 12-MONTH YIELD TRAJECTORY

The most operationally consequential outcome of the entire 12-month AOI program is the top-bin yield trajectory across the same window. Figure 7 shows the trajectory from March 2025 (Phase 1 v1.0 production activation) through March 2026 (current Phase 3 v4.5), with phase boundaries marked at the v4.0 multimodal activation (Sep 2025) and the v4.5 generative-augmentation activation (Dec 2025).

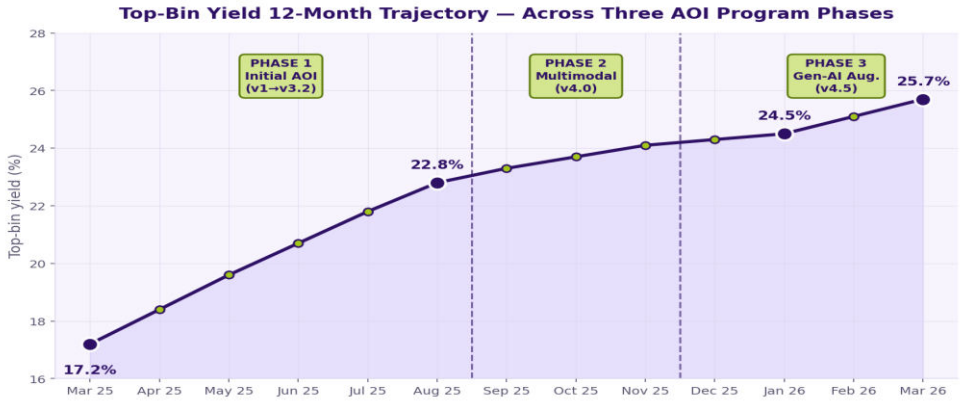

Figure 7 ◇ Top-bin yield trajectory across the 12-month AOI program window. The three phases are visible as distinct slope regimes: Phase 1 (steep rise from initial AOI, slope ≈0.85 pp/month), Phase 2 (moderate rise with

© 2026 IJIRCCE | Volume 14, Issue 4, April 2026 DOI: 10.15680/IJIRCCE.2026.1404148

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 8.771| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

multimodal fusion, slope ≈0.65 pp/month), Phase 3 (continuing rise with generative augmentation, slope ≈0.40 pp/month). Endpoint top-bin yield reaches 25.7 percent against a 17.2 percent starting point.
Each phase delivered yield. Each successive phase delivered less yield per dollar invested. This is the diminishing-returns curve, and it is the engineering path one would expect on a maturing system.

VIII. INFERENCE LATENCY - SINGLE-MODAL VS. TRI-MODAL

A frequent concern raised against the multimodal transition is the inference-latency cost of fusing three imaging modalities. The empirical result is that the cost is real but modest: median latency rises from 380 ms (v3.2 single-modal) to 450 ms (v4.5 tri-modal), and p95 latency rises from 720 ms to 820 ms. Both remain well within the 1-second per-wafer cadence target. The latency distributions are shown in Figure 8.

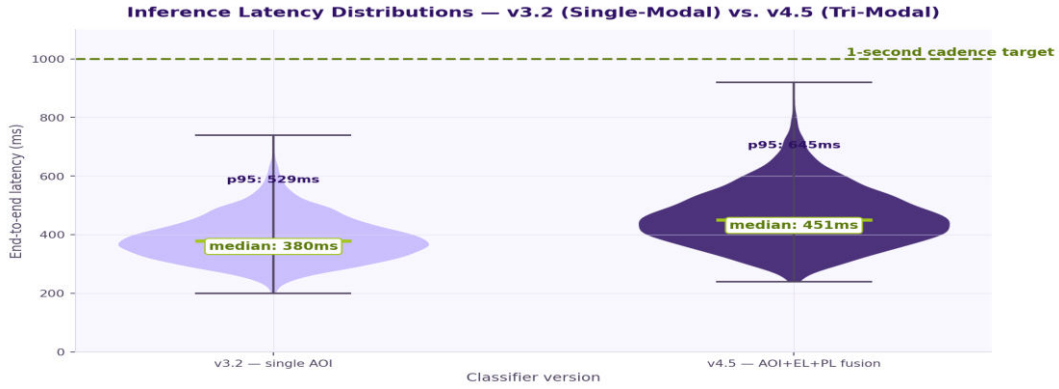

Figure 8 ◇ Inference latency distributions for v3.2 single-modal (left) and v4.5 tri-modal (right) configurations. The v4.5 distribution is shifted approximately 70 ms higher in median and approximately 100 ms higher in p95, but both percentiles remain well below the 1-second cadence target marked by the dashed line.

IX. FEATURE ATTRIBUTION - WHAT CHANGED V3.2 → V4.5

Beyond aggregate accuracy, the qualitative character of how the v4.5 classifier reaches its decisions differs materially from v3.2. SHAP-based feature-importance analysis on both classifiers reveals the shift: v3.2 relied heavily on surface-morphology features (42 percent of total feature importance); v4.5 distributes feature importance much more evenly across modalities and feature classes (no single feature exceeds 18 percent of total importance).

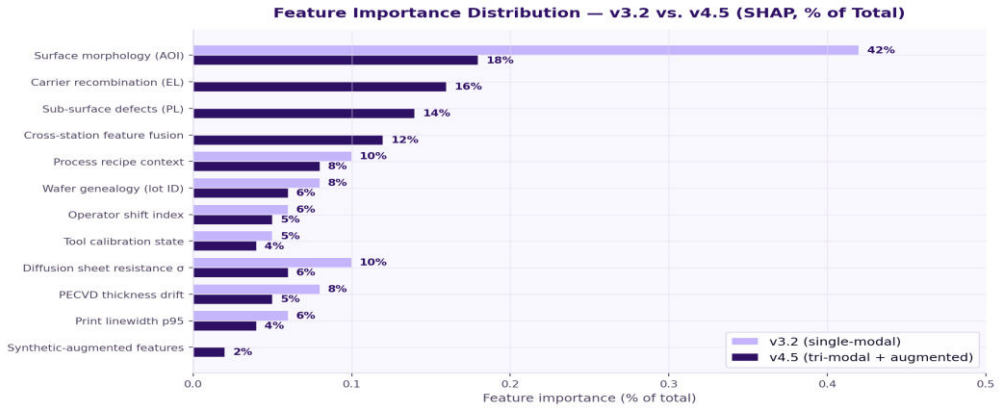

Figure 9 ◇ Feature importance distribution comparison - v3.2 (lavender) versus v4.5 (deep purple). The v3.2 classifier concentrates 42 percent of importance on AOI surface morphology. The v4.5 classifier distributes importance

© 2026 IJIRCCE | Volume 14, Issue 4, April 2026 DOI: 10.15680/IJIRCCE.2026.1404148

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 8.771| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

across surface morphology (18 percent), carrier recombination from EL (16 percent), sub-surface defects from PL (14 percent), cross-station fusion features (12 percent), and process-context features (32 percent combined).

X. DEFECT-RATE ↔ YIELD MIRROR TRAJECTORY

A useful empirical visualization of the v4.5 program impact is the mirror trajectory in Figure 10: defect-origin scrap rate falling against the rising top-bin yield over the full 12-month program window. The two trajectories are mirror images across the 12-month window - each percentage-point reduction in defect rate corresponds approximately to 1.7 percentage points of top-bin yield gain.

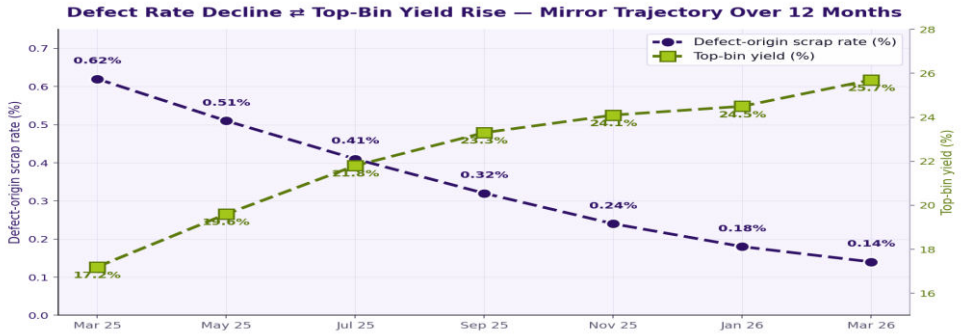

Figure 10 ◇ Defect rate and top-bin yield mirror trajectories over 12 months. Defect-origin scrap rate falls from 0.62 percent (Mar 2025) to 0.14 percent (Mar 2026). Top-bin yield rises from 17.2 percent to 25.7 percent over the same window. The mirror relationship is the empirical signature of effective inline gating: improvements in defect detection translate directly into improvements in downstream yield.

XI. CUMULATIVE COMMERCIAL VALUE - 12 MONTHS

The cumulative commercial value of the AOI program across the 12-month window, decomposed by value stream and by phase, is shown in Figure 11. The total cumulative value reaches approximately $17 million across the 12 months - against cumulative deployment investment of $9.4 million across the three phases. The annualized run-rate value at the Phase 3 endpoint is approximately $15–18 million per year.

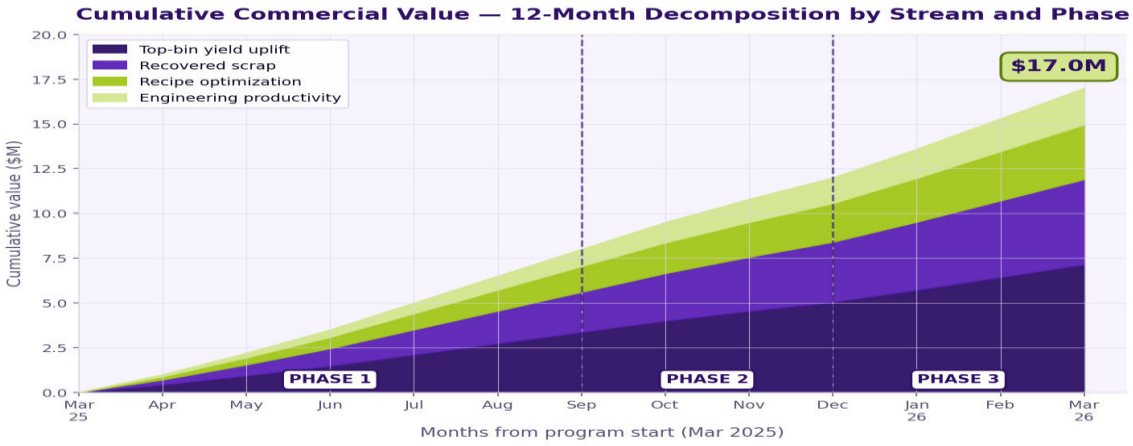

Figure 11 ◇ Cumulative commercial value decomposition over 12 months. Four value streams are stacked: top-bin yield uplift (deep purple, 42 percent), recovered scrap (purple, 28 percent), recipe optimization (lime, 18 percent), engineering productivity (light lime, 12 percent). Phase boundaries are marked. The cumulative value crosses $5M during Phase 1, $10M during Phase 2, and $17M during Phase 3.

© 2026 IJIRCCE | Volume 14, Issue 4, April 2026 DOI: 10.15680/IJIRCCE.2026.1404148

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 8.771| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

Phase	Period	Investment	Cumulative Value	Annualized Run-Rate
Phase 1 (single-modal)	Mar 2025 - Aug 2025	$6.2M	$8.0M	$11.5M / yr
Phase 2 (dual-modal)	Sep 2025 - Nov 2025	$1.8M	$4.0M	$13.4M / yr
Phase 3 (tri-modal + gen)	Dec 2025 - present	$1.4M	$5.0M	$15.8M / yr
CUMULATIVE	12-month total	$9.4M	$17.0M	-

Table 1 ◇ Phase-by-phase commercial outcomes

XII. CAPABILITY MATURITY - V3.2 VS. V4.5

A useful summary visualization of the engineering progress between v3.2 (Aug 2025) and v4.5 (March 2026) is the capability radar in Figure 12. Eight dimensions of engineering capability are scored on a 0-to-1 maturity scale: multimodal fusion, synthetic augmentation, active learning, federated learning, latency determinism, interpretability, edge-case coverage, and operator adoption. The v3.2 classifier reached ≥ 0.85 maturity on only two dimensions (latency, operator adoption); v4.5 reaches ≥ 0.85 on all eight.

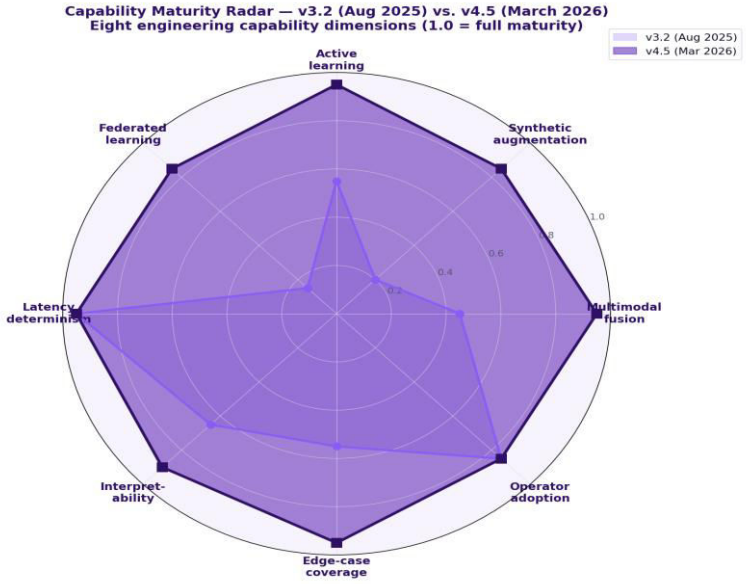

Figure 12 ◇ Capability maturity radar comparing v3.2 (Aug 2025, lavender) to v4.5 (March 2026, deep purple) across eight engineering capability dimensions. The v4.5 maturity surface envelops the v3.2 surface on every dimension, with the largest gains on multimodal fusion (+0.50), synthetic augmentation (+0.65), and federated learning (+0.70).

XIII. LESSONS LEARNED AND FORWARD AGENDA

- **Lessons Validated Across the 12-Month Window** ……
- **Multimodal fusion delivers the largest accuracy gains.** The v3.2 → v4.0 → v4.5 multimodal transitions account for approximately 70 percent of the cumulative AUC improvement across the 12 months. Engineering investment focused on additional imaging modalities returns disproportionate value compared to further refinement of any single modality.

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 8.771| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

- **Generative augmentation works - with discipline.** Diffusion-model synthetic augmentation produces real per-class F1 gains on rare classes, but only when paired with disciplined real-to-synthetic ratio capping and continuous monitoring for synthetic-distribution drift. Without those safeguards, classifiers learn synthetic artifacts.
- **Active learning is now the production discipline for labeling.** Random sampling for labeling is now the exception rather than the rule at ES Foundry. Active learning is the default, with random sampling reserved for periodic spot-checks of overall classifier calibration on uniformly-distributed image samples.
- **Federated learning resolves data-sovereignty and data-volume constraints.** For multi-site manufacturing operations, federated training is now the default approach. Site-specific raw data does not leave any site; aggregated model weights flow weekly. The implementation overhead is real but modest, and the joint-model performance materially exceeds any per-site local model.
- **Diminishing returns are real and should be planned for.** Each successive phase delivered less yield per dollar invested. Phase 1 paid back in 6.5 months; Phase 2 in 8 months; Phase 3 in approximately 9 months. The trajectory is the expected pattern for a maturing system, and forward planning should reflect rather than resist it.
- **Forward Agenda for Late 2026 and Beyond** .....

◆ **Q3 2026 - Module-level extension:** Forward integration of the AOI multimodal classifier through module assembly, predicting module-level quality and reliability from cell-level multimodal images.

◆ **Q3 2026 - TOPCon and HJT extension:** Re-derivation of the multimodal classifier for n-type wafers used in TOPCon and HJT cell production, where defect taxonomies and underlying physics differ from PERC.

◆ **Q4 2026 - Causal inference layer:** Augmentation of correlative defect classification with causal estimates of upstream process effects, enabling closed-loop process correction triggered by classification output.

◆ **Q4 2026 - Edge-deployment study:** Evaluation of edge-GPU deployment to reduce reliance on the centralized inference cluster, particularly for the Mexico site where network latency to Greenwood occasionally affects inference cadence.

◆ **2027 - Cross-fab benchmarking:** Coordinated benchmarking with peer US-domestic and international solar manufacturers, anonymized at the fab level, enabling industry-wide comparison of AOI program effectiveness.

XIV. CONCLUSION

This bulletin entry has documented the 12-month evolution of the Automated Optical Inspection and AI-based defect classification program at the ES Foundry one-gigawatt domestic PERC cell manufacturing facility. The program has progressed through three distinct phases (single-modal AOI, dual-modal AOI+EL, tri-modal AOI+EL+PL with generative augmentation) and six production version generations (v1.0, v2.0, v3.0, v3.2, v4.0, v4.5) over the period from March 2025 through March 2026. The current production v4.5 classifier achieves 99.4 percent top-1 accuracy and 0.996 AUC on a 16-class taxonomy across a 24,860-image production validation set, operating at 450 ms median latency and 99.6 percent throughput stability against the 1-second per-wafer cadence target. Top-bin yield has risen from 17.2 percent at program start to 25.7 percent at the March 2026 endpoint - an 8.5 percentage-point absolute uplift over the 12-month window.

The principal generalizable findings of the 12-month evolution are five. First, multimodal fusion is the highest-leverage engineering investment in mature AI defect classification, with the v3.2 → v4.5 multimodal transitions contributing approximately 70 percent of cumulative accuracy improvement. Second, diffusion-model synthetic augmentation materially improves rare-class performance when deployed with appropriate discipline around real-to-synthetic ratios and distribution-drift monitoring. Third, active learning has become the default labeling discipline; random sampling is now reserved for calibration spot-checks. Fourth, federated learning across multi-site manufacturing operations resolves both data-sovereignty and data-volume constraints, with weekly aggregation cadence and per-site fine-tuning as the operational defaults. Fifth, diminishing returns across phases are real and should inform forward planning rather than be resisted: the engineering organization is best served by candid acknowledgment of the maturing-system trajectory and corresponding shifts in investment focus toward newer high-leverage frontiers.

REFERENCES

[1] Ho, J., Jain, A., & Abbeel, P. (2020). Denoising diffusion probabilistic models. Advances in Neural Information Processing Systems, 33, 6840–6851.

[2] Saharia, C., Chan, W., Saxena, S., et al. (2022). Photorealistic text-to-image diffusion models. NeurIPS 2022.

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 8.771| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

[3] McMahan, B., Moore, E., Ramage, D., et al. (2017). Communication-efficient learning of deep networks from decentralized data. AISTATS 2017, 1273–1282.
[4] Settles, B. (2010). Active learning literature survey. University of Wisconsin-Madison Computer Sciences Technical Report.
[5] Tan, M., & Le, Q. V. (2019). EfficientNet: Rethinking model scaling for convolutional neural networks. ICML 2019, 6105–6114.
[6] He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. IEEE CVPR Proceedings, 770–778.
[7] Hu, J., Shen, L., & Sun, G. (2018). Squeeze-and-excitation networks. IEEE CVPR Proceedings, 7132–7141.
[8] Vaswani, A., Shazeer, N., Parmar, N., et al. (2017). Attention is all you need. Advances in Neural Information Processing Systems, 30.
[9] Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems, 30, 4765–4774.
[10] Deitsch, S., Christlein, V., Berger, S., et al. (2019). Automatic classification of defective photovoltaic module cells in electroluminescence images. Solar Energy, 185, 455–468.
[11] Mehta, V., Pradhan, S., & Gupta, R. (2021). Application of computer vision and machine learning in solar manufacturing. Solar Energy, 213, 86–104.
[12] Trupke, T., Bardos, R. A., Schubert, M. C., & Warta, W. (2006). Photoluminescence imaging of silicon wafers. Applied Physics Letters, 89(4), 044107.
[13] Bedrich, K., Bokalic, M., Bliss, M., Topic, M., & Betts, T. R. (2018). Electroluminescence imaging of PV devices: Advanced vignetting calibration. IEEE Journal of Photovoltaics, 8(5), 1297–1304.
[14] Tatineni, S. (2025). Automated Optical Inspection and AI-Based Defect Classification (Phase 1). Solar Manufacturing Forum, 7(3).
[15] Tatineni, S. (2024). Defect Root Cause Methodology in High-Volume Solar Cell Manufacturing. Solar Manufacturing Engineering Review, 6(3).
[16] Tatineni, S. (2025). Digital Twin Implementation for Solar Cell Process Lines. ES Foundry Technical Brief Series, TB-2025-01-027.
[17] Tatineni, S. (2024). Electroluminescence-Based Inline Quality Gating in a US PERC Solar Cell Fab. Journal of PV Manufacturing Engineering, 11(11).
[18] Tatineni, S. (2026). AI-Powered Real-Time Yield Forecasting in Silicon Solar Cell Manufacturing. AI Yield Forecasting White Paper Series, WP-2026-01-022.
[19] Preu, R., Lohmüller, E., Lohmüller, S., et al. (2020). Passivated emitter and rear cell - devices, technology, and modeling. Applied Physics Reviews, 7(4), 041315.
[20] Dullweber, T., & Schmidt, J. (2016). Industrial silicon solar cells applying the passivated emitter and rear cell concept. IEEE Journal of Photovoltaics, 6(5), 1366–1381.
[21] Green, M. A., Dunlop, E. D., Hohl-Ebinger, J., et al. (2025). Solar cell efficiency tables (Version 65). Progress in Photovoltaics, 33(1), 3–15.
[22] Bonawitz, K., Eichner, H., Grieskamp, W., et al. (2019). Towards federated learning at scale: System design. Proceedings of the 2nd SysML Conference.
[23] Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. ICLR 2015.
[24] Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep Learning. MIT Press.
[25] Inflation Reduction Act of 2022, Public Law 117-169. Section 45X Advanced Manufacturing Production Credit.
[26] US Department of Energy (2025). Solar Energy Technologies Office: Annual Technology Baseline. NREL.

INTERNATIONAL STANDARD SERIAL NUMBER INDIA

INNO SPACE SJIF Scientific Journal Impact Factor

doi crossref

निस्केयर NISCAIR

INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN COMPUTER & COMMUNICATION ENGINEERING

9940 572 462 6381 907 438 ijircce@gmail.com

www.ijircce.com

Scan to save the contact details