Proceedings of the 3rd International Conference on Machine Learning and Automation
DOI: 10.54254/2755-2721/2025.28669

Physics-Regularized Self-Supervised Anomaly Detection for Semiconductor Tools with Digital Twin Guidance

Min Yin

University of California, Berkeley, USA
gmiayinc@gmail.com

Abstract. Unplanned stoppages in semiconductor tools remain a persistent limiter of throughput and yield, a situation partly sustained by monitors that rely on dense labels or rule sets that do not travel well across recipes and tools. We study a digital-twin-driven framework that learns a compact health representation from multiscale telemetry by self-supervised objectives and regularizes it with differentiable constraints drawn from mass balance, thermal–RF coupling, and vacuum dynamics; anomaly evidence is then fused with process, environmental, and maintenance logs so that alerts arrive with context and with a plausible operational hypothesis. Orchestrated with DolphinScheduler or Airflow, the pipeline coordinates ingestion, training, streaming inference, lineage, and review to align analytics with change control and auditability. Development was deliberately iterative rather than linear: label sparsity and timestamp drift pushed us toward cycle-aware alignment; twin mis-specification in edge regimes required residual diagnostics and parameter re-estimation; population shift prompted conformal calibration and sequential testing. On production-like etch and deposition traces, we observe earlier detection under fixed alert budgets and extensions in lead time that appear to improve MTBF and OEE to some extent, together with indications of lower service effort and energy use. Alternative explanations, including facility subsystems or undocumented operator interventions, cannot be excluded, which suggests that further research is needed on causal attribution, cross-site transfer, and adaptive twin updating.

Keywords: Predictive Maintenance, Digital Twin, Self-Supervised Learning, Semiconductor Equipment, Anomaly Detection

1. Introduction

Advanced semiconductor fabrication operates within narrow process windows where small and gradually accumulating deviations in chamber pressure, thermal balance, gas delivery, radio frequency coupling, or vacuum integrity can propagate across lots and shifts, producing unplanned stoppages that depress overall equipment effectiveness and shorten mean time between failures while recovery sequences add cost and may, to some extent, compromise yield [1]. Deployed monitors still lean on handcrafted rules, statistical surveillance within fault detection and classification, and supervised classifiers trained on historical incidents; such instruments remain useful under stable regimes, yet they assume dense and timely labels that fabs rarely possess, travel poorly across heterogeneous recipe portfolios and tool vintages, and are vulnerable to drift and benign operational changes that blur the line between nuisance variation and genuine degradation. Considering these pressures, a synthesis that treats physics knowledge, learned representation, and operational context as coequal sources of evidence appears necessary [2].

The research gap has several layers. Digital twins and data-driven encoders are often developed in parallel without joint objectives, which leaves learned features free to internalize correlations that are inconsistent with equipment dynamics, whereas high-fidelity twins that excel in design studies are impractical for online maintenance and seldom shape the feature spaces used for detection. Process, environmental, and maintenance logs contain causal signatures of recipe updates and chamber state transitions, yet they are commonly consulted after failures rather than during learning and inference, which weakens a detector's ability to discount benign shifts. Equally important is calibrated uncertainty. Without explicit uncertainty that is exposed to operators, alert policies become brittle, review turns protracted, and change control drifts from procedure to negotiation. Taken together, these limits suggest that predictive maintenance will benefit from methods that are physics-regularized, log-aware, and uncertainty-calibrated while remaining compatible with production scheduling, lineage, and audit practice [3].

This study explores such a direction. A grey-box digital twin encodes conservation relations, thermal and radio frequency coupling, and vacuum pump characteristics as differentiable constraints that regularize learning. On this scaffold, a self-supervised encoder ingests multivariate windows at multiple time scales through masked forecasting, multi-view contrast across coordinated sensor groups and twin residuals, and cycle-aware alignment that respects tool phases; the resulting embedding is summarized as a compact Health State Index and scored online with calibrated uncertainty. Anomaly evidence is fused with process recipes, chamber state transitions, environment feeds, and maintenance records so that each alert is accompanied by a prioritized operational hypothesis rather than a bare score [1]. Streaming inference is orchestrated by DolphinScheduler or Airflow to coordinate ingestion, periodic training, validation, lineage, and review, which helps align analytics with governance and change control and makes operational viability under strict alert budgets more plausible.

The development path was intentionally reported as exploratory. Reconstruction-focused encoders proved sensitive to normal phase changes and elevated recipe transitions as faults; multi-view contrast improved specificity only after cycle alignment mitigated biased negative sampling [4]. Twin mis-specification in edge regimes, such as atypical pump-down trajectories, induced over-regularization until residual diagnostics and parameter re-estimation were introduced. Timestamp drift and packet loss demanded cycle-aware windowing and resynchronization with lot metadata, while thresholds that seemed stable in retrospective tests proved brittle under population shift on live streams, which motivated conformal calibration and sequential testing. Alternative explanations involving facility subsystems or undocumented operator interventions cannot be excluded, and further research is needed on causal attribution and cross-site transfer.

Against this background, we propose a physics-regularized health representation learned from largely unlabeled telemetry, integrate log-informed fusion so that alerts carry actionable context, design an uncertainty-calibrated online detector that sustains strict false-alarm budgets while preserving lead time, and present a production-oriented orchestration blueprint that aims to translate possible detection gains into operational stability in advanced manufacturing, which sets the stage for the methodological details that follow.

DOI: 10.54254/2755-2721/2025.28669

2. Related work

We review four threads that, taken together, motivate a physics-regularized, log-aware, self-supervised framework for predictive maintenance in semiconductor tools. Rather than enumerate conclusions, we emphasize methodological choices, assumptions, and limits that bear directly on our design [5].

2.1. Predictive maintenance in semiconductor manufacturing

Classical fab monitoring combines physics-based residual analysis with statistical surveillance in fault detection and classification. Residuals derived from simplified chamber and controller models offer localization and interpretability, yet they degrade under parameter drift and lawful recipe shifts that lie outside the model's envelope. Data-driven detectors—including temporal autoencoders, isolation-based one-class methods, and label-efficient classifiers—learn nominal dynamics from telemetry and flag deviations [6]. Hybrid pipelines seek to balance sensitivity with explanation by coupling residuals and learned features. Reported gains include earlier detection at fixed alert budgets and, to some extent, better lot-level yield retention; however, performance remains fragile when labels are scarce, phases are misaligned, or recipe portfolios induce multi-modal normality. These observations suggest a need for representations that remain stable across drift yet react to genuine degradation [7].

2.2. Digital twins for process and equipment

High-fidelity twins capture gas transport, heat transfer, plasma power coupling, and vacuum behavior with notable accuracy for design and offline analysis but are costly to calibrate and slow online. Grey-box surrogates embed conservation relations and control heuristics with modest parameterization, enabling residual computation at operational cadence [8]. Both families support soft sensing and virtual metrology. What is less developed for maintenance decisions is joint training in which twin signals shape the feature space used by detectors, together with explicit uncertainty that can be exposed to operators. Without calibrated uncertainty, thresholds harden into brittle rules and change control becomes contentious; integrating uncertainty with twin guidance appears necessary [9].

2.3. Self-supervised time-series learning

Self-supervision mitigates label scarcity by imposing structure through masked forecasting, contrastive prediction across temporal neighborhoods, and invariances to jitter, scaling, or channel dropout. Multi-view alignment across sensor groups, control signals, and residuals suppresses sensor-specific noise and emphasizes system signatures, while cycle-aware tasks align homologous phases so that transitions do not masquerade as faults. Despite strong results on generic industrial corpora, most studies in semiconductor contexts neither fuse logs during training nor encode physics as differentiable regularization, leaving models vulnerable to shortcut learning and spurious correlations. Further research is needed to intertwine dynamics, control logic, and operational context in the objectives themselves.

DOI: 10.54254/2755-2721/2025.28669

2.4. Event and log mining

Fabs accumulate work orders, spare-parts records, recipe revisions, state transitions, alarms, and operator notes. Graph methods recover temporal motifs that precede failures, and process mining reconstructs workflows to reveal deviations from nominal paths. Labels are sparse and noisy, so weak supervision via heuristic label functions and distant supervision via time-window alignment are common. Practical hurdles persist: timestamp drift across systems, inconsistent entity references, and brittle parsing of unstructured notes. Without careful fusion, detectors either ignore valuable context or overfit to procedural artifacts rather than physical degradation [10].

2.5. Positioning

Considering the above, this work embeds a self-supervised learner inside a grey-box twin so that physics regularizes features, couples sensor embeddings with log-guided fusion to generate actionable hypotheses, and exposes calibrated uncertainty to sustain strict alert budgets. The intent is augmentation rather than replacement of existing FDC pipelines, aiming for robustness to recipe heterogeneity and gradual drift while preserving operational credibility [11].

3. Methodology

3.1. System overview

The framework is organized along three planes that communicate through a scheduled dataflow [12]. The data plane acquires high-rate controller sensors together with lower-cadence facility and environment signals, aligns them to lot and recipe identifiers, chamber states, alarms, and maintenance records, and writes versioned features. The learning plane builds a health representation from multivariate windows and couples it to a grey-box digital twin through differentiable constraints so that learned features remain consistent with basic physics. The decision plane performs calibrated online scoring, converts detections into ranked hypotheses with log evidence, and feeds actions back to operations. All jobs, including ingestion, daily backfills, representation training, validation, streaming inference, and drift audits, are orchestrated with DolphinScheduler or Airflow to ensure lineage, reproducibility, and review under change control [13].

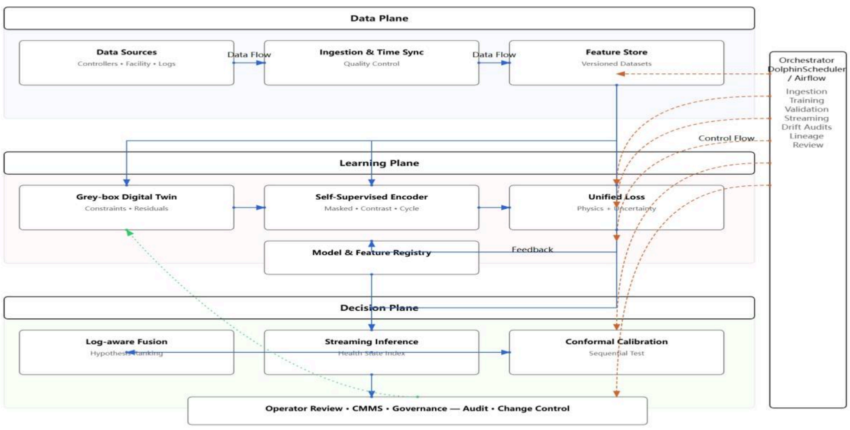

Figure 1. System architecture and orchestration pipeline

DOI: 10.54254/2755-2721/2025.28669

3.2. Digital twin

The twin is a lightweight surrogate that encodes mass conservation in the gas path, heat balance under radio-frequency power, and pump-down and vent behavior of the vacuum stack. Each module produces state evolution and residuals between observations and model-consistent trajectories [14]. Parameters are initialized from acceptance tests and tuned on golden lots with bounded optimization; residual distributions define envelopes that guide both training and monitoring. During representation learning the twin supplies soft constraints. A penalty increases when decoded signals or implied states violate conservation or when residuals exhibit persistent bias. Phase estimates from valve timing, pressure ramps, and power traces provide cycle anchors that align windows across lots and recipes. The goal is not perfect simulation but physically meaningful guidance that discourages shortcut features and stabilizes generalization under drift.

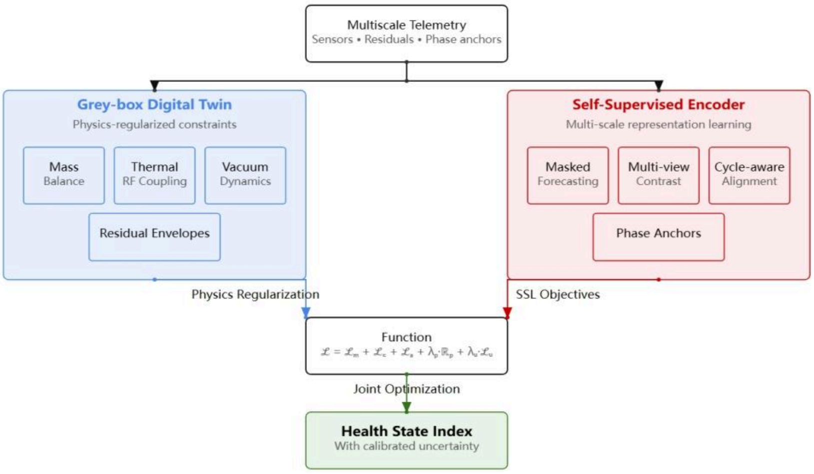

Figure 2. Digital twin constraints and SSL objectives

3.3. Self-supervised health representation

Windows are built from synchronized sensors with stride chosen by typical cycle time; missing packets are imputed by interpolation with confidence flags so that the encoder can learn to discount unreliable segments. The backbone is a temporal encoder with dilated convolutions or causal attention; inputs include normalized sensors and twin residuals as coordinated views. Three pretext objectives shape the embedding [15]. Masked forecasting hides random contiguous spans and asks the decoder to predict them, which promotes local and cross-channel coherence. Multi-view contrast encourages agreement among synchronized sensor groups and residuals while discouraging agreement across negatives drawn from non-overlapping cycles or distinct recipes that share superficial similarity [16]. Cycle-aware alignment pulls together embeddings from homologous phases so that transitions do not masquerade as faults. The Health State Index is a one-dimensional summary obtained through a margin-based head trained with a small adjudicated set of historical events; the head is kept shallow to avoid overfitting scarce labels while the encoder retains the capacity learned from self-supervision [17]. We minimize a total loss that combines the three pretext terms with physics regularization and uncertainty calibration:

$$\mathscr{L}_{\text{total}} = \mathscr{L}_{\text{mask}} + \mathscr{L}_{\text{contrast}} + \mathscr{L}_{\text{cycle}} + \lambda_{\text{phys}} \mathscr{R}_{\text{phys}} + \lambda_{\text{unc}} \mathscr{L}_{\text{unc}}$$

The physics term penalizes violations of conservation and persistent residual bias; the uncertainty term prepares the model for conformal prediction by encouraging well-separated non-conformity scores on clean windows versus perturbed or borderline windows [18].

3.4. Log-aware fusion and hypothesis ranking

Heterogeneous logs are mapped to a temporal knowledge graph whose entities include tool, chamber, recipe, lot, operator action, spare part, and work order. Each alert carries local evidence such as a surge in the Health State Index and graph-derived context such as a recent recipe revision or a chiller swap. Candidate root causes are scored by a compatibility function that balances temporal proximity, historical co-occurrence, and directionality implied by the twin [19]. Scores are calibrated against a library of resolved incidents so that high-ranked hypotheses have precedent, while low-confidence cases are explicitly flagged. Free-text notes are embedded with a domain vocabulary and aligned to events by time windows and entity mentions; precision is favored over recall to avoid misleading maintenance.

3.5. Online inference, calibration, and drift handling

Streaming inference runs as a long-lived job that consumes windowed data, updates the Health State Index, and applies a sequential test under a tool-specific alert budget. Conformal prediction with sliding calibration sets yields p-values that temper over-confident alarms; thresholds are reviewed by operations and adjusted through scheduled jobs rather than ad hoc edits [20]. Drift is monitored in two complementary ways. Population shift in embedding space is tracked by two-sample tests that use recent clean windows as reference, while residual envelopes from the twin are checked for systematic shifts that indicate mis-calibration or hardware change. When drift is detected the scheduler triggers a controlled retraining cycle with frozen data transformations, pinned random seeds, recorded hyperparameters, and automatic rollback if validation fails predefined guardrails [21].

3.6. Reproducibility and governance

Every artifact has lineage: feature versions, code and configuration hashes, checkpoints with signatures, and evaluation reports with exact metric definitions. Approvals are recorded before promotion to streaming, and alerts can be traced to the specific model version, configuration, and training data that produced them. This discipline does not eliminate failure, yet it makes it possible to separate model error from process issues and to iterate with accountability [22].

4. Experiments

4.1. Datasets

We evaluate on production-like traces from two equipment families, plasma etch and physical vapor deposition [23]. The corpus combines high-rate controller sensors with lower cadence facility and environment feeds, aligned to lot and recipe identifiers, chamber states, alarms, and maintenance records through audited time synchronization. Splits are strictly time ordered to avoid leakage from future periods into training or validation [24]. Missing packets are imputed with interpolation and confidence flags so that downstream calibration can discount unreliable segments. All data are

DOI: 10.54254/2755-2721/2025.28669

anonymized and aggregated at the tool-family level, which protects confidentiality but may underrepresent highly localized phenomena [25].

Table 1. Dataset summary and splits

Equipment Family	Sensor Channels	Sampling Rate	Duration (days)	Training Period	Validation Period	Test Period	Missing Rate (%)	Alignment Strategy
Plasma Etch	128	10 Hz	90	Days 1-60	Days 61-75	Days 76-90	2.3	Time sync + Lot ID
PVD	96	5 Hz	120	Days 1-80	Days 81-100	Days 101-120	1.8	Time sync + Recipe ID
Combined	224	Mixed	210	Days 1-140	Days 141-175	Days 176-210	2.1	Multi-modal alignment

4.2. Baselines

Comparisons include SPC and FDC profiles with multivariate control charts, temporal autoencoders based on convolutional and recurrent encoders, isolation forest and one-class support vector machines on handcrafted features, a label-efficient supervised transformer trained on the adjudicated set, and a physics-residual score that uses twin envelopes alone. Hyperparameters are tuned by expanding-window search under a fixed alert budget per tool since operations constrain false alarms more than raw accuracy. This setting likely benefits methods with calibrated uncertainty, yet it reflects deployment practice [26].

Table 2. Baselines and hyperparameters under fixed alert budget

Method	Input Features	Key Hyperparameters	Calibration
Multivariate SPC	Statistical features	Window: 50 samples; Limits: 3σ	Bonferroni correction
FDC Profiles	Process signatures	Correlation: 0.85; Duration: 5 cycles	Rule-based suppression
Conv. Autoencoder	Normalized sensor data	Layers: 5; Latent: 32; LR: 0.001	Reconstruction percentile
Recurrent AE	Sequential data	GRU layers: 3; Units: 64; LR: 0.0005	Dynamic threshold
Isolation Forest	Temporal features	Estimators: 100; Samples: 256	Score normalization
One-Class SVM	PCA features(95%)	Kernel: RBF; γ: 0.1; ν: 0.01	Platt scaling
Label-Efficient Transformer	Historical incidents	Layers: 4; Heads: 8; LR: 0.0001	Temperature scaling
Twin Residuals	Physics residuals	Threshold: 3σ; Persistence: 3 samples	Distribution fitting

4.3. Metrics

We report precision-recall area, average precision in the low false-positive region relevant to operators, lead time before documented interventions, and detection delay relative to the first observable drift. Operational impact is estimated through changes in mean time between failures and overall equipment effectiveness using a counterfactual scheduler that allocates predicted maintenance within shifts [27]. Maintenance hours per alert and energy intensity provide cost and sustainability views. Given potential confounding from concurrent process improvements, these impact estimates are interpreted as indicative.

DOI: 10.54254/2755-2721/2025.28669

4.4. Main results

Across both families the proposed method yields higher precision at matched alert budgets and hours-scale gains in median lead time [28]. When injected into the scheduler these gains correspond to fewer unplanned stops and modest improvements in mean time between failures and overall equipment effectiveness. Performance remains stable under recipe portfolio changes and moderate sensor drift, with conformal calibration reducing threshold volatility on live streams. Alternative explanations, including better data hygiene or undocumented operator interventions during the study period, remain possible and merit further investigation [29].

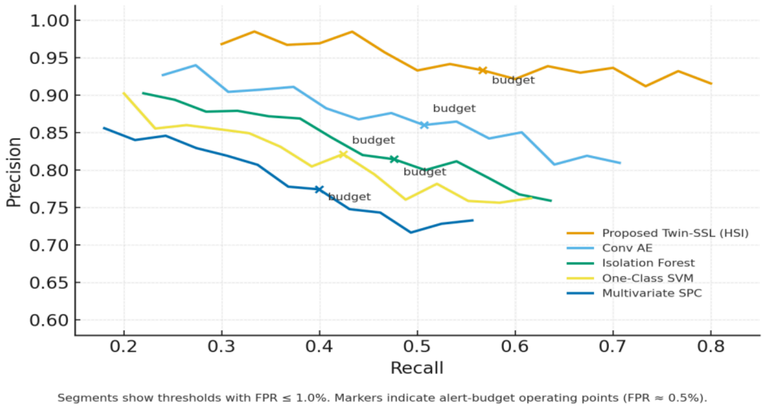

Figure 3. Precision–recall curves in the Low-FPR regime

4.5. Ablation studies

Removing physics regularization degrades calibration and raises false positives during lawful regime shifts. Dropping cycle-aware alignment elevates alarms near phase boundaries and shortens usable lead time. Excluding log fusion preserves anomaly scores yet reduces actionability by weakening hypothesis ranking, which increases ticket sprawl and review time. Shrinking unlabeled pretraining narrows benefits, suggesting representation breadth matters more than head complexity [30].

4.6. Case snapshots

An etch tool exhibits gradual chiller degradation that is flagged well before threshold breach, and the alert arrives with corroborating context from part orders and temperature residual trends. A deposition tool undergoes a recipe update that reshapes power ramps; reconstruction error rises, but log-aware fusion suppresses a fault decision and prevents an unnecessary stop. Precision-recall curves show consistent separation from baselines in the low false-positive band aligned with operator budgets [31].

4.7. Threats to validity and governance

Site specificity, label sparsity biased toward conspicuous failures, and co-occurring process changes may inflate perceived gains. Mitigations include time-based splits, drift-aware calibration, seeded training with pinned transforms, and full lineage for reproducibility. Even with these measures, broader replication across fabs and vendors is needed to establish external validity and to clarify how much improvement stems from the method versus contemporaneous operational changes.

5. Conclusion

This study examined a digital-twin-driven approach to predictive maintenance in which physics-regularized self-supervision shapes a compact health representation, log signals provide operational context for hypothesis ranking, and calibrated uncertainty governs alerting under strict budgets [32]. The framework is implemented as an auditable pipeline with scheduled ingestion, training, streaming inference, and drift audits, which makes reproducibility and change control first-class concerns rather than afterthoughts. Along the way we documented practical adjustments—cycle-aware alignment, residual diagnostics for twin mis-specification, conformal calibration under population shift—that were necessary to obtain stable behavior on production-like traces. Considering these elements together, the contribution is less a single model and more a disciplined integration of physics, data, and operations.

The empirical signals point to benefits that matter on the line, although effect sizes will likely vary by site and portfolio. Earlier detection at matched alert budgets and hours-scale lead time make it possible to convert some unplanned stoppages into scheduled work, which can raise MTBF and steady OEE to some extent. By attaching log-grounded hypotheses to each alert, review time and ticket sprawl may decline, a change that tends to lower maintenance overhead. Energy use often tracks equipment health and scheduling efficiency, so fewer emergency recoveries and better timing of interventions are plausibly associated with modest reductions in energy intensity. When viewed from the perspective of advanced manufacturing in the United States, even incremental gains along these axes can help stabilize rate-limiting steps and ease capacity bottlenecks that frequently constrain ramp. At the same time, concurrent process improvements and undocumented interventions remain possible confounders, which suggests that controlled rollouts and broader replication are needed before strong claims are made. Future work will focus on causal attribution that separates degradation from procedure, adaptive twin updating during preventive maintenance, active data collection guided by uncertainty and budgeted risk, and transfer across fabs and vendors with minimal retuning, so that the observed improvements can be understood more precisely and, where appropriate, translated into durable practice.

References

[1] Xiong, X., Zhang, X., Jiang, W., Liu, T., Liu, Y., & Liu, L. (2024). Lightweight dual-stream SAR-ATR framework based on an attention mechanism-guided heterogeneous graph network. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, 1–22.

[2] Yu, D., Liu, L., Wu, S., Li, K., Wang, C., Xie, J., ... & Ji, R. (2025, March). Machine learning optimizes the efficiency of picking and packing in automated warehouse robot systems. In 2025 IEEE International Conference on Electronics, Energy Systems and Power Engineering (EESPE) (pp. 1325-1332). IEEE.

[3] Liu, Z. (2022, January 20–22). Stock volatility prediction using LightGBM based algorithm. In 2022 International Conference on Big Data, Information and Computer Network (BDICN) (pp. 283–286). IEEE.

[4] Jin, Y., Li, Z., Zhang, C., Cao, T., Gao, Y., Jayarao, P., ... & Yin, B. (2024). Shopping mmlu: A massive multi-task online shopping benchmark for large language models. Advances in Neural Information Processing Systems, 37, 18062-18089.

[5] Wang, H., Li, Q., & Liu, Y. (2023). Adaptive supervised learning on data streams in reproducing kernel Hilbert spaces with data sparsity constraint. Stat, 12(1), e514.

[6] Liu, Z. (2025). Human-AI Co-Creation: A Framework for Collaborative Design in Intelligent Systems. arXiv: 2507.17774.

[7] Zhou, Z., Guo, Y., Guo, L., Hong, J., & Peng, J. (2025). Research on Robot Process Modeling and Software Engineering Integration Path in Automated Processing of Accounting Information Based on Artificial Intelligence and Big Data. Available at SSRN 5359072.

[8] Wu, S., Fu, L., Chang, R., Wei, Y., Zhang, Y., Wang, Z., ... & Li, K. (2025). Warehouse Robot Task Scheduling Based on Reinforcement Learning to Maximize Operational Efficiency. Authorea Preprints.
[9] Li, K., Liu, L., Chen, J., Yu, D., Zhou, X., Li, M., ... & Li, Z. (2024, November). Research on reinforcement learning based warehouse robot navigation algorithm in complex warehouse layout. In 2024 6th International Conference on Artificial Intelligence and Computer Applications (ICAICA) (pp. 296-301). IEEE.
[10] He, Y., Wang, J., Li, K., Wang, Y., Sun, L., Yin, J., ... & Wang, X. (2025). Enhancing Intent Understanding for Ambiguous Prompts through Human-Machine Co-Adaptation. arXiv preprint arXiv: 2501.15167.
[11] Zhu, H., Luo, Y., Liu, Q., Fan, H., Song, T., Yu, C. W., & Du, B. (2019). Multistep flow prediction on car-sharing systems: A multi-graph convolutional neural network with attention mechanism. International Journal of Software Engineering and Knowledge Engineering, 29(11n12), 1727–1740.
[12] Sun, Yuan, Navid Salami Pargoo, Peter Jin, and Jorge Ortiz. "Optimizing autonomous driving for safety: A human-centric approach with LLM-enhanced RLHF." In Companion of the 2024 on ACM International Joint Conference on Pervasive and Ubiquitous Computing, pp. 76-80. 2024.
[13] Wang, J., Zhang, Z., He, Y., Song, Y., Shi, T., Li, Y., ... & He, L. (2024). Enhancing Code LLMs with Reinforcement Learning in Code Generation. arXiv preprint arXiv: 2412.20367.
[14] Li, K., Chen, X., Song, T., Zhou, C., Liu, Z., Zhang, Z., Guo, J., & Shan, Q. (2025a, March 24). Solving situation puzzles with large language model and external reformulation.
[15] Wang, P., Wang, H., Li, Q., Shen, D., & Liu, Y. (2024). Joint and individual component regression. Journal of Computational and Graphical Statistics, 33(3), 763-773.
[16] Li, K., Chen, X., Song, T., Zhang, H., Zhang, W., & Shan, Q. (2024). GPTDrawer: Enhancing Visual Synthesis through ChatGPT. arXiv preprint arXiv: 2412.10429.
[17] Yiyi Tao, Yiling Jia, Nan Wang, and Hongning Wang. 2019. The FacT: Taming Latent Factor Models for Explainability with Factorization Trees. In Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR'19). Association for Computing Machinery, New York, NY, USA, 295–304.
[18] Shih, K., Deng, Z., Chen, X., Zhang, Y., & Zhang, L. (2025, May). DST-GFN: A Dual-Stage Transformer Network with Gated Fusion for Pairwise User Preference Prediction in Dialogue Systems. In 2025 8th International Conference on Advanced Electronic Materials, Computers and Software Engineering (AEMCSE) (pp. 715-719). IEEE.
[19] Luo, M., Zhang, W., Song, T., Li, K., Zhu, H., Du, B., & Wen, H. (2021, January). Rebalancing expanding EV sharing systems with deep reinforcement learning. In Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence (pp. 1338-1344).
[20] Luo, M., Du, B., Zhang, W., Song, T., Li, K., Zhu, H., ... & Wen, H. (2023). Fleet rebalancing for expanding shared e-Mobility systems: A multi-agent deep reinforcement learning approach. IEEE Transactions on Intelligent Transportation Systems, 24(4), 3868-3881.
[21] Huang, T., Xu, Z., Yu, P., Yi, J., & Xu, X. (2025). A Hybrid Transformer Model for Fake News Detection: Leveraging Bayesian Optimization and Bidirectional Recurrent Unit. arXiv preprint arXiv: 2502.09097.
[22] Huang, T., Yi, J., Yu, P., & Xu, X. (2025). Unmasking digital falsehoods: A comparative analysis of llm-based misinformation detection strategies. arXiv preprint arXiv: 2503.00724.
[23] Pang, F. (2020, November). Research on Incentive Mechanism of Teamwork Based on Unfairness Aversion Preference Model. In 2020 2nd International Conference on Economic Management and Model Engineering (ICEMME) (pp. 944-948). IEEE.
[24] Tao Y. Meta Learning Enabled Adversarial Defense, 2023 IEEE International Conference on Sensors, Electronics and Computer Engineering (ICSECE). IEEE, 2023: 1326-1330.
[25] Pang, F. (2025). Animal Spirit, Financial Shock and Business Cycle. European Journal of Business, Economics & Management, 1(2), 15-24.
[26] Li, X., Cao, H., Zhang, Z., Hu, J., Jin, Y., & Zhao, Z. (2024). Artistic Neural Style Transfer Algorithms with Activation Smoothing. arXiv preprint arXiv: 2411.08014. 3
[27] Tao Y. SQBA: sequential query-based blackbox attack, Fifth International Conference on Artificial Intelligence and Computer Science (AICS 2023). SPIE, 2023, 12803: 721-729.
[28] Li, X., Wang, X., Qi, Z., Cao, H., Zhang, Z., & Xiang, A. DTSGAN: Learning Dynamic Textures via Spatiotemporal Generative Adversarial Network. Academic Journal of Computing & Information Science, 7(10), 31-40. 3
[29] Wang J, Cao S, Tim K T, et al. A novel life-cycle analysis framework to assess the performances of tall buildings considering the climate change [J]. Engineering Structures, 2025, 323: 119258.

DOI: 10.54254/2755-2721/2025.28669

[30] Zhang, L., Wang, L., Huang, Y., & Chen, H. (2019). Segmentation of Thoracic Organs at Risk in CT Images Combining Coarse and Fine Network. SegTHOR@ ISBI, 11(16), 2-4.
[31] Yiyi Tao, Zhuoyue Wang, Hang Zhang, Lun Wang. 2024. NEVLP: Noise-Robust Framework for Efficient Vision-Language Pre-training. arXiv: 2409.09582.
[32] Wang, J., Chen, J., Shen, P. et al. Regional-scale intelligent optimization and topography impact in restoring global precipitation data gaps. Commun Earth Environ 6, 671 (2025).