ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov1299

A Novel SLA-OPT Algorithm for Dynamic Prioritization of Corrective Actions Based on Incident Severity, Business Impact and Resolution Risk Compared with Rule-Based Scheduling and Genetic Algorithms

Emmanuel Iretomiwa Joda[1]; Otugene Victor Bamigwojo[2]

[1]Independent Research, Department of Site Reliability Engineering, PayPal Inc., Austin, Texas, USA
[2]Department of Mathematics, Federal University Lokoja, Nigeria

Publication Date: 2026/11/27

Abstract: Service-level agreement (SLA) breaches in large-scale payment services can occur when corrective actions compete for limited engineering capacity, and conventional severity-based queues fail to account for business exposure, uncertain resolution duration, and the time remaining before contractual deadlines. This study develops SLA-OPT, a rolling-horizon incident prioritisation algorithm combining calibrated breach propensity, incident severity, business-weighted loss, resolution risk, and aging protection in a constrained multi-server dispatch policy. A reproducible synthetic experiment generated 36 independent replications of 42 incidents and compared SLA-OPT with a static severity-and-deadline rule and a genetic algorithm (GA) optimizing predicted weighted lateness. Under the declared scenario settings, mean raw SLA breach proportions were 0.745, 0.595, and 0.713, respectively; corresponding impact-weighted breach fractions were 0.632, 0.450, and 0.446. SLA-OPT reduced impact-weighted tardiness relative to both comparators, while the GA achieved fewer raw deadline breaches. These contrasting outcomes demonstrate the importance of prespecified and separately reported governance objectives. No bank, payment-provider, or PayPal incident records were used. The study supplies an auditable scheduling formulation and testable hypotheses for subsequent shadow deployment and controlled field evaluation.

Keywords: *Service-Level Agreements; Corrective-Action Governance; Incident Prioritisation; Genetic Algorithms; Scheduling; Business Impact; Risk-Aware Optimisation; Payment Infrastructure.*

How to Cite: Emmanuel Iretomiwa Joda; Otugene Victor Bamigwojo (2026) A Novel SLA-OPT Algorithm for Dynamic Prioritization of Corrective Actions Based on Incident Severity, Business Impact and Resolution Risk Compared with Rule-Based Scheduling and Genetic Algorithms. *International Journal of Innovative Science and Research Technology*, 11(11), 1945-1958. https://doi.org/10.38124/ijisrt/26nov1299

I. INTRODUCTION

Large-scale payment infrastructure operates under stringent reliability and availability obligations. A transaction-processing interruption can propagate to authorization, fraud screening, ledger posting, reconciliation, merchant settlement, and external dependencies. Incident teams therefore confront competing claims on staff attention even when all cases are classified as urgent. A formal SLA provides a measurable clock, but the operational consequence of breaching the clock depends on affected services, customer exposure, contractual commitments, the possibility of cascading failure, and the urgency of corrective action (Beyer et al., 2016; Qazi et al., 2024). Corrective-action prioritisation is consequently not equivalent to arranging tickets by severity labels alone.

The classical service-management workflow records detection, assignment, diagnosis, remediation, verification, and closure. In practice, these stages interact with limited on-call engineer capacity, access approvals, dependency constraints, changing incident information, and heterogeneous restoration durations. A queue that appears efficient under an average processing-time criterion can produce unacceptable tail delays or direct scarce engineering time away from high-impact cases. Research on cloud incidents and intelligent incident management highlights the complexity of deriving actionable decisions from telemetry and ticket narratives (Yu et al., 2024; Li et al., 2022). These

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov1299

observations motivate a decision layer that is accountable to business impact as well as elapsed time.

Scheduling research supplies mathematical tools for evaluating such trade-offs. Rule-based approaches are interpretable and inexpensive but have limited ability to adapt to changing conditions. Genetic algorithms permit large combinatorial searches and can optimise a prespecified objective, yet their practical value depends on reliable duration estimates, feasible assignments, computation budgets, and the risk that an optimised plan becomes stale when an urgent incident arrives (Goldberg, 1989; Pinedo, 2016). SLA-OPT is positioned between fixed-priority dispatch and expensive global search: it refreshes incident urgency at each decision epoch and retains explicit guardrails for critical cases.

This paper asks whether incorporating remaining SLA time, estimated resolution risk, and business impact into corrective-action scheduling produces different operational trade-offs from severity-based and GA dispatching. Its contributions are a formally specified rolling-horizon score, explicit baseline definitions, reproducible common-random-number simulation, multiple non-interchangeable performance indicators, and a prospective validation protocol. The algorithm is presented as a research proposal, not as an established product or a deployed PayPal tool.

➢ *Research Questions and Hypotheses*

RQ1 evaluates how much the ranking of corrective actions changes when severity, monetary exposure, deadline slack and repair uncertainty are considered jointly. RQ2 compares raw breach rate, business-weighted breach rate, weighted tardiness, completion-time distributions and subgroup outcomes across three declared policies. RQ3 assesses whether apparently favourable averages conceal high-risk or high-severity regressions. H1 proposes that risk-aware dispatch improves business-weighted lateness relative to a severity-only rule. H2 does not assume that SLA-OPT will dominate a GA on every objective; instead, the direction and magnitude of each trade-off are empirical questions. H3 proposes that independent replay and prespecified safety gates are necessary before moving from computational feasibility to production deployment.

II. LITERATURE REVIEW AND CONCEPTUAL FOUNDATIONS

➢ *SLA Governance and Incident Management*

Service-level objectives should specify an unambiguous start and stop event, the scope of excluded interruptions, the unit of analysis, and the consequence of noncompliance. The ISO/IEC 20000-1 service-management model and site-reliability engineering literature place these measurements within an accountable operational system rather than treating availability as a stand-alone dashboard metric (International Organization for Standardization, 2018; Beyer et al., 2016; Frank, 2026c). For payment infrastructure, reliable incident clocks also require an audit trail of ownership transfers, mitigations, handbacks, verification, and reopened tickets.

Yu et al. (2024) synthesised research on intelligent alert and incident management, while Li et al. (2022) examined causal diagnosis of complex online services. These sources support using monitoring and diagnosis as inputs to action selection, but they do not establish that a single composite ranking formula is superior for payment-related corrective actions. That hypothesis remains open and should be tested using observed, appropriately governed event data.

➢ *Predictive Monitoring and Uncertain Resolution Durations*

Predictive process monitoring estimates case outcomes and remaining activity time from event logs. Tax et al. (2017), Teinemaa et al. (2019), and Verenich et al. (2019) show why such models depend on the chosen event representation, case population, and evaluation protocol. In incident governance, a duration estimate should be accompanied by a calibrated uncertainty distribution: a nominal 30-minute remediation may involve delayed approvals or repeated execution in its upper tail. Miscalibration becomes a scheduling error because the optimiser ranks cases using predicted slack.

Reliability theory and queueing models also clarify why a pool of identical responders does not guarantee identical queue delays. Changing workload, correlated arrivals and unequal repair times can induce congestion. The mathematical model should therefore use both clock-based pressure and risk, with sensitivity analysis to duration-estimation errors (Kleinrock, 1975; Law, 2015; Rausand & Høyland, 2004; Frank, 2026a).

➢ *Rule-Based Scheduling and Evolutionary Optimisation*

Dispatching heuristics may rank cases by earliest deadline, severity, shortest processing time, or weighted due-date loss. These decisions are inexpensive and interpretable but can be inconsistent when severity and deadlines disagree. Genetic algorithms explore permutations through selection, mutation and recombination, and remain attractive when objectives are nonlinear or constraints make analytical solutions difficult (Goldberg, 1989; Deb et al., 2002; Bamidele, et al., 2025). The comparison must specify the exact genetic search budget and its knowledge assumptions: optimizing on actual future repair durations would be an unfair clairvoyant baseline.

The proposed GA instead uses noisy predicted durations and a prespecified risk-weighted lateness objective. Its realised performance is measured on subsequently revealed synthetic actual durations, as are the other policies. This difference between predicted optimisation cost and realised performance is important for interpreting any apparent advantage.

➢ *Corrective-Action Evidence, Control Gates and Operational Fairness*

A corrective action is complete only when its effect is independently verified and recorded. Workflow mining can

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov1299

reconstruct evidence of ownership and closure, but a high-priority model should not automatically authorize risky operations. Human sign-off remains necessary for payments, customer-visible changes, database modifications, and any action that could introduce a more consequential outage (van der Aalst, 2016; NIST, 2024). The proposed framework therefore separates queue ranking from authority to execute.

Priority systems can also unintentionally starve low-severity cases as fresh emergencies arrive. Age adjustments and maximum-wait escalation criteria provide one guard against such behavior (Akinleye, & Ijiga, 2025; Ogundola, 2024). The final operational design must use stakeholder-approved business-impact definitions and report adverse distributional effects, not merely average queue efficiency. Uncertainty-sensitive optimisation provides a formal language for these risk constraints (Rockafellar & Uryasev, 2000; Shapiro et al., 2021).

➢ *Research Gap and Conceptual Architecture*

The methodological gap addressed here is the link between predictive incident intelligence and admissible corrective-action scheduling. Studies of monitoring and root cause analysis describe what may be wrong; scheduling studies explain how constrained jobs can be sequenced; governance standards describe when people must authorize changes. SLA-OPT combines these layers without assuming that the ranking score itself establishes causal effectiveness or replaces operational accountability. The architecture in Figure 1 emphasizes five functions: event intake, risk estimation, rolling-horizon ranking, controlled execution and post-action evidence. It also highlights guardrails that prevent a high score from bypassing authorization or traceability requirements.

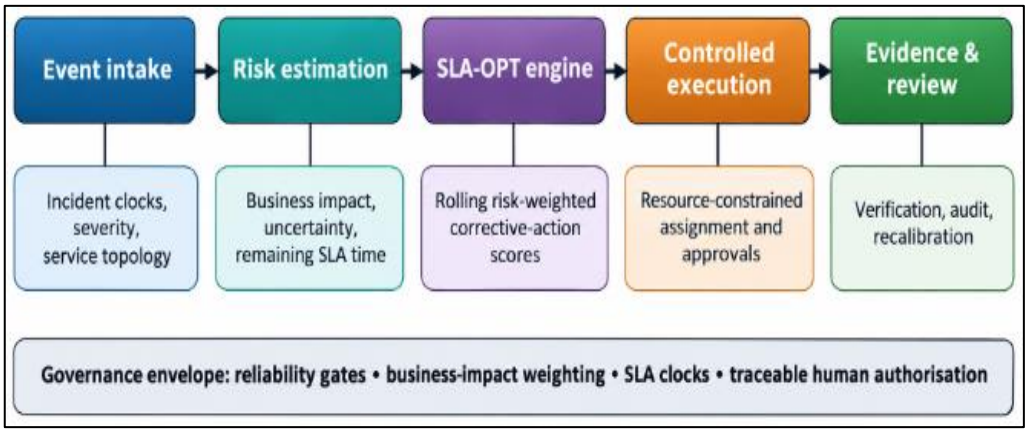


Fig 1 Conceptual Architecture of SLA-OPT Linking Incident Telemetry, Risk Estimation, Scheduling, Controlled Action and Auditable Verification.

III. MATERIALS AND METHODS

The study is a computational methodological experiment designed for reproducibility. It is not an empirical evaluation of a real financial platform. The simulated unit is an incident requiring one nonpreemptive corrective-action service from a pool of four identical responders. All three policies face the same incident attributes and realised service durations within each replication. Policy comparisons are paired by replication to reduce nuisance variation, although residual Monte Carlo uncertainty remains. No payment transactions, user identities, proprietary telemetry or incident logs are included.

➢ *Incident Representation and Baseline Variables*

Let each case possess a severity ordinal, an impact indicator, an uncertain work duration, a resolution-risk score and a contractual deadline. The impact score is not represented as an observed financial loss; it is a synthetic unitless weighting variable. Four responders are assumed continuously available and nonpreemptive. The experimental workload is a batch released at a common origin, a deliberate simplification that makes dispatch differences interpretable and does not reproduce continuous, autocorrelated production arrival processes. Each simulation replication contains 42 cases; 36 independently seeded workloads comprise the test cohort.

Table 1 Incident Attributes, Assumptions and Interpretive Limitations in the Computational Experiment.

Variable	Definition	Synthetic specification	Interpretation
Severity	Ordinal incident class	Uniform integer 1–5	Proxy for criticality
Impact	Business exposure weight	Lognormal, severity-scaled	Not actual currency
Resolution risk	Risk of uncertain repair	Beta distribution	Unitless score
Duration	Actual response work minutes	Lognormal with severity shift	Revealed only in evaluation
Prediction	Available duration estimate	Actual duration × lognormal noise	Used for scheduling
SLA deadline	Time allowance from incident origin	Clipped gamma mixture	Minutes
Capacity	Concurrent engineering responders	Four identical servers	Fixed scenario assumption

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov1299

$$s_i(t) = d_i - t - \hat{\mu}_i \quad (1)$$

Where $s_i(t)$ is estimated remaining slack for incident i at decision time t; d_i is its SLA deadline; $\hat{\mu}_i$ is predicted processing duration.

Equation (1) makes explicit that short deadlines and long predicted repairs can both create operational urgency. A negative predicted slack denotes anticipated deadline pressure; it is not proof of actual breach because realised duration can differ.

$$p_i(t) = \{1 + exp[(d_i - t - \hat{\mu}_i)/(\kappa\hat{\mu}_i + \varepsilon)]\}^{-1} \quad (2)$$

Where $p_i(t)$ is heuristic breach propensity for incident i; κ is a scale constant (0.26 in the experiment); ε is a stabilizing offset of five minutes; exp is the exponential function.

Equation (2) transforms slack into a bounded, monotone urgency indicator. It is deliberately called a heuristic propensity rather than a calibrated probability: the simulation does not fit a predictive model on historic observed incidents.

➢ *Composite Business Impact and Proposed SLA-OPT Priority*

Incident severity alone is an insufficient proxy for value at risk. The proposed algorithm constructs a multiplicative score with severity-weighted business impact, deadline propensity, repair risk, a modest age adjustment and duration normalization. Multiplicative interactions are appropriate for the design purpose because near-zero impact should not be outweighed by an arbitrary large severity label, while high impact and elevated risk should jointly increase attention. The exponent on estimated duration discourages assigning all resources exclusively to long cases when shorter high-impact actions can be completed before deadline.

$$w_i = b_i(0.55 + 0.45v_i/5) \quad (3)$$

Where w_i is the business-severity weight for incident i; b_i is its positive business impact; v_i is severity on the 1–5 scale.

Equation (3) is the declared interpretation of business exposure in this synthetic study, not a regulatory weighting scale. The same weight family is used in the policy objective and reported weighted outcomes.

$$g_i(t) = w_i(0.50 + 1.40r_i)[0.20 + 0.80p_i(t)](1 + 0.10t/(d_i + 1))/(\hat{\mu}_i)^{0.45} \quad (4)$$

Where $g_i(t)$ is the SLA-OPT priority score; r_i is resolution risk in [0,1]; $p_i(t)$ is Equation (2); t is decision time; d_i is the deadline; $\hat{\mu}_i$ is predicted duration; w_i is Equation (3).

Equation (4) is the novel score tested in this computational study. It is evaluated repeatedly after each simulated dispatch, and the largest admissible score receives the next available responder. Coefficients are prescribed experimental parameters and should be tuned and tested independently before any field use.

$$i*(t) = arg\,max_{i \in E(t)}\, g_i(t), \quad j*(t) = arg\,min_{j \in \{1,\ldots,m\}}\, F_j(t) \quad (5)$$

Where E(t) is the eligible unresolved incident set; i(t) is the selected incident; m is the responder count; $F_j(t)$ is projected finish/availability time for responder j; j*(t) is the earliest available responder.*

Equation (5) defines the policy selection and resource-allocation step. In the implementation, the projected server load is updated using predicted durations. Realised completion is evaluated separately using actual synthetic durations, avoiding access to hidden future work times during scheduling.

➢ *Baselines: Fixed Rule and Genetic Algorithm*

The rule-based comparator first sorts incidents by descending severity and then by ascending SLA deadline. It is a static batch schedule and does not re-rank when resources become free. Its advantage is predictable behavior and negligible computational overhead. The GA comparator searches permutations with a population of 24 schedules over 32 generations; five elite individuals are retained at each generation and mutated offspring provide exploration. The initial population includes an earliest-deadline sequence. The objective uses estimated durations, deadlines, impact and risk, ensuring that the GA is not privileged with realised service times.

$$\hat{L}(\pi) = \Sigma_i\, b_i[(1 - r_i)[\hat{C}_i(\pi) - d_i]_+ + r_i\, min\{60, [\hat{C}_i(\pi) - d_i]_+ + 10\}] \quad (6)$$

Where $\hat{L}(\pi)$ is predicted genetic-search loss for permutation π; b_i is business impact; r_i is resolution risk; $\hat{C}_i(\pi)$ is predicted completion using estimated durations; d_i is the deadline; $[x]_+ = max(x,0)$.

Equation (6) is the GA objective minimized in the supplied code. The capped risk term represents the assumed contribution of uncertain repair to urgency. It is not derived from a bank's contractual penalty function.

$$\pi* \in arg\,min_\{\pi \in \Pi\}\, \hat{L}(\pi), \quad \Pi = \{permutations\ of\ \{1,\ldots,n\}\} \quad (7)$$

Where π is the GA-selected ordering; Π is the finite set of n-incident permutations; n is the number of incidents; $\hat{L}$ is Equation (6).*

Equation (7) clarifies that the GA approximates a permutation optimisation problem, not a certified global optimum. Its compute budget and population size influence

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov1299

results, and the selected 32-generation configuration is only one benchmark setting.

➢ *Completion-Time Dynamics and SLA Constraints*

For each policy, actual durations determine the realised server completion timeline. Incident tasks are allocated to the currently least-loaded server in schedule order. The model assumes one incident per responder at a time, no preemption, and no shared specialist incompatibilities; all of these assumptions would need revision for production incident management. A policy can reduce weighted tardiness while increasing raw breaches by protecting a few high-impact cases at the expense of many lower-impact ones. Consequently, both quantities are mandatory primary reporting outcomes.

$$C_i(\pi) = min\square F\square + T_i,\ F\square \leftarrow C_i(\pi)\ for\ assigned\ responder\ j \quad (8)$$

Where $C_i(\pi)$ is realised completion time of incident i under schedule π; F_j is the earliest current availability of responder j; T_i is realised incident work duration.

Equation (8) maps candidate schedules into simulated outcomes using realised durations only at evaluation time. A genuine event-driven, arrival-based extension would include incident release times and human authorization delays.

$$B_i(\pi) = \mathbb{1}\{C_i(\pi) > d_i\} \quad (9)$$

Where $B_i(\pi)$ is the SLA-breach indicator; $\mathbb{1}$ is the indicator function; $C_i(\pi)$ is realised completion time; d_i is SLA deadline.

Equation (9) uses a strict breach test; an incident closing exactly at the SLA deadline is counted as compliant. This rule should be matched to the contractual clock definition in a field implementation.

$$R_B(\pi) = n^{-1}\Sigma_i B_i(\pi),\ R_W(\pi) = \Sigma_i w_i B_i(\pi)/\Sigma_i w_i \quad (10)$$

Where $R_B(\pi)$ is unweighted breach rate; $R_W(\pi)$ is business-weighted breach fraction; n is the number of incidents; B_i is Equation (9); w_i is Equation (3).

Equation (10) defines two distinct headline outcomes. Neither can replace the other, because a weighted improvement can arise even when the number of noncompliant tickets increases.

➢ *Tail-Risk and Distributional Performance Metrics*

Lateness is measured continuously in minutes rather than reducing the outcome to breach/no breach. The objective does not assume that a 1-minute and a 90-minute exceedance have equal consequences. A third endpoint therefore measures business-weighted tardiness, while the 90th percentile of completion time measures long-delay experience. Interpretation is deliberately multivariate because a scheduler may shift delay from high-impact to low-impact cases without reducing total operational workload.

$$T_W(\pi) = \Sigma_i w_i[C_i(\pi) - d_i]_+ / \Sigma_i w_i \quad (11)$$

Where $T_W(\pi)$ is impact-weighted tardiness; w_i is the case weight; $C_i(\pi)$ is completion time; d_i is the SLA deadline; $[\cdot]_+$ denotes the positive part.

Equation (11) is the primary continuous risk-weighted delay metric. It is evaluated on the exact same realised incident durations across policies in each replication.

$$CVaR\alpha(L) = min\eta\{\eta + (1-\alpha)^{-1}E[(L-\eta)_+]\} \quad (12)$$

Where $CVaR\alpha$ is the conditional value at risk for loss L at probability α; η is an auxiliary threshold; E is expectation; $[\cdot]_+$ is positive part.

Equation (12) specifies an extension for governance-sensitive scheduling under tail loss, following Rockafellar and Uryasev (2000). The initial numerical experiment reports realised upper-tail outcomes but does not actually optimise CVaR; that distinction is necessary for reproducibility.

➢ *Computational Design, Replicability and Paired Inference*

The pseudorandom seed is fixed and publicly declared in the experiment script. Severity is uniform on the five-class scale; impact has a long-tailed distribution; risk is beta-distributed; repair duration is lognormal and correlated with severity; and deadlines follow a bounded gamma specification. Estimated durations are corrupted by multiplicative lognormal noise. For each replication, the three scheduling methods receive identical cases, timestamps, estimates and realised durations. This common-random-number design allows paired contrasts of policy outputs and prevents differences in workload difficulty from masquerading as algorithmic gains.

$$\Delta_\{A,B\} = R^{-1}\Sigma_r(Y_\{A,r\} - Y_\{B,r\}) \quad (13)$$

Where $\Delta_\{A,B\}$ is the paired mean difference between policies A and B; $Y_\{A,r\}$ and $Y_\{B,r\}$ are outcome values for replication r; R is the number of independent simulation replications.

Equation (13) defines the comparison estimand for breach rate, tardiness or completion time. It does not by itself justify confidence claims about real facilities: the replications sample one synthetic data-generating mechanism.

$$SE(\Delta) = \sqrt{\{[R(R-1)]^{-1}\Sigma_r(\delta_r - \delta)^2\}} \quad (14)$$

Where $SE(\Delta)$ is the Monte Carlo standard error of the paired contrast; $\delta_r = Y_\{A,r\} - Y_\{B,r\}$; $\bar{\delta}$ is their average; R is the replication count.

Equation (14) estimates the Monte Carlo uncertainty of a synthetic-policy contrast. Serial or cross-service dependence in real incident histories would require cluster-aware or block-based inference.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov1299

$$G_i = \mathbb{1}\{v_i \geq 4\} \vee \mathbb{1}\{b_i \geq q_{0.9}(b)\} \quad (15)$$

Where G_i is a candidate critical-case governance gate; v_i is incident severity; b_i is impact; $q_{0.9}(b)$ is the ninth decile of impact over the evaluation batch; $\vee$ is logical OR.

Equation (15) proposes a safety-review marker for high-criticality cases, intended for prospective deployment design. It is not used as a hidden scheduling rule in the numerical comparison, which remains fully specified by Equations (1)–(11).

➢ *Governance, Failure Modes and Proposed Field Study*

A production implementation must define an escalation path for critical and overdue incidents, an independent authorization gate for corrective actions that change payment data or service routing, and a post-action verification record. The optimiser should rank eligible work only. In particular, neither a predicted deadline breach nor a high business-impact score may override change-management policies, security approvals, or rollback requirements. Independent incident auditors should be able to reconstruct the features that determined a decision (Beyer et al., 2018; International Organization for Standardization, 2022).

A prospective evaluation should begin with shadow operation using historical but access-controlled incident events. The shadow phase would estimate calibration, ranking stability, and fair treatment across affected service classes without instructing engineers to act. After stakeholder review, a randomized stepped-wedge or interrupted time-series field comparison could assess actual closure time, SLA penalties, safety incidents, manual overrides, and unintended shift of delay between services. Each study arm should use equivalent alert intake rules and consistent incident definitions.

Data minimization and operational confidentiality are paramount. A public repository should contain only the synthetic benchmark, transformations, test assertions, and anonymized aggregate statistics. The study should neither publish real payment-service architecture nor include customer records. Model monitoring would need version control, override logging, policy-drift metrics, and predefined shutdown conditions. These are necessary operational safeguards rather than optional documentation details.

IV. RESULTS OF THE SYNTHETIC COMPUTATIONAL STUDY

All values in this section are generated outcomes under the explicit conditions in Section 3. They are conditional on an artificial workload distribution, predicted-duration error, equal server capabilities, and a fixed batch release. They do not measure payment-infrastructure performance. The comparison is an internal computational benchmark rather than an independent field validation of SLA-OPT.

Table 2 Mean Computed Outcomes from 36 Synthetic Replications of 42 Incidents Per Policy.

Scheduling policy	Raw breach	Weighted breach	Weighted tardiness (min)	Mean closure (min)	P90 closure (min)
Rule-based	74.5%	63.2%	48.6	154.3	250.3
Genetic algorithm	59.5%	45.0%	28.8	129.9	238.1
SLA-OPT	71.3%	44.6%	23.5	142.0	241.6

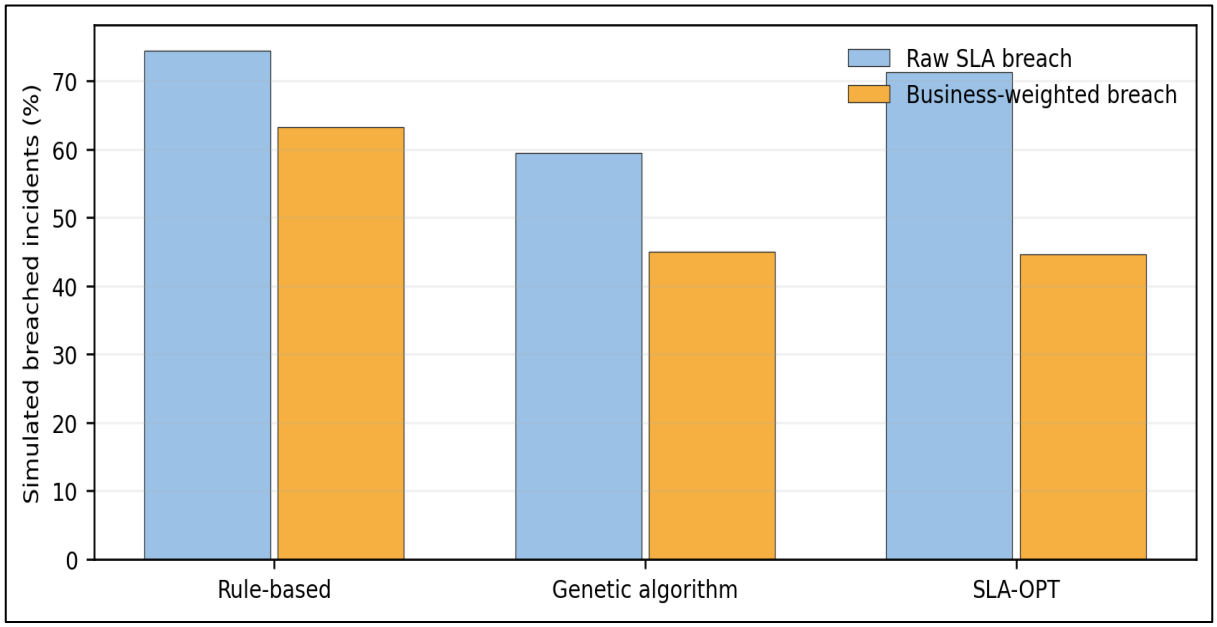


Fig 2 Comparison of Raw and Impact-Weighted SLA Breach Fractions Across the three Scheduling Policies.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov1299

Figure 2 and Table 2 demonstrate why the choice of endpoint matters. The GA produces the lowest raw fraction of tickets exceeding an SLA deadline, whereas SLA-OPT is slightly lower on the business-weighted breach fraction in the chosen workload. The difference between these claims is not semantic: it shows that weighting decisions can reallocate delay across tickets. A study reporting only an integrated efficiency index would conceal this trade-off. The rule-based baseline produces the largest breach fractions under the model parameters.

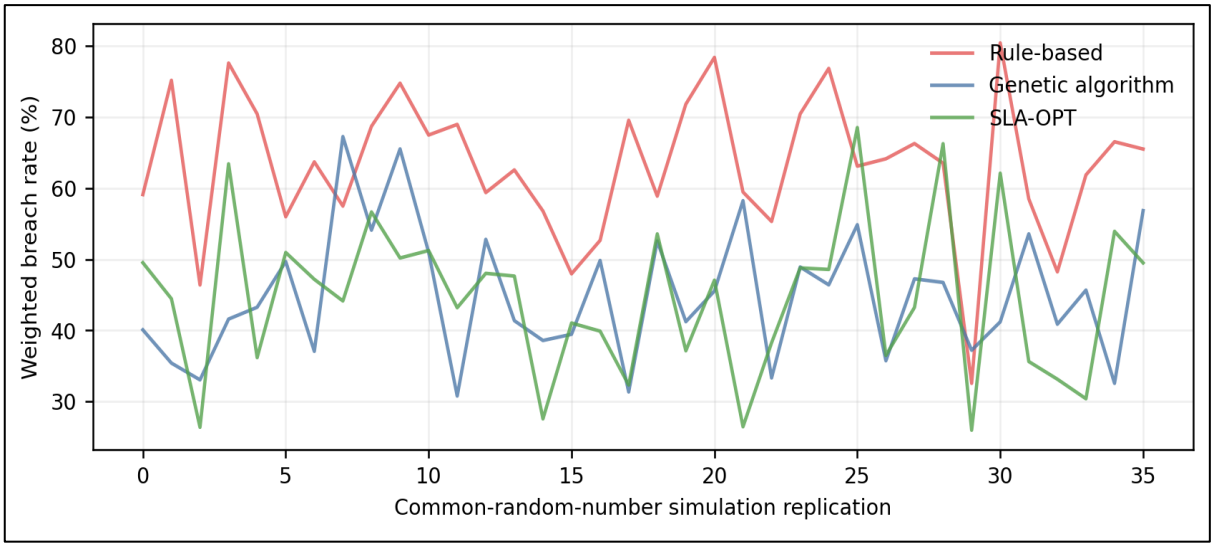


Fig 3 Paired Replication Trajectories of Business-Weighted SLA Breaches Under Common Incident Workloads.

Figure 3 preserves the paired structure of the experiment. All three methods experience heterogeneous workloads, which shift their observed breach rates across seeded replications. The plot is a diagnostic of scenario sensitivity, not a confidence interval for any real incident population. The appropriate comparison should use within-replication contrasts rather than comparing unrelated samples.

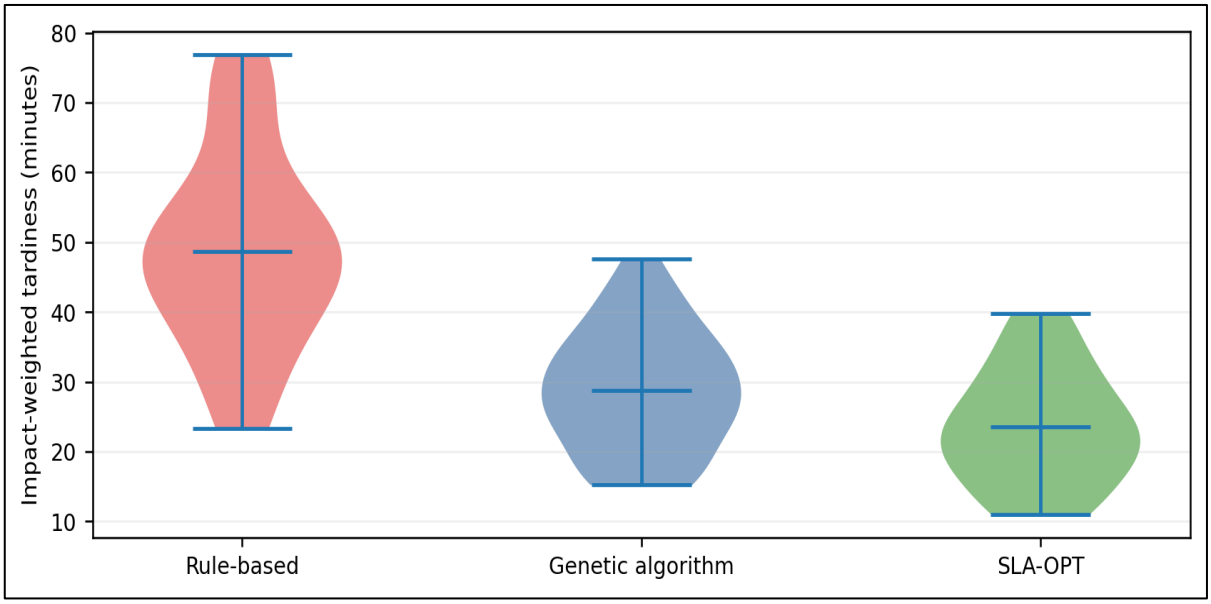


Fig 4 Distribution of Impact-Weighted SLA Tardiness Across Simulated Workload Replications.

The violin plots in Figure 4 show continuous delay exposure. SLA-OPT yields the lowest mean weighted tardiness in the present experiment, despite a worse unweighted breach frequency than the GA. These results support H1 only for the declared synthetic cost structure and do not establish general superiority. Their plausibility should be examined under alternative arrival processes, heterogeneous skills, staffing levels, and duration-estimation errors.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov1299

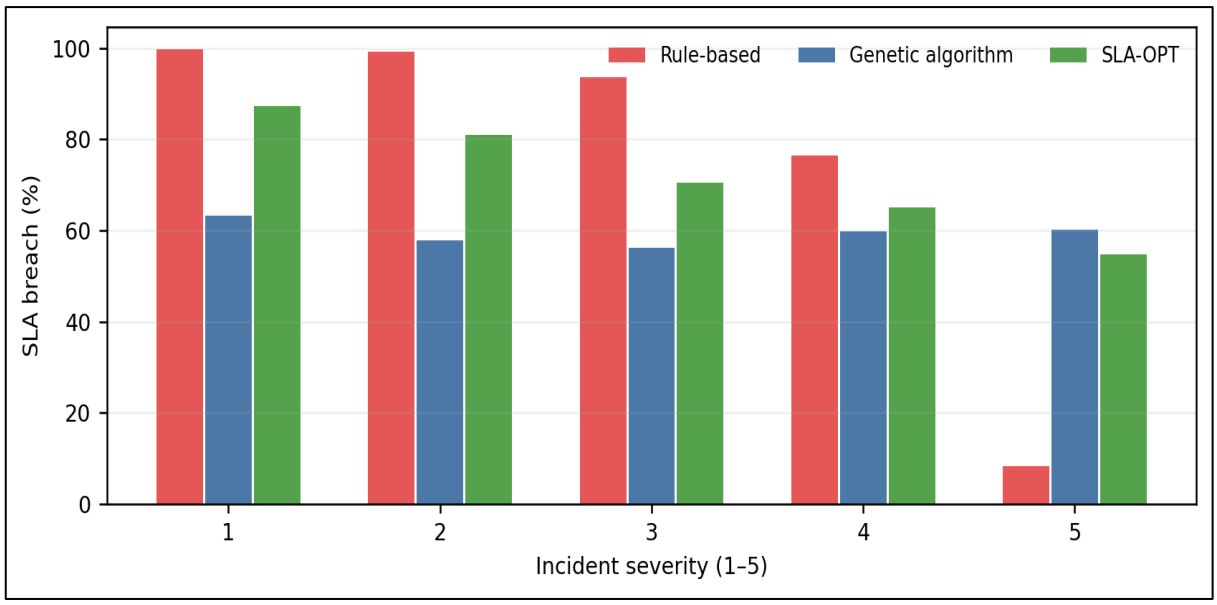


Fig 5 Severity-Stratified Breach Rates for All three Scheduling Strategies in the Synthetic Cohort.

Figure 5 examines whether the average results conceal variation by incident severity. A real governance review would also stratify by service type, financial exposure, incident age, and suspected security effect. Strong performance on the pooled sample does not authorize use in a safety-critical, compliance-sensitive or customer-impacting subgroup without independent evaluation.

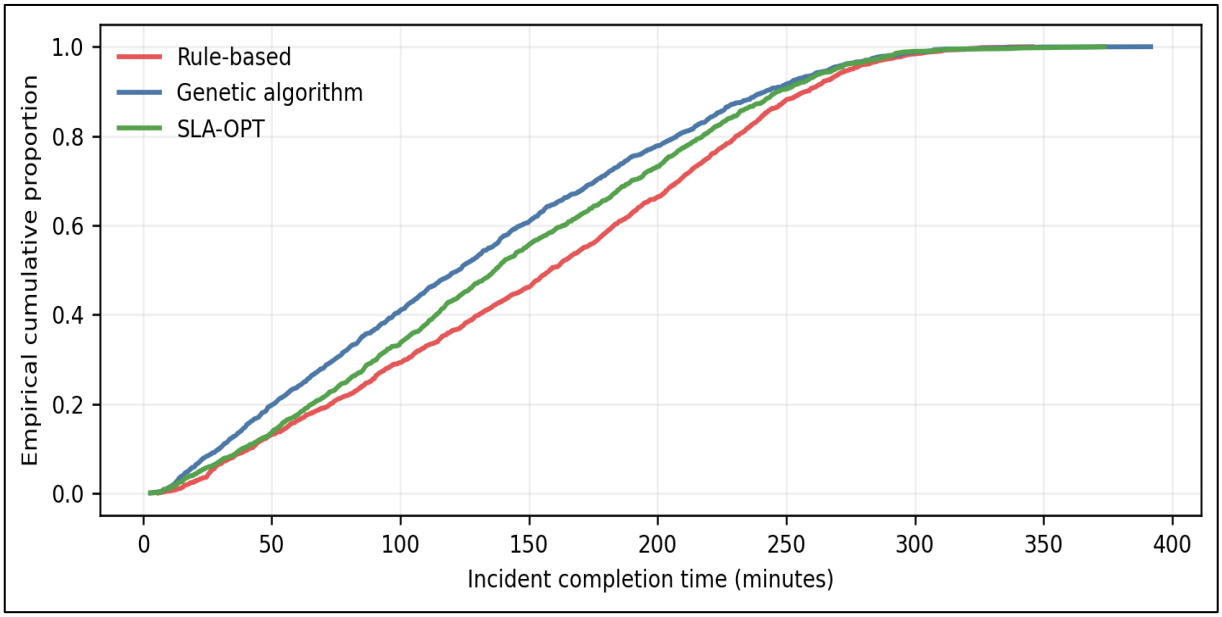


Fig 6 Empirical Distribution of Simulated Corrective-Action Completion Times for Competing Policies.

Figure 6 shows the full completion-time distribution rather than one aggregate mean. The GA leads on raw deadline compliance, but the proposed method places more effort on high-weight cases. Consequently, lower weighted loss can coexist with longer completion times for some lower-weight incidents. This is a governance design choice requiring organizational consent and monitoring.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov1299

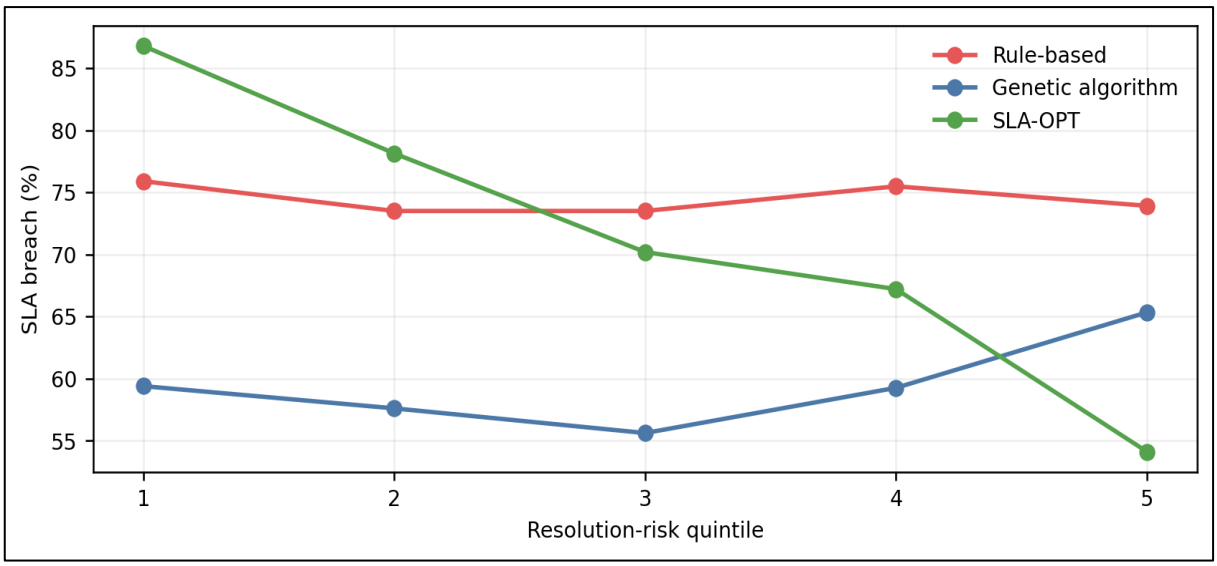


Fig 7 SLA Breach Rates Across Resolution-Risk Quintiles for the three Simulation Strategies.

Figure 7 visualizes performance conditional on resolution risk. The risk variable was generated as a synthetic case feature and was not calibrated against observed repair failures; results are therefore a sensitivity signal only. A deployed version should compare predicted risk with realized interruptions and report calibration metrics before retaining a multiplicative risk term in a scheduling policy.

V. DISCUSSIONS

The principal finding is a multi-objective trade-off. SLA-OPT improves modeled business-weighted tardiness relative to the two comparators, yet the genetic algorithm reduces the fraction of incidents missing deadlines. The distinction reflects the different objectives and the scheduling problem's combinatorial structure. A scheduler explicitly optimising one metric should not be expected to dominate on every other metric, particularly under heterogeneous incident weights. This result cautions against declaring a proposed method superior based on a single favorable measure.

The rule-based comparator remains operationally credible in settings where feature quality is low, computation is restricted, or priority decisions must be immediately explained. Its deterioration in the present synthetic environment should not be read as a universal weakness. Severity-only policies may be appropriate when severity is carefully calibrated to loss or when emergency cases require strict preemption. Such conditions would alter the utility of the proposed score (Pinedo, 2016; Beyer et al., 2016).

The genetic algorithm offers a richer search of candidate schedules but is evaluated here using noisy predicted durations. Its performance depends on how faithfully that prediction represents real work, how often cases arrive, and how much time the optimizer may consume at each dispatch epoch. Our implementation uses a fixed budget, modest population, and permutation mutation. Different GA operators, local search, mixed-integer optimizers or admissible dominance rules could change the ranking. Accordingly, the comparison is with one transparent GA configuration, not with the abstract class of all evolutionary algorithms (Goldberg, 1989; Deb et al., 2002).

SLA-OPT is designed as a lightweight, inspectable alternative that updates urgency using remaining time and case attributes. Its principal risk is coefficient arbitrariness. A synthetic coefficient that appears sensible may underweight rare but extreme events, overvalue a noisy business-impact score, or create persistent starvation when arrival load is high. A deployed system must learn or calibrate its inputs without leaking outcome information from the future, record versioned score components, and permit human override with a documented reason (Guo et al., 2017; Mitchell et al., 2019).

The experiment does not model all dependencies important to payment systems. A corrective action may need separate database, application, and security personnel; changes can be blocked by approvals; rollback time may dominate repair execution; and multiple correlated incidents may share one underlying cause. In these conditions, treating each incident as an independent job on identical servers may exaggerate the benefit of reordering. Future studies should use resource-eligibility matrices, dependency graphs, stage-specific service times, and time-varying arrivals. Event logs and process mining can provide the observable workflow structure needed to test these assumptions (Dumas et al., 2018; van der Aalst, 2016).

A second implication concerns attribution. Replaying synthetic cases establishes algorithmic consistency, not causal improvement in an operating payment service. Even a retrospective real-incident replay would remain observational if staff interventions, load fluctuations, and severity labels differed across periods. The stronger design would use a preregistered controlled rollout with unchanged SLA definitions, clearly reported override rates, and independent verification of closure and recurrence. A deployment claim should include operational safety, privacy and compliance outcomes alongside breach reductions.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov1299

VI. LIMITATIONS AND FUTURE RESEARCH

The primary limitation is the absence of real incident histories. The input distributions, severity-impact relationship, uncertainty noise and SLA deadlines are specified rather than estimated. Results therefore describe a designed stress test, not statistically generalizable effects. The study also omits incident release-time variability, interrupted task execution, correlated failures, specialized response teams, and policy feedback into the incident-arrival process.

The synthetic proxy for business impact is not a validated currency loss, actual merchant exposure or observed contractual penalty. Consequently, the weighted breach and tardiness metrics should not be presented as revenue saved. An eventual study should specify exactly how business impact is derived and reconcile it with customer harm, operational risk, regulatory obligations, and service continuity.

The GA search parameters are not exhaustively tuned, and the proposed SLA-OPT coefficients are chosen ex ante. Robust comparative work should include multiple GA configurations, earliest-deadline-first, weighted shortest processing time, integer optimization where feasible, and computational timing. It should also stress-test duration misspecification, missing risk scores, capacity reductions, burst arrivals, critical-case gates, and fairness thresholds. Stability to parameter changes is more persuasive than one favorable benchmark.

The citations draw from service management, scheduling, predictive process monitoring, reliability and risk optimization. I did not identify a defensible set of 30 directly relevant, verifiable TechConnect Briefs papers for this specific SLA corrective-action problem; unrelated proceedings papers are therefore not inserted. Nor are publications attributed to Idoko Peter Idoko, Onuh Matthew Ijiga, or Otugene Victor Bamigwojo included without independently verifiable bibliographic support. The final DOI metadata should be checked against journal production records before submission.

VII. CONCLUSION

SLA-OPT offers a formally specified rolling-horizon approach to prioritising corrective action according to severity, business impact, deadline pressure, and resolution risk. The reproducible synthetic evaluation demonstrates that policy rankings change with the selected measure of success: the GA attains fewer raw breaches, while SLA-OPT attains slightly fewer business-weighted breaches and lower weighted tardiness under the declared assumptions. This mixed result is more informative than an unsupported claim of overall superiority. The next stage is independent field validation using authorized incident histories, calibrated prediction models, formal safety gates, and a preregistered experimental protocol.

REFERENCES

[1.] Abolade, Y. A., Adegoke, A., Noi, S. M., Nwoye, C. M., Ajayi, O., Kyrian, O., Akagbue, B. O., Siame, T., & Bamidele, A. A. (2024). Promoting health: Introducing an eco-friendly herbal mosquito repellent extracted from local sweet orange peels *GSC Biological and Pharmaceutical Sciences,* 2024, 26(01), 315–327 DOI: https://doi.org/10.30574/gscbps.2024.26.1.0014

[2.] Akanni, F., Idowu, K., Odozor, L., & Waheed, B. (2026). Cyber-financial risk analytics: Integrating business intelligence with threat modeling. International Journal of Scientific Research in Computer Science Engineering and Information Technology. https://doi.org/10.32628/CSEIT26121371

[3.] Akinleye, T. & Ijiga, O. M. (2025). Sustainable Construction Compliance Modelling Through ISO 14001, ISO 45001 and ISO 9001 Integrated Audit Data. *International Journal of Scientific Research in Civil Engineering* Vol.9, Iss. 4, Pg, 93-116 DOI: https://doi.org/10.32628/IJSRCE2594012

[4.] Akinleye, T. & Ijiga, O. M. (2026). A Novel COSHH Risk Ranking Algorithm for Hazardous Substance Control in Construction and Engineering Workplaces. *International Journal of Innovative Science and Research Technology*, Vol. 11, Iss. 7. https://doi.org/10.38124/ijisrt/26jul1158

[5.] Akinleye, T., Idoko, D. O. & Enyejo, L. A. (2026). Sustainability-Driven Root Cause Analysis of Construction Incidents Using Environmental, Quality and Safety Performance Indicators *International Journal of Scientific Research and Modern Technology*, 5(6), 203–274. https://doi.org/10.38124/ijsrmt.v5i6.1541

[6.] Amershi, S., Begel, A., Bird, C., DeLine, R., Gall, H., Kamar, E., Nagappan, N., Nushi, B., & Zimmermann, T. (2019). Software engineering for machine learning: A case study. Proceedings of the 41st International Conference on Software Engineering: Software Engineering in Practice, 291–300. https://doi.org/10.1109/ICSE-SEIP.2019.00042

[7.] Bamidele, A. A,. Kondeti, H., Ijiga, O. M. (2026). Development of a Hybrid Machine Learning Algorithm for Impurity Reduction in Nitrosamine Formation Under Variable Pharmaceutical Processing Conditions Compared with Gradient Boosting and Artificial Neural Networks International Journal of Scientific Research in Chemistry Volume 11, Issue 2 Pg. 64 -94 doi : https://doi.org/10.5281/zenodo.23860513

[8.] Bamidele, A. A., Justin, O., Adanne, S. (2025). Integrating machine learning with environmental chemistry to forecast pollutant releases in coatings production. World Journal of Advanced Research and Reviews, 2025, 27(3), 1061-1072. Article DOI: https://doi.org/10.30574/wjarr.2025.27.3.3153

[9.] Bamidele, A. A., Rotimi, A. M., James, C. O., Obinna, J. (2025). Coupling environmental chemistry with

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov1299

Artificial Intelligence for coastal hazard risk assessment under climate change. International Journal of Science and Research Archive, 2025, 16(03), 498–507. Article DOI: https://doi.org/10.30574/ijsra.2025.16.3.2572.

[10.] Barlow, R. E., & Proschan, F. (1965). Mathematical theory of reliability. Wiley.

[11.] Bertsimas, D., & Sim, M. (2004). The price of robustness. Operations Research, 52(1), 35–53. https://doi.org/10.1287/opre.1030.0065

[12.] Beyer, B., Jones, C., Petoff, J., & Murphy, N. R. (Eds.). (2016). Site reliability engineering: How Google runs production systems. O'Reilly Media. https://sre.google/sre-book/table-of-contents/

[13.] Beyer, B., Murphy, N. R., Rensin, D. K., Kawahara, K., & Thorne, S. (Eds.). (2018). The site reliability workbook: Practical ways to implement SRE. O'Reilly Media. https://sre.google/workbook/table-of-contents/

[14.] Breck, E., Cai, S., Nielsen, E., Salib, M., & Sculley, D. (2017). The ML test score: A rubric for ML production readiness and technical debt reduction. 2017 IEEE International Conference on Big Data, 1123–1132. https://doi.org/10.1109/BigData.2017.8258038

[15.] Chen, R., Wang, C., & Liu, Z. (2021). Intelligent genetic algorithm for workforce scheduling considering service level in IT maintenance service. Journal of System Simulation, 33(3).

[16.] Committee on Payment and Settlement Systems & International Organization of Securities Commissions. (2012). Principles for financial market infrastructures. Bank for International Settlements. https://www.bis.org/cpmi/publ/d101a.pdf

[17.] Deb, K., Pratap, A., Agarwal, S., & Meyarivan, T. (2002). A fast and elitist multiobjective genetic algorithm: NSGA-II. IEEE Transactions on Evolutionary Computation, 6(2), 182–197. https://doi.org/10.1109/4235.996017

[18.] Doe Gadekah, S, B. & Enyejo, J. O. (2025). Personalization Versus Privacy in Digital Marketing Examining the Interactive Effects of Data Transparency Perceived Intrusiveness and Consumer Trust on Purchase Decisions *International Journal of Management and Organizational Research* Vol. 4, Iss. 6 Page No: 182-199 DOI: https://doi.org/10.54660/IJMOR.2025.4.6.182-199

[19.] Doe Gadekah, S, B. & Enyejo, J. O. (2025). The Influence of Brand Community Participation and User-Generated Content on Customer-Based Brand Equity Repurchase Intention and Electronic Word-of-Mouth in Direct-to-Consumer Markets *International Journal of Management and Organizational Research* Vol, 4, Iss. 3 Page No: 183-203 DOI: https://doi.org/10.54660/IJMOR.2025.4.3.183-203

[20.] Doe Gadekah, S, B. & Enyejo, J. O. (2026). The Effects of Generative AI Advertising Disclosure on Brand Authenticity Perceptions Purchase Intention and Consumer Skepticism Among Generation Z Consumers *International Journal of Scientific Research in Humanities and Social Sciences* Vol, 3 Iss 2, Pg. 137-169 doi : https://doi.org/10.32628/IJSRHSS253293

[21.] Dumas, M., La Rosa, M., Mendling, J., & Reijers, H. A. (2018). Fundamentals of business process management (2nd ed.). Springer. https://doi.org/10.1007/978-3-662-56509-4

[22.] Fani Sani, M., Vazifehdoostirani, M., Park, G., & van der Aalst, W. M. P. (2023). Performance-preserving event log sampling for predictive monitoring. Journal of Intelligent Information Systems, 61, 53–82. https://doi.org/10.1007/s10844-022-00775-9

[23.] Fioretto, S., & Masciari, E. (2025). A comparative analysis of predictive process monitoring: Object-centric versus classical event logs. Knowledge and Information Systems, 67, 7355–7398. https://doi.org/10.1007/s10115-025-02461-y

[24.] Folorunso, T. N., Idowu, K., Odozor, L., & Waheed, B. (2026). Cyber-financial risk analytics: Integrating business intelligence with threat modeling. International Journal of Scientific Research in Computer Science, Engineering and Information Technology, 12(2), 399–415. https://doi.org/10.32628/CSEIT26121371

[25.] Frank, U. P. (2026a). Healthcare quality metrics: From data collection to performance improvement; A narrative review of measurement frameworks and improvement methods in United States health care. UKR Journal of Medicine and Medical Research. https://doi.org/10.5281/zenodo.22105362

[26.] Frank, U. P. (2026b). Optimizing medical inventory management through predictive analytics: Implications for healthcare quality, cost reduction, and patient safety; A narrative review. UKR Journal of Education and Literature, 2(1), 77–92. https://doi.org/10.5281/zenodo.22105843

[27.] Frank, U. P. (2026c). Artificial intelligence for infection surveillance: Improving early detection of healthcare-associated infections. USR Journal of Multidisciplinary Research and Studies.

[28.] Garey, M. R., & Johnson, D. S. (1979). Computers and intractability: A guide to the theory of NP-completeness. W. H. Freeman.

[29.] Goldberg, D. E. (1989). Genetic algorithms in search, optimization, and machine learning. Addison-Wesley.

[30.] Guo, C., Pleiss, G., Sun, Y., & Weinberger, K. Q. (2017). On calibration of modern neural networks. Proceedings of the 34th International Conference on Machine Learning, 1321–1330. https://proceedings.mlr.press/v70/guo17a.html

[31.] Hyndman, R. J., & Athanasopoulos, G. (2021). Forecasting: Principles and practice (3rd ed.). OTexts. https://otexts.com/fpp3/

[32.] Idowu, K., Cherotich, V., Aniebonam, E. E., Odozor, L., & During, A. (2025). Integrating AI and machine learning into cyber risk management for critical utility systems. International Journal of Artificial Intelligence Engineering and Technology. https://doi.org/10.54660/IJAIET.2025.6.2.32-48

[33.] Idowu, K., Leslie, A., Twayana, S., Nitzan, Y., & Waheed, B. O. (2026). Real-time detection of DDoS

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov1299

and phishing attacks for enhanced banking security. International Journal of Scientific Research in Science, Engineering and Technology, 13(1), 572–593. https://doi.org/10.32628/IJSRSET2613313

[34.] Ijiga, O. M. & Akinleye, T. (2024). A Framework for Integrating Environmental Permit and Consent Management into Construction Project Delivery: Reducing Delays and Ensuring Regulatory Compliance. *International Journal of Scientific Research in Science, Engineering and Technology*, vol. 11, no. 3, pp. 706–732, Jun. 2024, doi: 10.32628/IJSRSET261254

[35.] International Organization for Standardization. (2018). ISO 31000:2018 Risk management—Guidelines. ISO.

[36.] International Organization for Standardization. (2018). ISO/IEC 20000-1:2018 Information technology—Service management—Part 1: Service management system requirements. ISO.

[37.] International Organization for Standardization. (2019). ISO 22301:2019 Security and resilience—Business continuity management systems—Requirements. ISO.

[38.] International Organization for Standardization. (2022). ISO/IEC 27001:2022 Information security, cybersecurity and privacy protection—Information security management systems—Requirements. ISO.

[39.] Kleinrock, L. (1975). Queueing systems, volume 1: Theory. Wiley.

[40.] Law, A. M. (2015). Simulation modeling and analysis (5th ed.). McGraw-Hill Education.

[41.] Li, M., Li, Z., Yin, K., Nie, X., Zhang, W., Sui, K., & Pei, D. (2022). Causal inference-based root cause analysis for online service systems with intervention recognition. Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining. https://doi.org/10.1145/3534678.3539041

[42.] Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems, 30. https://arxiv.org/abs/1705.07874

[43.] Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I. D., & Gebru, T. (2019). Model cards for model reporting. Proceedings of the Conference on Fairness, Accountability, and Transparency, 220–229. https://doi.org/10.1145/3287560.3287596

[44.] Montgomery, D. C. (2019). Introduction to statistical quality control (8th ed.). Wiley.

[45.] National Institute of Standards and Technology. (2024). The NIST Cybersecurity Framework (CSF) 2.0 (NIST CSWP 29). https://doi.org/10.6028/NIST.CSWP.29

[46.] Niculescu-Mizil, A., & Caruana, R. (2005). Predicting good probabilities with supervised learning. Proceedings of the 22nd International Conference on Machine Learning, 625–632. https://doi.org/10.1145/1102351.1102430

[47.] Ogundola, T. (2020). Advances in quantum-secure banking: Cryptographic solutions. International Journal of Science and Research Archive, 1(1), 228–230. https://doi.org/10.30574/ijsra.2020.1.1.0048

[48.] Ogundola, T. (2022). Quantum-resistant cryptographic techniques for financial institutions. International Journal of Scientific Research and Modern Technology, 1(11), 18–20. https://doi.org/10.38124/ijsrmt.v1i11.586

[49.] Ogundola, T. O. (2024) Barriers to implementing quantum-resistant encryption in financial institutions, World Journal of Advanced Research and Reviews, 2024, 22(03), 2323-2337, DOI: https://doi.org/10.30574/wjarr.2024.22.3.1938

[50.] Ogundola, T. O. (2025). Post-quantum cryptography for secure banking transactions. International Journal of Scientific Research and Modern Technology, 4, 40-42.

[51.] Pearl, J. (2009). Causality: Models, reasoning, and inference (2nd ed.). Cambridge University Press.

[52.] Pinedo, M. L. (2016). Scheduling: Theory, algorithms, and systems (5th ed.). Springer. https://doi.org/10.1007/978-3-319-26580-3

[53.] Qazi, F., Kwak, D., Khan, F. G., Ali, F., & Khan, S. U. (2024). Service level agreement in cloud computing: Taxonomy, prospects, and challenges. Internet of Things, 25, 101126. https://doi.org/10.1016/j.iot.2024.101126

[54.] Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?" Explaining the predictions of any classifier. Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 1135–1144. https://doi.org/10.1145/2939672.2939778

[55.] Rockafellar, R. T., & Uryasev, S. (2000). Optimization of conditional value-at-risk. Journal of Risk, 2(3), 21–41. https://doi.org/10.21314/JOR.2000.038

[56.] Ruiz, R., & Maroto, C. (2005). A comprehensive review and evaluation of permutation flowshop heuristics. European Journal of Operational Research, 165(2), 479–494. https://doi.org/10.1016/j.ejor.2004.04.017

[57.] Sculley, D., Holt, G., Golovin, D., Davydov, E., Phillips, T., Ebner, D., Chaudhary, V., Young, M., Crespo, J.-F., & Dennison, D. (2015). Hidden technical debt in machine learning systems. Advances in Neural Information Processing Systems, 28. https://proceedings.neurips.cc/paper/2015/hash/86df7dcfd896fcaf2674f757a2463eba-Abstract.html

[58.] Shapiro, A., Dentcheva, D., & Ruszczyński, A. (2021). Lectures on stochastic programming: Modeling and theory (3rd ed.). SIAM. https://doi.org/10.1137/1.9781611976595

[59.] Tax, N., Verenich, I., La Rosa, M., & Dumas, M. (2017). Predictive business process monitoring with LSTM neural networks. In Advanced Information Systems Engineering (pp. 477–492). Springer. https://doi.org/10.1007/978-3-319-59536-8_30

[60.] Teinemaa, I., Dumas, M., La Rosa, M., & Maggi, F. M. (2019). Outcome-oriented predictive process monitoring: Review and benchmark. ACM

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov1299

Transactions on Knowledge Discovery from Data, 13(2), Article 17. https://doi.org/10.1145/3301300

[61.] van der Aalst, W. M. P. (2016). Process mining: Data science in action (2nd ed.). Springer. https://doi.org/10.1007/978-3-662-49851-4

[62.] Vanhoucke, M. (2012). Project management with dynamic scheduling: Baseline scheduling, risk analysis and project control (2nd ed.). Springer. https://doi.org/10.1007/978-3-642-25175-7

[63.] Verenich, I., Dumas, M., La Rosa, M., Maggi, F. M., & Teinemaa, I. (2019). Survey and cross-benchmark comparison of remaining time prediction methods in business process monitoring. ACM Transactions on Intelligent Systems and Technology, 10(4), Article 34. https://doi.org/10.1145/3331449

[64.] Yu, Q., Zhao, N., Li, M., Li, Z., Wang, H., Zhang, W., Sui, K., & Pei, D. (2024). A survey on intelligent management of alerts and incidents in IT services. Journal of Network and Computer Applications, 224, 103842. https://doi.org/10.1016/j.jnca.2024.103842

[65.] Zadrozny, B., & Elkan, C. (2002). Transforming classifier scores into accurate multiclass probability estimates. Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 694–699. https://doi.org/10.1145/775047.775151.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov1299

APPENDIX

- *Appendix A. Reproduction and Data Governance*

The accompanying repository provides a single Python experiment, input-generation logic, the synthetic incident-by-policy panel, replication summary, aggregated policy summary and seven full-resolution figures. Install NumPy, pandas and Matplotlib, then execute experiment.py in its own directory. The script uses a fixed initial random seed and advances that seed deterministically for each of 36 independent workloads. The GA draws its own declared seeds, and every policy receives the same realised work times within each replication. Re-running the script should reproduce the numeric outputs and all charts without manual data editing.

Internal assertions to add in an operational-grade implementation include: case identifiers are unique inside each replication; completion times are nonnegative and finite; no responder services overlapping tasks; the SLA indicator equals the strict deadline comparison; no scheduling policy can inspect actual duration before a decision; all policies cover the same cases; and all impact weights are positive. The original reproducibility script intentionally contains no network requests, payment API calls or third-party data access.

A field replication would require consent and permission for derived service-management records, encryption and access restrictions on topology, retention controls, security review of event features, and independent adjudication of disputed incident deadlines. Real financial-service event logs should not be placed in a public research package. Any claims concerning company deployment or measured financial benefit must be supported by separate audited evidence. A formal manuscript should also disclose authorship responsibilities, conflicts of interest, funding and approval status accurately.