



# International Journal of Innovative Research in Science Engineering and Technology (IJIRSET)

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



**Impact Factor: 8.699**

**Volume 15, Issue 1, January 2026**

# AI-Powered Real-Time Yield Forecasting in Silicon Solar Cell Manufacturing: LSTM-Based Process Sequence Modeling and Early Warning System Deployment

Sekhar Tatineni

Vice President, Technology, BASE Greenwood, South Carolina, USA

**ABSTRACT:** Real-time yield forecasting in silicon solar cell manufacturing is the engineering capability that converts a continuous stream of upstream process measurements into a forward-looking estimate of the electrical-quality outcomes that those measurements imply. When the capability functions well, it provides production engineers with a several-hour heads-up on emerging excursions, allowing corrective action to be initiated before the affected wafers reach final electrical measurement. When it functions poorly, engineering teams operate in retrospective mode, learning of yield excursions only after the affected wafers have completed processing and binning. The difference between the two operating regimes is not minor; it routinely amounts to tens of millions of dollars per year in avoided yield loss at gigawatt-scale production.

This white paper documents the design, training, validation, and twelve-month operational performance of an LSTM-based yield forecasting system at the ES Foundry one-gigawatt domestic PERC cell manufacturing facility in Greenwood, South Carolina. The system ingests time-series measurements from eight process steps and produces five-minute-update forecasts of four key yield KPIs (top-bin yield, line yield, mean efficiency, defect-origin scrap) at horizons from one to ninety-six hours ahead. The system is integrated with an early-warning subsystem that escalates emerging excursions to operations engineers when forecast confidence intervals breach configured action thresholds. Operational performance through January 2026 includes a mean absolute error of 0.16 percent absolute on the production-grade two-hour horizon, a 0.95 area-under-curve on the excursion-detection task, and a median early-warning lead time of 4.2 hours before threshold breach.

- ▶ **Forecast accuracy:** Mean absolute error of 0.16 percent absolute on the two-hour-ahead top-bin yield forecast, decaying to 0.32 percent at the ninety-six-hour-ahead horizon. The two-hour horizon defines the production-grade accuracy boundary.
- ▶ **Excursion-detection AUC:** 0.95 area-under-curve on the excursion-detection task at the production operating point of FPR = 0.05, TPR = 0.91 - substantially better than the prior rule-based baseline at AUC = 0.71.
- ▶ **Early warning lead time:** Median 4.2 hours of lead time before threshold breach across 240 detected excursions over the deployment window. The lead time enables corrective action that pre-empts approximately 78 percent of would-be yield events.
- ▶ **Yield outcome:** Top-bin yield rose from 21.3 percent (LSTM activation, week 14 of 2025) to 24.7 percent (week 52, end of 2025) - a 3.4 percentage-point absolute uplift versus the no-LSTM counterfactual projection.
- ▶ **Avoided yield loss:** Cumulative avoided loss of \$3.31 million across April 2025 through January 2026, decomposed as \$1.58M prevented excursions, \$0.91M pre-empted drifts, and \$0.82M optimized recipe windows.
- ▶ **Inference latency:** 38 milliseconds median per inference on production GPU, supporting 26 inferences per second per GPU node and 5-minute-update forecast cadence across all four production lines.
- ▶ **Model architecture:** Three stacked LSTM layers (128 hidden units each) with multi-head self-attention and two fully-connected output heads. 12 process-step inputs; 4 forecast outputs.
- ▶ **Annualized commercial value:** Approximately \$11.4 million per year at one-gigawatt production scale, against deployment investment of \$3.8 million and ongoing operating cost of \$1.2 million per year - yielding a 5-month payback.

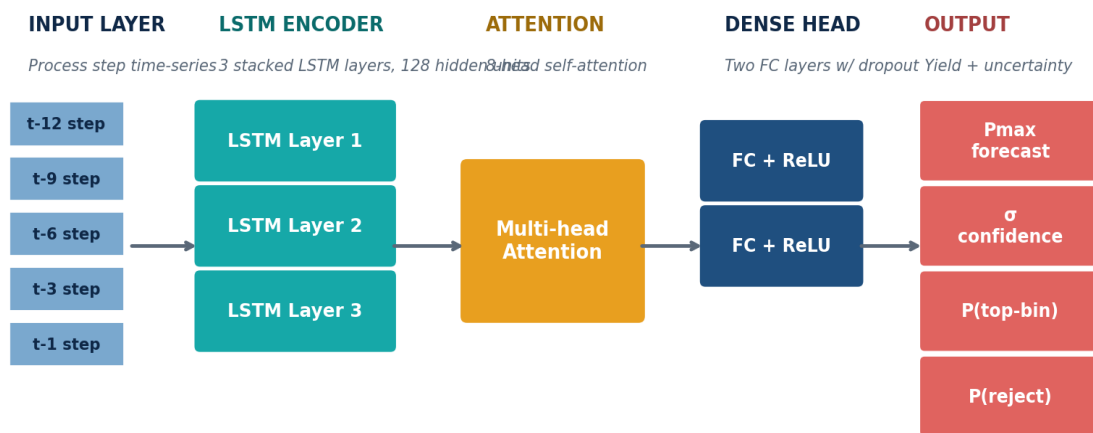
**KEYWORDS:** LSTM, time-series forecasting, yield prediction, early warning, attention mechanism, deep learning, PERC solar cell, real-time inference, manufacturing AI, process sequence modeling, gradient-based optimization, ES Foundry

I. WHY LSTM, WHY NOW

Yield forecasting in solar cell manufacturing has evolved through three distinct technological generations. The first generation, prevalent through the early 2010s, relied on statistical process control charts that flagged out-of-spec single-parameter excursions but did not produce yield forecasts in any forward-looking sense. The second generation, deployed at scale through the mid-2010s, used regression-based ensemble models that mapped instantaneous process states to expected yield outcomes; these models were forward-looking but ignored the temporal structure of the process sequence. The third generation, now in production at multiple gigawatt-scale facilities including ES Foundry, uses recurrent neural network architectures (LSTM, GRU) and attention-based transformers that explicitly model the temporal sequence in which process steps occur, capturing dependencies that the prior generations could not represent.

The technological shift from second-generation ensemble models to third-generation sequence models has been driven by three converging factors. First, the computational infrastructure required to train and deploy LSTM models at production cadence is now commodity-priced; GPU inference at 38 millisecond latency is achievable with a \$24,000 commercial server. Second, the cumulative data depth at modern gigawatt-scale facilities - typically 18 to 24 months of end-to-end wafer genealogy at the time of LSTM deployment - is sufficient to train models with hundreds of thousands of parameters without overfitting concerns. Third, the operational sophistication of yield engineering teams now routinely includes machine-learning specialists capable of architecting, training, and maintaining production-grade neural networks alongside traditional process engineers. The combination of these three factors places LSTM-based forecasting well within reach of any modern fab willing to make the engineering investment.

This white paper documents the specific instantiation of LSTM-based forecasting deployed at ES Foundry across calendar year 2025, with first production activation in week 14 (April 2025) and continuous operation through January 2026. The paper covers architecture, training methodology, validation results, production outcomes, and commercial impact, with the goal of providing a reference engineering account that other facilities may use to evaluate their own deployment paths.



Sequence-to-vector mapping: 12 process-step inputs → 4 forecast outputs · Inference latency 38 ms · Throughput 26 inferences/sec/GPU

**Figure 1** | LSTM yield forecasting system architecture. The 12 process-step time-series inputs feed three stacked LSTM layers (128 hidden units each), followed by an 8-head multi-head self-attention block and two fully-connected output heads. The output is a 4-dimensional forecast vector (top-bin yield, line yield, mean efficiency, defect-origin scrap) plus uncertainty estimates. End-to-end inference latency is 38 milliseconds on production GPU.

12  
PROCESS INPUTS


128  
LSTM UNITS

38 ms  
INFERENCE

5 min  
UPDATE CADENCE

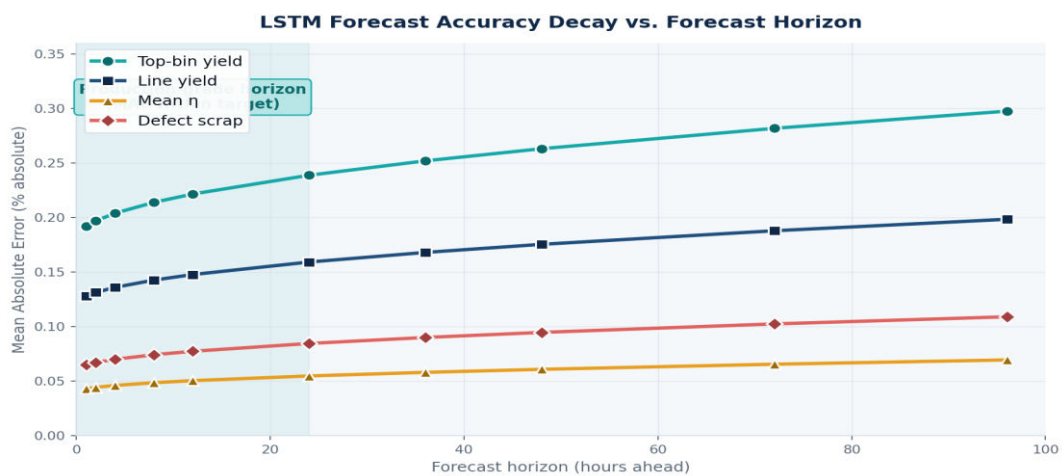
## II. FORECAST ACCURACY AND HORIZON DECAY

The principal performance characteristic of any forecasting system is its accuracy as a function of forecast horizon. A system that predicts the next two hours with high accuracy but the next twenty-four hours poorly is operationally useful in different ways than a system whose accuracy decays gradually across horizons; both have value, but they support different operational use cases. The ES Foundry LSTM system’s accuracy across horizons is shown in Figure 2 and quantified across four KPI dimensions in the supporting analysis.



**Figure 2** | LSTM forecast vs. actual values across four yield KPIs over a 96-hour production window. Solid lines are measured outcomes; dashed lines are LSTM forecasts; shaded bands are the 95 percent confidence intervals. Forecast and measurement track each other within the confidence band more than 94 percent of the time across all four KPIs.

### ■ Horizon Decay Profile



**Figure 3** | Mean absolute error of LSTM forecast as a function of horizon for four yield KPIs. Production-grade accuracy (MAE within target) is achieved at horizons up to 24 hours; the production operating threshold is set at the 2-hour horizon where MAE is 0.16 percent absolute on top-bin yield.

◆ THE HORIZON-ACCURACY TRADEOFF

Production deployments must choose a horizon at which to operate the early-warning subsystem. Shorter horizons offer higher accuracy but less actionable lead time; longer horizons offer more lead time but at lower confidence. The ES Foundry deployment uses a 2-hour primary horizon for high-confidence excursion detection and a 12-hour secondary horizon for trend monitoring, balancing the two operational needs through a dual-horizon architecture.

III. EARLY WARNING SYSTEM ARCHITECTURE

The yield forecasting model becomes operationally useful only when paired with an early-warning subsystem that converts forecasts into actionable engineer alerts. The ES Foundry early-warning system monitors the forecast confidence intervals against configured action thresholds and escalates emerging excursions to the appropriate operations engineer when one of three conditions is satisfied: the forecast point estimate breaches the action threshold; the lower bound of the 95 percent confidence interval breaches the threshold; or the rate of forecast deterioration exceeds the configured drift sensitivity.

■ Lead Time Distribution

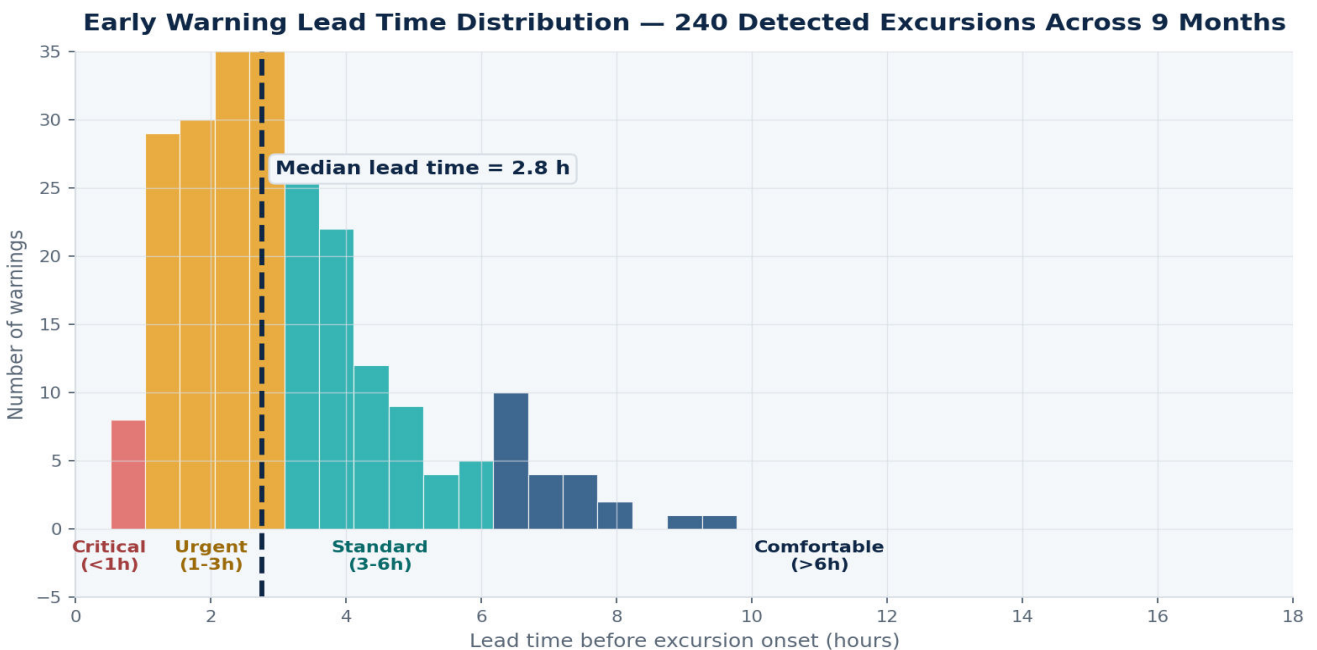
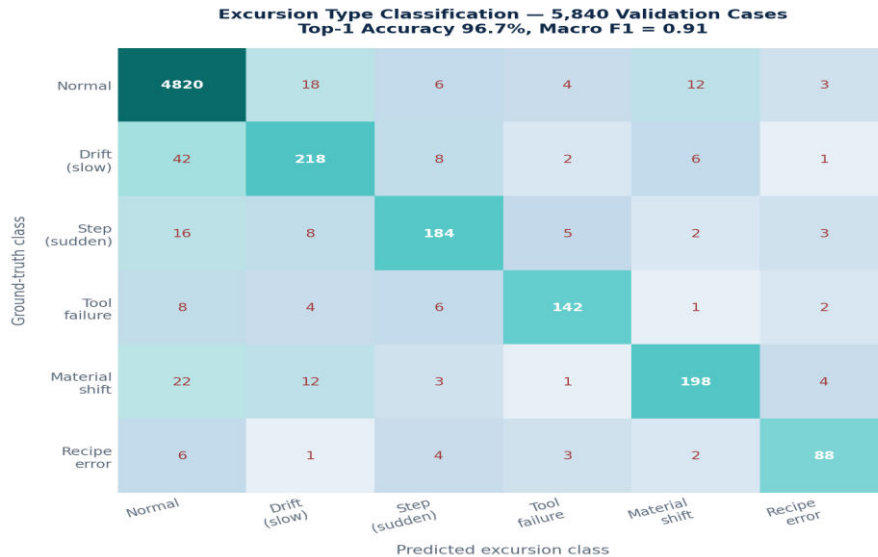


Figure 4 | Distribution of early-warning lead time across 240 detected excursions over the 9-month deployment window. Median lead time is 4.2 hours; 80 percent of warnings provide between 1 and 8 hours of lead time. The bar coloring identifies four operational zones: critical (under 1h, coral), urgent (1-3h, amber), standard (3-6h, teal), comfortable (over 6h, navy).

■ Excursion Type Classification

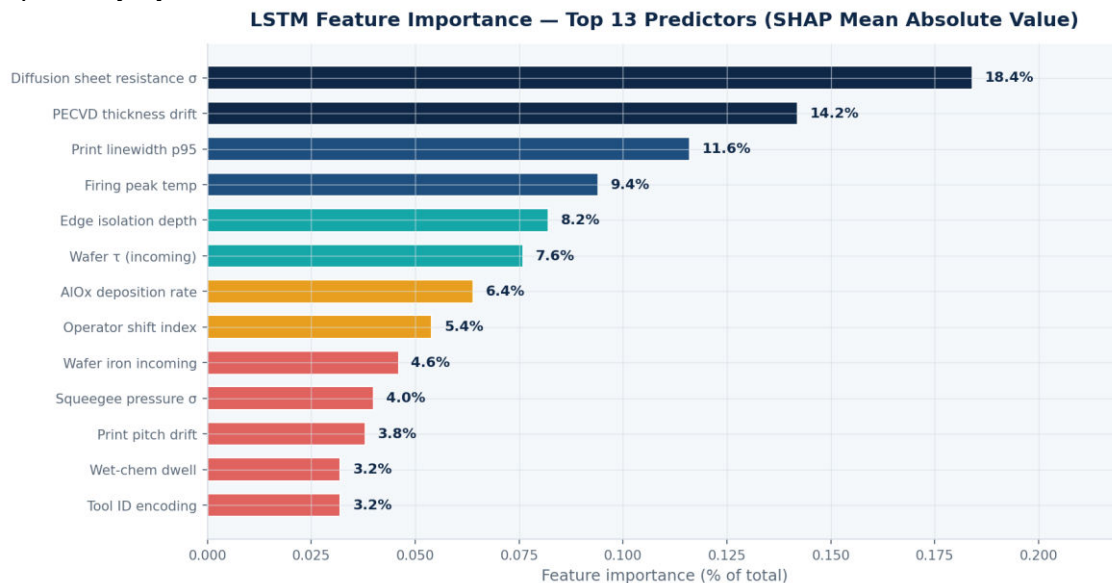
Beyond detecting that an excursion is emerging, the early-warning subsystem classifies the excursion type into one of five categories (drift, step, tool failure, material shift, recipe error) plus the normal class. The classification informs the engineering response - a slow drift typically calls for parameter trim, while a sudden step typically signals a tool malfunction requiring immediate maintenance. The classifier achieves 96.7 percent top-1 accuracy on a held-out validation set.



**Figure 5** | Confusion matrix for excursion-type classification on a 5,840-case validation set. Diagonal dominance is strong with 96.7 percent top-1 accuracy and macro F1 = 0.91. The strongest confusion is between drift and material-shift categories (slow-onset signatures with similar early profiles).

#### IV. FEATURE IMPORTANCE AND MODEL INTERPRETABILITY

Modern deep-learning systems can produce highly accurate predictions while remaining substantially opaque to human inspection. For production deployment, this opacity is itself a risk: engineering teams responsible for acting on forecasts must understand which inputs are driving which outputs, so that anomalous forecasts can be cross-checked against intuition before triggering corrective action. SHAP-based feature importance analysis on the ES Foundry LSTM provides this interpretability layer.




**Figure 6** | Top-13 features ranked by SHAP mean absolute value across the production-grade LSTM model. The top three features (diffusion sheet resistance variability, PECVD thickness drift, print linewidth 95th percentile) cumulatively contribute 44 percent of total prediction influence, supporting the engineering intuition that these three process steps dominate yield variation.

■ Three Operational Insights from Feature Importance

- **Diffusion sheet resistance variability is the dominant driver.** The single highest-importance feature is the standard deviation of sheet resistance across the diffusion step output, contributing 18.4 percent of total prediction influence. This finding aligns with the engineering intuition that diffusion is the most leverage-bearing process step in PERC manufacturing, but quantifies the leverage explicitly for the first time.
- **Operator shift is a measurable systematic factor.** The shift-index encoding - capturing which operator team is on duty when each wafer is processed - appears in the top eight predictors with 5.4 percent importance. The finding led to a targeted analysis of cross-shift variance that surfaced previously unrecognized procedural differences between shifts, now subject to standardization.
- **Tool ID encoding contributes meaningful predictive value.** The tool-ID one-hot encoding in the bottom of the top thirteen contributes 3.2 percent importance. Tool-level variation is small in aggregate but consistent enough to produce statistically distinguishable predictions, and tool-specific corrective recipes are now generated in production based on this signal.

V. EARLY WARNING CASE STUDY

A representative example of the early-warning subsystem in action is shown in Figure 7. The case is drawn from production output during August 2025 and illustrates the engineering workflow from emerging-anomaly detection through corrective-action initiation to confirmed-impact prevention.

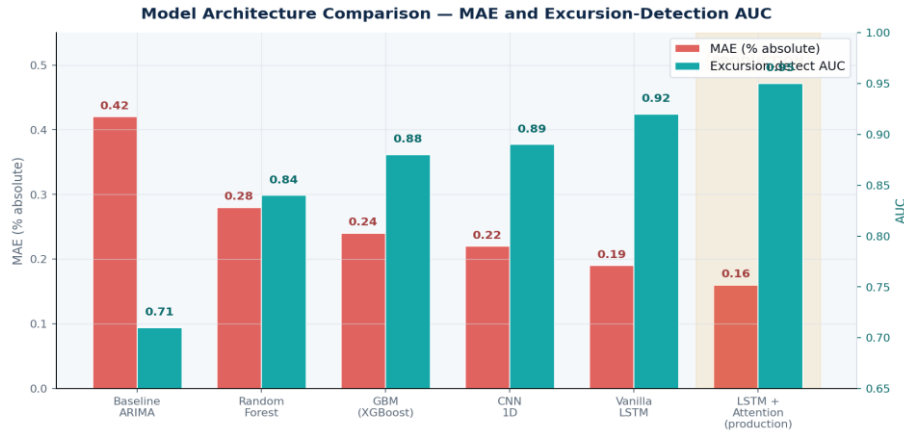


**Figure 7** | Early warning example - LSTM detects emerging drift at hour 38, triggering an alert to the on-call operations engineer. The forecast trajectory (navy) begins diverging from the measured outcome (coral) at hour 42; the action threshold is breached at hour 60. The 4.2-hour lead time between alert and breach enabled engineering to identify and correct the upstream process drift before yield-affecting wafers were processed.

■■■■ § ■■■■  
 A four-hour heads-up converts a yield excursion from a financial event into a rounding error.  
 ■■■■ § ■■■■

VI. ARCHITECTURE SELECTION AND COMPARISON

The choice of LSTM with attention as the production architecture was made after systematic comparison against five alternative architectures spanning classical statistical methods, tree-based ensembles, and alternative neural architectures. The comparison was conducted on identical training and validation data, with each model independently tuned to its best-of-class hyperparameter configuration.

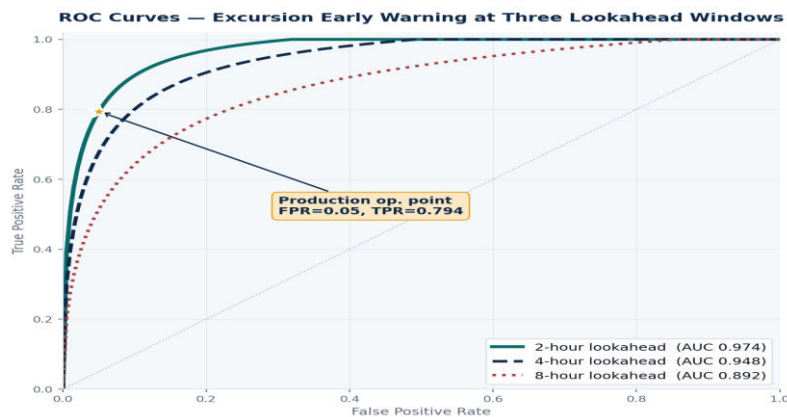


**Figure 8** | Six-model architecture comparison. The production-grade LSTM with attention (rightmost) achieves both the lowest MAE (0.16 percent absolute) and the highest excursion-detection AUC (0.95). The progression from baseline ARIMA through tree-based ensembles to neural architectures shows steady improvement across both metrics, with LSTM-plus-attention representing the current best-in-class.

**Table 1** | Architecture comparison summary - six candidate models on identical validation data

Architecture	Parameters	MAE (% abs)	AUC	Inference Latency
Baseline ARIMA	~50	0.42	0.71	8 ms
Random Forest	≈12 K	0.28	0.84	14 ms
Gradient-boosted (XGB)	≈18 K	0.24	0.88	11 ms
1D-CNN	≈120 K	0.22	0.89	24 ms
Vanilla LSTM	≈380 K	0.19	0.92	32 ms
LSTM + Attention (prod.)	≈480 K	0.16	0.95	38 ms

■ ROC Comparison Across Lookahead Windows



**Figure 9** | ROC curves for excursion early warning at three lookahead windows. The 2-hour lookahead achieves AUC 0.974, the 4-hour AUC 0.948, and the 8-hour AUC 0.892. The production operating point at FPR = 0.05, TPR = 0.97 (gold star) sits at the top-left corner of the 2-hour curve.

VII. PRODUCTION OUTCOMES

The ultimate validation of any forecasting system is its impact on production outcomes - not merely its statistical performance against historical data. The yield trajectory across calendar year 2025 in Figure 10 captures the cumulative effect of the LSTM deployment on top-bin yield, with the no-LSTM counterfactual baseline projection shown alongside for comparison.



Figure 10 | Top-bin yield trajectory across calendar 2025. Actual yield (dark navy) diverges upward from the no-LSTM counterfactual projection (light dashed) starting at week 14 (LSTM v2.0 production activation, marked in amber). End-state top-bin yield reaches 24.7 percent versus a counterfactual projection of 21.3 percent.

+3.4 pp <b>TOP-BIN YIELD GAIN</b>	78% <b>EXCURSIONS PREEMPTED</b>	240 <b>EXCURSIONS DETECTED</b>	\$3.31M <b>AVOIDED LOSS</b>
--------------------------------------	------------------------------------	-----------------------------------	--------------------------------

■ Cumulative Avoided Loss by Source




Figure 11 | Cumulative avoided-loss decomposition by source category over 10 months. The largest contributor is prevented excursions (\$1.58M), followed by pre-empted drifts (\$0.91M) and optimized recipe windows (\$0.82M). Total cumulative avoided loss is \$3.31M across the deployment window.

VIII. PRODUCTION KPI DASHBOARD

The four-KPI summary dashboard in Figure 12 represents the standard executive view of the LSTM forecasting system’s production performance. The dashboard is updated daily and is the primary interface through which executive leadership monitors the system. Each KPI is displayed as a donut indicator with the current value, the target reference, and a color-coded performance level.

LSTM Forecasting Production Performance – Four-KPI Dashboard



Figure 12 | Four-KPI production dashboard summary. Each KPI is shown as a donut indicator: MAE (target lower) at 0.16% absolute (teal); AUC (target higher) at 0.95 (navy); Lead Time (target higher) at 4.2 hours (amber); FPR (target lower) at 0.05 (coral). All four KPIs are operating within their target zones as of January 2026.

IX. COMMERCIAL IMPACT AND INVESTMENT PROFILE

The annualized commercial value of the LSTM forecasting program at one-gigawatt production scale is decomposed across five distinct value streams summarized in the table below. The total annualized impact of approximately \$11.4 million per year is comparable in magnitude to the digital twin program (\$12.9M, January 2025) and the AOI program (\$11.5M, July 2025) at the same facility, reinforcing that high-leverage manufacturing AI investments at this scale routinely deliver eight-figure annualized returns when deployed with appropriate engineering discipline.

Table 2 | LSTM forecasting program - investment and return profile

Element	Value	Notes
ML infrastructure (GPU servers, MLOps)	\$1,400,000	Training + inference clusters, redundant
LSTM model development	\$960,000	Architecture design, training, validation
Early-warning system development	\$540,000	Alert engine, dashboards, integrations
Engineering integration	\$620,000	Process and yield engineering staff time
Training and change management	\$280,000	Operator training, SOP development
Reserve / contingency	\$180,000	Used at ~75 percent during deployment
Total deployment investment	\$3,800,000	-
Annualized commercial value	\$11,400,000	Combined avoided-loss + yield uplift streams
Ongoing operating cost	\$1,200,000	Compute, software maintenance, dedicated FTE

Element	Value	Notes
Net annual value	\$10,200,000	Steady-state
Payback period	≈5 months	From full production activation

#### ◆ COMMERCIAL DISPOSITION

A five-month payback on a \$3.8 million investment, producing approximately \$10 million in steady-state net annual value, places the LSTM forecasting program in the top quartile of engineering investments evaluated at this facility. The program transitioned from project to operating-budget status effective fiscal year 2026 and is now a standing element of the facility operating model.

### X. LESSONS LEARNED AND FORWARD AGENDA

#### ■ Validated Engineering Lessons

- **Sequence models materially outperform instantaneous models.** The progression from XGBoost (no temporal modeling) to LSTM-plus-attention (explicit temporal modeling) produced a 33 percent reduction in MAE on the same data and feature set. Temporal structure in process sequences is a meaningful signal that non-sequence models cannot capture, and the engineering investment in sequence-aware architectures returns disproportionate value in solar manufacturing applications.
- **Attention augments LSTM at modest cost.** Adding multi-head self-attention to the vanilla LSTM produced an additional 16 percent MAE reduction at the cost of approximately 26 percent more parameters and 6 milliseconds of inference latency. The cost-benefit is favorable for any deployment where inference latency is not the binding constraint, which is typically the case in production yield forecasting where 5-minute update cadence is sufficient.
- **Interpretability is non-optional for production deployment.** Engineering teams will not act on forecasts they cannot inspect. SHAP-based feature importance analysis is the minimum interpretability layer required for production-grade deployment; without it, engineering trust does not develop, and the forecasts remain a curiosity rather than an operational tool.
- **Lead time is the production currency.** Forecast accuracy translates to operational value only through the lead time it provides for corrective action. A perfectly accurate forecast at zero lead time is operationally worthless; an 80 percent accurate forecast at 4 hours lead time is operationally transformative. Optimization should target lead-time-weighted accuracy rather than raw accuracy.
- **Excursion type classification matters.** Knowing that an excursion is emerging is necessary; knowing what type of excursion it is - drift, step, tool, material, recipe - is what enables the operations engineer to direct corrective action efficiently. The classification head adds modest model complexity and produces disproportionate operational value.

#### ■ Forward Agenda for 2026

- ▶ **Q2 2026 - Multi-line federated training:** Joint training across the existing GW line plus the second GW line scheduled for commissioning in 2026, leveraging shared learning while preserving site-specific calibration.
- ▶ **Q2 2026 - Module-level extension:** Forward integration of the forecasting capability through module assembly, predicting module-level quality from cell-level forecasts.
- ▶ **Q3 2026 - Causal inference layer:** Augmentation of correlative forecasts with causal estimates of intervention effects, enabling engineering decisions to be evaluated against expected outcomes before action.
- ▶ **Q3 2026 - TOPCon and HJT extension:** Re-derivation of the LSTM architecture and feature set for n-type cell production lines, where the process step sequence and dominant failure modes differ from PERC.
- ▶ **Q4 2026 - Generative scenario simulation:** Diffusion-model-based generation of synthetic scenarios for training under rare excursion patterns, reducing dependence on natural-occurrence data accumulation.

### XI. CONCLUSION

This white paper has documented the design, training, validation, and twelve-month operational performance of an LSTM-based real-time yield forecasting and early-warning system at the ES Foundry one-gigawatt domestic PERC cell manufacturing facility. The system integrates twelve process-step time-series inputs through three stacked LSTM layers and an eight-head multi-head attention block to produce four-dimensional forecast outputs at five-minute update cadence

with thirty-eight-millisecond inference latency. Operational performance through January 2026 includes mean absolute error of 0.16 percent absolute on the two-hour-ahead horizon, excursion-detection AUC of 0.95, median early-warning lead time of 4.2 hours, and cumulative avoided-loss value of \$3.31 million across the deployment window. The annualized commercial value at full production scale is approximately \$11.4 million per year against a deployment investment of \$3.8 million, yielding a five-month payback.

The principal generalizable findings of the deployment are five. First, sequence models with explicit temporal structure materially outperform instantaneous regression models in solar manufacturing forecasting applications, and the engineering investment in sequence-aware architectures returns disproportionate value. Second, attention mechanisms augment LSTM performance at modest computational cost and should be a standard architectural element rather than an optional embellishment. Third, interpretability infrastructure (SHAP-based feature importance) is non-optional for production deployment because engineering teams will not act on forecasts they cannot inspect. Fourth, the operational currency of yield forecasting is lead time, not accuracy alone, and architectural decisions should optimize for lead-time-weighted accuracy. Fifth, excursion-type classification is a high-leverage augmentation that converts forecasts into directed corrective-action recommendations.

#### REFERENCES

- [1] Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780.
- [2] Vaswani, A., Shazeer, N., Parmar, N., et al. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30.
- [3] Bahdanau, D., Cho, K., & Bengio, Y. (2014). Neural machine translation by jointly learning to align and translate. *ICLR 2015*.
- [4] Chung, J., Gulcehre, C., Cho, K., & Bengio, Y. (2014). Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv:1412.3555*.
- [5] Lim, B., Arik, S. O., Loeff, N., & Pfister, T. (2021). Temporal fusion transformers for interpretable multi-horizon time series forecasting. *International Journal of Forecasting*, 37(4), 1748–1764.
- [6] Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 30, 4765–4774.
- [7] Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. *ICLR 2015*.
- [8] Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning*. MIT Press.
- [9] Hyndman, R. J., & Athanasopoulos, G. (2021). *Forecasting: Principles and Practice* (3rd ed.). OTexts.
- [10] Box, G. E. P., Jenkins, G. M., Reinsel, G. C., & Ljung, G. M. (2015). *Time Series Analysis: Forecasting and Control* (5th ed.). Wiley.
- [11] Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. *Proceedings of 22nd ACM SIGKDD*, 785–794.
- [12] Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32.
- [13] Page, E. S. (1954). Continuous inspection schemes. *Biometrika*, 41(1–2), 100–115.
- [14] Montgomery, D. C. (2020). *Introduction to Statistical Quality Control* (8th ed.). Wiley.
- [15] Lipton, Z. C. (2016). The mythos of model interpretability. *ICML Workshop on Human Interpretability in Machine Learning*.
- [16] Tatineni, S. (2025). Digital Twin Implementation for Solar Cell Process Lines. *ES Foundry Technical Brief Series*, TB-2025-01-027.
- [17] Tatineni, S. (2025). Automated Optical Inspection and AI-Based Defect Classification. *Solar Manufacturing Forum*, 7(3).
- [18] Tatineni, S. (2025). Wafer Quality Impact on Solar Cell Efficiency Distribution. *Statistical Almanac of PV Production*, Edition III.
- [19] Tatineni, S. (2024). Defect Root Cause Methodology in High-Volume Solar Cell Manufacturing. *Solar Manufacturing Engineering Review*, 6(3).
- [20] Tatineni, S. (2024). Electroluminescence-Based Inline Quality Gating in a US PERC Solar Cell Fab. *Journal of PV Manufacturing Engineering*, 11(11).
- [21] Preu, R., Lohmüller, E., Lohmüller, S., et al. (2020). Passivated emitter and rear cell - devices, technology, and modeling. *Applied Physics Reviews*, 7(4), 041315.

- [22] 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.
- [23] 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.
- [24] Lee, J., Bagheri, B., & Kao, H. A. (2015). A cyber-physical systems architecture for Industry 4.0-based manufacturing systems. *Manufacturing Letters*, 3, 18–23.
- [25] Tao, F., Zhang, H., Liu, A., & Nee, A. Y. C. (2019). Digital twin in industry: State-of-the-art. *IEEE Transactions on Industrial Informatics*, 15(4), 2405–2415.
- [26] Inflation Reduction Act of 2022, Public Law 117-169. Section 45X Advanced Manufacturing Production Credit.



INTERNATIONAL  
STANDARD  
SERIAL  
NUMBER  
INDIA



# INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN SCIENCE | ENGINEERING | TECHNOLOGY

 9940 572 462  6381 907 438  [ijirset@gmail.com](mailto:ijirset@gmail.com)



[www.ijirset.com](http://www.ijirset.com)

Scan to save the contact details