Power System Protection and Control
ISSN:1674-3415

REDEFINING BUSINESS PROCESS MANAGEMENT: THE SYNERGY OF AI AND INTELLIGENT AUTOMATION

Ronakkumar shah

58, SOMNATH SOCI-, NR. BHAGAT VADI
BODELI, CHHOTAUDEPURPIN: 391135, GUJARAT ,INDIA.
43142 WHELPLEHILL TER ASHBURN VA 20148 UNITED STATES.

Received: 07/05/2026 **Revised**: 10/06/2026 **Accepted**: 02/07/2026

***ABSTRACT*:**

Business process management (BPM) has, for two decades, organised enterprise work around explicit process models, deterministic routing rules, and human judgement at every point the model does not anticipate. Robotic process automation (RPA) extended this paradigm to the desktop, mechanising well-defined, rule-based tasks, but left the harder, judgement-dependent steps — interpreting an unstructured claim narrative, weighing ambiguous evidence, drafting a context-sensitive response — to human operators, capping the share of any given process that automation could realistically absorb. This paper argues that the recent maturation of machine learning for process mining and prediction, combined with the reasoning and language capability of large language models (LLMs), enables a qualitatively different architecture for BPM, in which prediction, judgement, and language understanding are embedded directly into the process runtime rather than bolted on as isolated point solutions. We formalise an AI-augmented process as a conventional workflow extended with a decision-intelligence layer that scores each pending step for predicted outcome and model confidence, a continuous process-learning loop that retrains on operator feedback rather than remaining static after initial deployment, and a human-in-the-loop governance mechanism that routes low-confidence or high-stakes decisions to a human reviewer without silently automating around them. A reference architecture composed of a predictive-monitoring service, a cognitive task-automation layer built on retrieval-augmented LLM inference, an adaptive routing engine, and a review console is described, together with the data-lifecycle, legacy-integration, and explainability considerations required to deploy it responsibly. We evaluate the framework through a nine-month case study of intelligent automation in insurance claims processing, presented through two figures: a comparison of process performance across traditional BPM, RPA-augmented BPM, and the proposed AI-augmented configuration, and a learning-curve and confidence-threshold sensitivity analysis. Straight-through processing rose from 22% under traditional BPM to 79% under the full framework, median cycle time fell from 46 to 6.5 hours, and residual error rate at a high confidence threshold fell below 1.5%, while the sensitivity analysis quantifies the trade-off between automation coverage and human-escalation volume. We conclude with a discussion of the organisational, technical, and governance limitations of the approach and outline directions for adaptive, increasingly autonomous process management that nonetheless preserves meaningful human oversight.

***Keywords**: Business Process Management, Intelligent Automation, Robotic Process Automation, Large Language Models, Process Mining, Predictive Process Monitoring, Human-In-The-Loop, Hyper Automation*

INTRODUCTION

Business process management emerged as a discipline on the premise that organisational work, once made explicit as a process model, could be measured, optimised, and automated with the same rigour applied to any other engineered system [1]. Two decades of BPM practice have delivered substantial value along exactly this premise: processes that were once tacit and inconsistent across branches or shifts became documented, comparable, and improvable. The premise's limitation, however, has always been visible at the boundary of what a process model can specify in advance. Deterministic routing rules handle the steps whose logic can be enumerated — route to underwriting if the policy value exceeds a threshold, require a supervisor signature above a spend limit — but hand off to a human the moment a step requires interpreting something the model's designers could not fully anticipate: an ambiguous claim narrative, a document that does not match its expected template, a customer request phrased in a way the intake form did not foresee.

DOI: 10.46121/pspc.54.3.14

This boundary is not a minor edge case at the margins of an otherwise fully automated process; in many judgement-intensive industries — insurance, healthcare administration, financial-services back-office operations, government benefits processing — the judgement-dependent steps are frequently the steps that determine the majority of a process's total cycle time and cost, precisely because they are the steps an organisation has historically been unable to automate. A process that automates ninety percent of its steps but leaves the single most time-consuming, judgement-dependent step untouched may show little practical improvement in end-to-end cycle time, a pattern repeatedly observed in BPM maturity assessments across industries [14].

Robotic process automation extended the reach of deterministic automation into the individual desktop applications a process touches, scripting the clicks, keystrokes, and data transcription that previously required a human operator to bridge one system to another [2]. RPA delivered real, measurable efficiency gains, particularly in high-volume, low-variability back-office processes, but it inherited BPM's fundamental limitation rather than resolving it: an RPA bot automates the mechanical execution of a step whose logic is already fully specified, and simply fails, or silently misroutes, the moment a case falls outside the pattern its scripts anticipated. Industry surveys of RPA programs have repeatedly found that bot failure rates rise sharply with input variability, and that the resulting exception volume routed back to human operators can offset a large share of the labour savings the program was built to capture [3].

Two more recent technical developments change this picture. First, machine-learning-based process mining and predictive process monitoring can now forecast, from a partially completed case trace, the probability that a case will breach a service-level target, require escalation, or ultimately be denied, days before a rule-based system would have surfaced any exception [4], [5]. Second, large language models have reached a level of general-purpose language understanding and generation capability sufficient to draft, summarise, classify, and reason over unstructured text — claim narratives, medical notes, customer correspondence, contract clauses — with an accuracy that, for a meaningful share of cases, approaches or exceeds that of a junior human reviewer on narrowly scoped tasks [6], [7]. Neither development alone redefines BPM: predictive process monitoring without a language-understanding capability still cannot act on unstructured case content, and an LLM without embedding in a governed process runtime is a point solution rather than an operating model.

This paper's central argument is that the synergy of these two developments — predictive process intelligence and language-capable cognitive automation, composed within a single governed architecture — enables a redefinition of BPM in which the process model is no longer a static specification of deterministic routing but a living system that predicts outcomes, drafts and evaluates judgement-dependent steps, learns continuously from the corrections its human reviewers make, and escalates to a human precisely when, and only when, its own confidence indicates it should. Four contributions follow from this position. First, we formalise an AI-augmented process as an extension of the conventional workflow model in which each transition additionally carries a predicted-outcome distribution and a model-confidence score, rather than treating prediction as an external, disconnected analytics exercise. Second, we describe a continuous process-learning loop in which operator corrections to model-drafted or model-routed decisions are captured as first-class training signal, closing the loop between production operation and model improvement that most current automation deployments leave open. Third, we specify a human-in-the-loop governance mechanism, structurally similar in spirit to graduated enforcement mechanisms proposed for workflow governance more broadly [8], that routes decisions to a human reviewer based on an explicit, auditable confidence threshold rather than an implicit, undocumented judgement embedded in a model's deployment configuration. Fourth, we evaluate the resulting framework not only for automation coverage but for the trade-off between that coverage and the residual error and human-escalation burden it produces, since a framework that reports high straight-through processing at an undisclosed cost in downstream error would materially overstate its practical value.

The remainder of the paper proceeds as follows. Section 2 reviews the evolution of BPM, the limits of RPA, machine learning for process mining, and the emerging role of LLMs in enterprise automation. Section 3 presents the proposed AI-augmented BPM architecture and its formal model. Section 4 develops the specific mechanisms — predictive monitoring, cognitive task automation, adaptive routing, and human-in-the-loop governance — that compose the architecture. Section 5 addresses implementation considerations spanning data lifecycle, legacy integration, and explainability. Section 6 reports a nine-month case study in insurance claims processing, including two figures summarising process-performance comparisons and learning dynamics. Section 7 discusses implications and limitations, and Section 8 concludes.

DOI: 10.46121/pspc.54.3.14

BACKGROUND AND RELATED WORK

The Evolution of Business Process Management

Classical BPM formalises processes using notations such as BPMN and executes them through workflow engines offering well-understood soundness guarantees over control flow — the absence of deadlock, the guarantee that every started case eventually reaches a terminal state — but treats the content of a task's execution as opaque to the process model itself [1], [9]. Process mining subsequently made it possible to reconstruct the process actually executed from event logs and compare it against the reference model, surfacing the divergence between designed and executed process that motivates a large share of contemporary BPM tooling investment [10].

More recent BPM scholarship has begun to characterise a maturity progression from process documentation, through process automation of individual steps, to what industry analysts term hyperautomation — the systematic combination of RPA, process mining, and increasingly AI-based decision capability across an end-to-end process rather than at isolated points within it [15]. This progression frames the contribution of the present paper: rather than proposing hyperautomation as an aspirational label, we specify the formal architecture, the governance mechanism, and the evaluation methodology by which such a combination can be built, deployed, and audited as a coherent system.

2.2 Robotic Process Automation and Its Limits

RPA tools automate interaction with existing application interfaces by scripting the sequence of UI actions a human operator would otherwise perform, and have been widely adopted precisely because they require no change to the underlying systems a process touches [2]. The approach's core limitation is structural rather than a matter of tooling maturity: an RPA script encodes a fixed interaction pattern, and any input variation outside that pattern — a differently formatted invoice, a system dialog the script's author did not anticipate — causes the bot to fail or misbehave, typically requiring the case to be routed to a human exception queue [3]. Because the exception rate scales with the variability of the underlying task content rather than with the volume of cases processed, RPA programs applied to judgement-dependent processes have repeatedly plateaued at a ceiling well short of full automation, a pattern documented across multiple industry deployments and one of the principal motivations for augmenting RPA with a genuine content-understanding capability [3], [11].

A further limitation of RPA specific to judgement-dependent processes is that bots are typically maintained by an automation centre of excellence separate from the business team that owns the underlying process, so that a bot script encoding an outdated business rule can persist in production well after the rule itself has changed, since the automation team may not be aware of a policy update that the business team considers routine. This organisational separation between process ownership and automation maintenance is a recurring source of the configuration drift discussed in Section 1, and one that the architecture proposed in Section 3 addresses in part by keeping the governing logic — the confidence threshold, the retrieved policy corpus — visible and versioned within the same governance console business stakeholders already use, rather than embedded in a script only the automation team can read.

2.3 Machine Learning for Process Mining and Predictive Monitoring

A substantial body of process-mining research has applied sequence models, survival analysis, and more recently deep-learning architectures to the problem of predicting, from a partial case trace, the remaining time to completion, the probability of a specific outcome, or the likelihood of a compliance violation before it occurs [4], [12]. This predictive-process-monitoring literature demonstrates that meaningful early-warning signal exists in the structure and timing of a case's trace alone, independent of the unstructured content of any individual task, and supplies the formal basis for the decision-intelligence layer proposed in Section 3.3. Predictive process monitoring, however, has generally been developed and evaluated as a standalone analytics capability layered on top of an existing process, rather than as a component embedded in the runtime routing decision of the process itself, which is the integration this paper's architecture performs.

Comparative benchmarking studies across multiple predictive-process-monitoring techniques have found that model accuracy varies considerably by process type and by how much of the case trace has been observed at prediction time, with accuracy for most techniques improving substantially once a case is at least partially through its expected lifecycle [17]. This finding directly informs the confidence-scoring design of Section 3.3: rather than committing to a single prediction accuracy figure independent of case progress, the decision-intelligence layer's

DOI: 10.46121/pspc.54.3.14

confidence value is computed fresh at each pending transition, allowing the same underlying model to correctly express lower confidence early in a case's life and higher confidence as more of its trace becomes observable.

2.4 Large Language Models in Enterprise Automation

Large language models, particularly when combined with retrieval-augmented generation over an organisation's own document corpus, have demonstrated strong performance on enterprise-relevant language tasks including document classification, summarisation, structured-field extraction from unstructured narrative text, and draft generation of routine correspondence [6], [7], [13]. Retrieval augmentation in particular addresses one of the central risks of applying a general-purpose LLM to enterprise content — that the model's fluent output may not be grounded in the organisation's specific policies or the specific case's own documentation — by conditioning generation on retrieved passages from an authoritative internal corpus rather than relying solely on the model's parametric knowledge [13]. Early enterprise deployments of LLM-based automation have tended to treat the LLM as a standalone assistant or a single automation step within an otherwise conventional RPA pipeline, rather than as one component of a jointly governed decision-intelligence architecture that also incorporates the predictive-monitoring capability of Section 2.3, which is the composition this paper proposes.

2.5 Gap Analysis and Positioning

Reviewed together, these four strands each supply a necessary but individually insufficient piece of the picture this paper assembles. BPM supplies sound process execution and control-flow guarantees but treats task content as opaque; RPA extends automation to the desktop but is structurally capped by input variability; predictive process monitoring supplies early-warning forecasting but has generally remained a standalone analytics layer; and LLM-based cognitive automation supplies content understanding and generation but has generally been deployed as an isolated point solution. This paper's contribution is architectural: composing predictive monitoring and LLM-based cognitive automation directly into the process runtime's routing decision, governed by an explicit confidence-based human-in-the-loop mechanism and a continuous learning loop, rather than introducing a fifth isolated capability alongside the other four.

PROPOSED FRAMEWORK: AI-AUGMENTED BPM ARCHITECTURE

3.1 System Architecture

The proposed architecture extends a conventional BPM or case-management runtime with four additional components. The decision-intelligence layer (Section 3.3) scores each pending task transition with a predicted-outcome distribution and a confidence value, computed from the case's structured trace history and, where the task involves unstructured content, from an LLM-based content assessment. The cognitive task-automation layer (Section 4.2) drafts or classifies the specific task content — a claim disposition, a customer response, a structured data extraction from a submitted document — using retrieval-augmented LLM inference conditioned on the organisation's policy corpus and the specific case's documentation. The adaptive routing engine (Section 4.3) consumes the decision-intelligence layer's confidence score and either commits the cognitive layer's drafted output automatically, routes it to a human reviewer for confirmation or correction, or escalates the case to a specialist queue, according to the human-in-the-loop governance mechanism of Section 4.4. The review console is the human-facing interface through which reviewers confirm, correct, or override automated drafts and dispositions, and through which the corrections captured become training signal for the continuous process-learning loop described in Section 3.4.

3.2 Formal Model of an AI-Augmented Process

We extend the labelled-transition-system representation of a workflow, $W = (S, s0, A, \rightarrow, L)$, with a decision-intelligence function $d : S \times A \rightarrow (\Theta \times [0,1])$, mapping a pending state–action pair to a predicted-outcome distribution Θ over the action's possible results and a scalar confidence value in [0,1]. Unlike the deterministic guard conditions of a conventional workflow's routing logic, d is a learned function, retrained periodically as described in Section 3.4, and its output is consumed by the adaptive routing engine rather than by the workflow engine's native control flow, preserving the underlying process engine's own soundness guarantees over case progression while layering a probabilistic recommendation on top. This separation is deliberate: d can be wrong, retrained, or temporarily disabled without altering the workflow's structural correctness, since a disabled or maximally uncertain d simply routes every case to a human reviewer, recovering the conventional BPM behaviour as a limiting case of the framework rather than a failure mode outside it.

DOI: 10.46121/pspc.54.3.14

3.3 Decision Intelligence Layer

The decision-intelligence layer computes d(s, a) from two complementary sources depending on the nature of the pending action. For actions whose relevant signal is primarily structural — sequencing, timing, and case metadata — d is estimated by a sequence model trained on historical case traces, in the tradition of predictive process monitoring [4], [12], predicting quantities such as the probability that the case will ultimately be approved, denied, or require escalation. For actions whose relevant signal is primarily the unstructured content of a document or narrative associated with the case, d additionally incorporates an LLM-based content assessment, in which the language model is prompted to classify or score the content against the same outcome categories, conditioned on relevant retrieved policy passages, with the model's own expressed uncertainty (via calibrated output probabilities or an explicit self-assessment prompt) contributing to the confidence value. The two sources are combined via a weighted ensemble whose weights are themselves tuned against the continuous learning loop's feedback, so that the framework can learn, over time and per task type, whether structural or content-based signal is the more reliable predictor for that specific task.

3.4 Continuous Process Learning Loop

Every case routed through the adaptive routing engine, whether committed automatically or routed to a human reviewer, generates a labelled training example once its outcome is known: the case's trace history and content, the decision-intelligence layer's prediction and confidence at the time, and the eventual ground-truth outcome, including any correction a human reviewer made to an automatically drafted disposition. These examples accumulate in a training-data store and are used to retrain both the structural sequence model and, where applicable through fine-tuning or updated retrieval corpora, the LLM-based content-assessment component, on a fixed retraining cadence rather than continuously in real time, so that model updates themselves pass through the same staged validation and rollout discipline recommended for any production model change before being promoted to influence live routing decisions. This closes a loop that most current point-solution automation deployments leave open: an RPA bot or a standalone LLM classifier deployed without a feedback mechanism will continue making the same category of misclassification indefinitely, whereas the framework's decision-intelligence layer is designed to improve specifically on the categories of case its human reviewers are actually correcting.

MECHANISMS FOR INTELLIGENT AUTOMATION

4.1 Predictive Process Monitoring

Within the architecture, predictive process monitoring serves two distinct functions. Proactively, it forecasts, from a case's partial trace, the probability of a service-level breach or an eventual compliance exception, surfacing the case to a specialist queue before the breach occurs rather than after, addressing the same preventive-versus-detective distinction that motivates enforcement automata in workflow-governance frameworks more broadly [8]. Diagnostically, its confidence output feeds directly into the adaptive routing engine's decision of whether a given transition may be committed automatically, providing the case-level signal that, combined with the content-based signal of Section 4.2, determines the routing decision of Section 4.3.

4.2 Cognitive Task Automation via Retrieval-Augmented LLM Inference

Judgement-dependent tasks whose content is primarily unstructured text — interpreting a claim narrative, drafting a coverage determination, extracting structured fields from a submitted document — are automated through an LLM inference pipeline that retrieves relevant passages from the organisation's policy corpus, prior similar case dispositions, and the specific case's own documentation, and conditions its classification or draft generation on that retrieved context rather than on the model's unconditioned parametric knowledge [13]. This retrieval-augmented design is intended to keep the model's output traceable to a specific, citable internal source, which both improves accuracy relative to unconditioned generation and gives the review console a concrete basis — the specific retrieved passage the model relied on — for a human reviewer to evaluate the draft's justification rather than treating the model's output as an unexplained black box.

4.3 Adaptive Routing and Exception Handling

The adaptive routing engine consumes the combined confidence value from Section 3.3 and applies a per-task-type threshold, itself a versioned, auditable configuration parameter rather than an implicit property of the deployed model, to determine whether a transition is committed automatically, routed to a human reviewer for confirmation, or escalated to a specialist. Exception handling for cases the automation genuinely cannot resolve — malformed input, a case type outside any trained model's coverage, an explicit customer request for human handling — is routed through the same review console used for confidence-based escalation, so that an operator

DOI: 10.46121/pspc.54.3.14

sees a single, unified queue of cases requiring attention rather than a fragmented set of separate exception mechanisms inherited from RPA-era tooling.

4.4 Human-in-the-Loop Governance

Consistent with governance mechanisms proposed for automated workflow enforcement more broadly [8], the framework treats the confidence threshold governing automatic commitment as a staged, monitored rollout parameter rather than a value fixed once at deployment. A newly introduced task-automation capability begins with every decision routed to a human reviewer regardless of confidence, in an advisory configuration analogous to detective-mode enforcement, allowing the accuracy of the decision-intelligence layer's confidence calibration to be characterised against reviewer agreement before any transition is committed without review. Once calibration accuracy at a candidate threshold is validated over a minimum observation window, the threshold is promoted to govern live automatic commitment for that task type, and the promotion decision, together with the validation evidence supporting it, is recorded for subsequent audit.

IMPLEMENTATION CONSIDERATIONS

5.1 Data and Model Lifecycle Management

Because the decision-intelligence layer's structural and content-assessment models are retrained on production feedback, the framework requires the same data-lineage and model-versioning discipline expected of any production machine-learning system: every prediction logged with the specific model version that produced it, every retraining run evaluated against a held-out validation set before promotion, and every promotion recorded in a manner that allows a later audit to reconstruct exactly which model version was active for any historical case. This lifecycle discipline is what allows the continuous learning loop of Section 3.4 to improve the framework's accuracy over time without introducing an untracked, silently shifting decision boundary that no one could explain if a regulator or an internal auditor later asked why a specific case was routed the way it was.

5.2 Integration with Legacy BPM and RPA Stacks

Enterprises adopting this framework rarely do so as a green-field replacement of existing BPM and RPA investment; the architecture is instead designed to attach to an existing process engine's native transition-request interface, in the same engine-agnostic adapter pattern used for workflow-governance enforcement layers [8], so that the decision-intelligence and cognitive-automation layers augment rather than replace the underlying process execution. Existing RPA bots continue to handle the mechanical, low-variability steps they already automate well, while the framework's cognitive layer is introduced specifically at the judgement-dependent steps that RPA's structural limitations, discussed in Section 2.2, have historically left to human operators.

5.3 Risk, Explainability, and Compliance

Every automatically committed decision retains, in the audit trail, the specific retrieved passages and confidence score that supported it, giving a reviewer or auditor a concrete basis for evaluating an individual case after the fact, rather than an unexplained model output. Task types operating in regulated domains — claims adjudication, credit decisioning, medical-necessity determination — are subject to a higher confidence threshold and a mandatory periodic sampling review regardless of the automation layer's measured accuracy, reflecting the position that a well-calibrated model operating at high volume can still produce a materially harmful error pattern that only a human-reviewed sample, rather than aggregate accuracy statistics alone, will reliably surface.

5.4 Change Management and Workforce Transition

Deploying the framework changes the nature of the work performed by the human operators who remain in the loop, shifting a substantial share of their time from originating a decision from scratch to reviewing, confirming, or correcting a model-drafted recommendation. This shift carries genuine organisational risk if handled poorly: reviewers asked to rubber-stamp a high volume of automated drafts without adequate time or incentive to scrutinise them may develop the same complacency documented in other high-automation review settings, quietly eroding the human-in-the-loop safeguard the framework depends on. Mitigating this risk is treated as an explicit part of the deployment, not an afterthought: review workload is monitored and capped per reviewer, a sample of automatically committed decisions (not only escalated ones) is routed for retrospective quality review regardless of confidence score, and reviewer compensation and evaluation criteria are adjusted to reward accurate correction of the model rather than throughput of confirmations alone.

DOI: 10.46121/pspc.54.3.14

Power System Protection and Control
ISSN:1674-3415

CASE STUDY: INTELLIGENT AUTOMATION IN INSURANCE CLAIMS PROCESSING

6.1 Task Specification

We evaluate the framework on a first-notice-of-loss and claims-adjudication process at a mid-size property and casualty insurer, spanning claim intake, coverage verification, damage assessment review, and settlement determination, historically processed by a claims operations team of sixty-one adjusters handling approximately 8,200 claims per month. The process had previously been augmented with RPA for structured data entry and system-to-system transcription but retained a fully manual review for coverage determination and settlement drafting, the two steps most dependent on interpreting the claim narrative and supporting documentation.

6.2 System Deployed

The full AI-augmented framework was deployed over a nine-month period, following the staged rollout discipline of Section 4.4: an initial three-month advisory period in which the decision-intelligence layer's predictions and the cognitive layer's drafted coverage determinations were logged and shown to adjusters for reference but not automatically committed, followed by a phased promotion of the confidence threshold governing automatic commitment for the coverage-verification and settlement-drafting steps specifically.

6.3 Process Performance Across Automation Stages

Figure 1 compares five process-performance metrics across three configurations: traditional BPM (the insurer's process prior to any RPA or AI augmentation), RPA-augmented BPM (structured data entry and transcription automated, judgement steps manual), and the proposed AI-augmented configuration following the full nine-month deployment. Straight-through processing — the proportion of claims resolved without any manual touch beyond a routine confirmation — rose from 22% under traditional BPM to 51% once RPA automated the structured steps, and to 79% under the full AI-augmented configuration once the cognitive-automation layer addressed coverage verification and settlement drafting. Median cycle time fell correspondingly from 46 hours to 6.5 hours, and the exception rate — claims requiring routing outside the standard path — fell from 31% to 9%, reflecting both the decision-intelligence layer's earlier identification of claims likely to require special handling and the cognitive layer's ability to resolve narrative ambiguity that would previously have generated an automatic RPA exception. Cost per case fell by roughly two-thirds relative to the traditional baseline, and customer satisfaction, measured through the insurer's existing post-settlement survey instrument, rose from 68 to 88 on the insurer's standard hundred-point CSAT scale, consistent with the combined effect of faster resolution and a lower rate of claims requiring a frustrating second round of manual review.

Figure 1: Process performance across automation stages
(Insurance claims processing, 9-month evaluation)

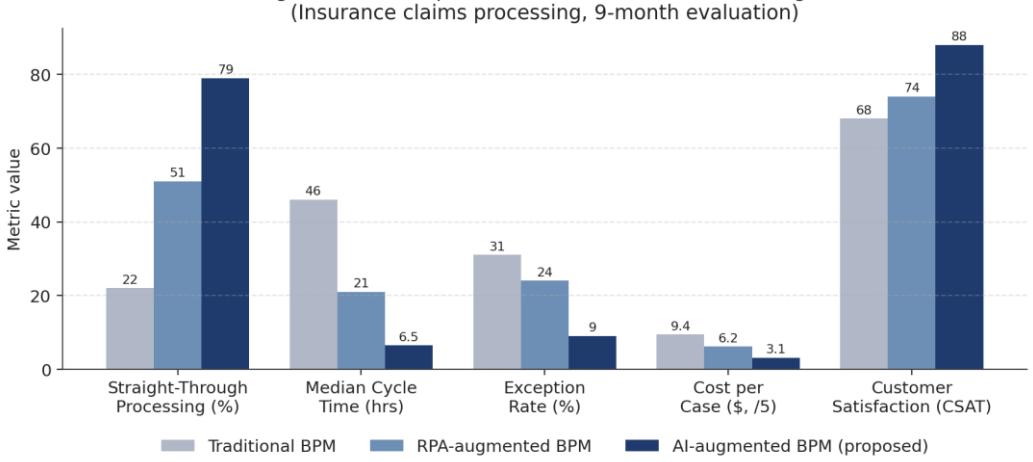

Figure 1: Process performance compared across traditional BPM, RPA-augmented BPM, and the proposed AI-augmented configuration, evaluated over a nine-month deployment in insurance claims processing.

6.4 Learning Dynamics and Confidence-Threshold Sensitivity

Figure 2(a) tracks the straight-through-processing rate over the nine months since initial deployment, showing a learning curve that rises from 34% in the first month to a plateau of 79% by month nine, with the steepest gains occurring during months two and three, coinciding with the first scheduled retraining cycle of the continuous process-learning loop on the corrections accumulated during the initial advisory period; this pattern is consistent

DOI: 10.46121/pspc.54.3.14

Power System Protection and Control
ISSN:1674-3415

with the framework's premise that a meaningful share of its accuracy improvement comes specifically from learning on production feedback rather than from the initial model's out-of-the-box capability alone. Figure 2(b) reports a sensitivity analysis over the confidence threshold governing automatic commitment, ranging from a permissive 0.5 to a conservative 0.95. Lowering the threshold from 0.95 to 0.5 reduced the human-escalation rate from 55% to 6% of claims but raised the residual error rate — the proportion of automatically committed decisions subsequently found, on downstream audit sampling, to have been incorrect — from 1.1% to 14.8%, quantifying the same automation-coverage-versus-quality trade-off observed in comparable graduated-enforcement mechanisms [8] and giving the insurer's claims operations leadership an empirical, auditable basis for selecting a threshold appropriate to the regulatory risk tolerance of each task type rather than adopting a single default value across all claim categories.

Figure 2: Learning dynamics and confidence-threshold sensitivity

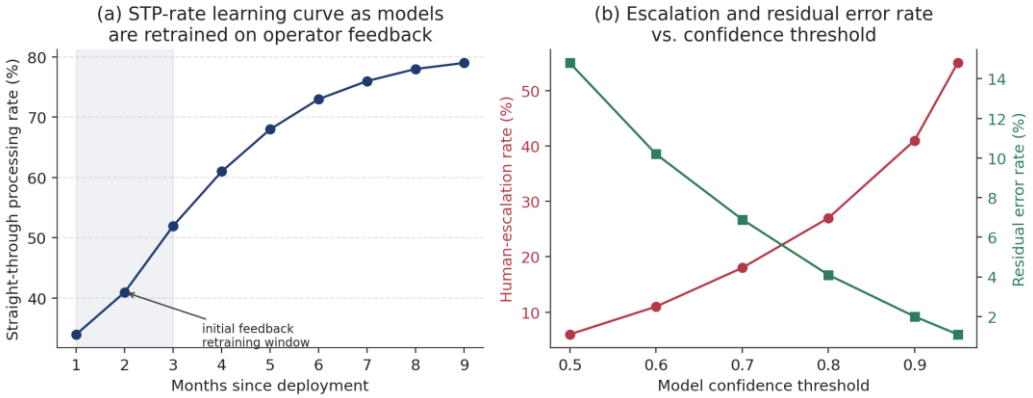

Figure 2: (a) Straight-through-processing learning curve over the nine months following deployment, showing the effect of the first scheduled feedback-retraining cycle. (b) Human-escalation rate and residual error rate as a function of the automatic-commitment confidence threshold.

DISCUSSION

7.1 Implications for the Practice of BPM

The case study's central finding — that straight-through processing continued rising substantially after initial deployment, driven by the continuous learning loop rather than by the initial model alone — suggests that the appropriate way to evaluate an AI-augmented BPM deployment is longitudinally rather than at a single post-launch snapshot, since a measurement taken during the initial advisory period would have understated the framework's eventual performance by a wide margin. This has a direct implication for how organisations should structure procurement and internal success criteria for intelligent-automation programs: a fixed accuracy bar evaluated at go-live is a poor proxy for a framework whose central value proposition is that it improves specifically because it is deployed into live operation rather than despite being deployed.

A second implication concerns how the framework changes the unit of BPM improvement itself. Traditional BPM practice treats process improvement as a discrete project: a process is analysed, redesigned, and re-deployed, and the improved version then runs largely unchanged until the next redesign cycle, often years later. The continuous learning loop of Section 3.4 collapses this cycle into an ongoing, incremental one, in which the process's effective behaviour shifts continuously as the underlying models retrain. This is a genuine benefit — the process improves without waiting for the next formal redesign project — but it also means that the conventional BPM practice of periodically re-certifying a process against its documented design needs to be adapted to a system whose behaviour is, by design, not fixed between certification cycles, an adaptation this paper's governance console and versioned model-promotion record are intended to support.

A third implication touches procurement and vendor strategy: because the framework is designed to attach to existing BPM and RPA investment rather than replace it, organisations evaluating vendors for the decision-intelligence and cognitive-automation layers should weight interoperability with their existing process-engine adapters and audit-ledger export formats at least as heavily as raw model accuracy benchmarks, since a highly

DOI: 10.46121/pspc.54.3.14

accurate model that cannot be integrated through the thin-adapter pattern of Section 5.2 will incur substantially higher deployment cost than a marginally less accurate one that can.

7.2 Operational and Organisational Considerations

The RPA-augmented configuration in Figure 1 already captured a substantial share of the cycle-time improvement relative to traditional BPM, confirming that RPA's structural automation of low-variability steps remains valuable in its own right and should not be discarded in favour of the AI-augmented layer; the framework is additive to, rather than a replacement for, existing RPA investment, consistent with the integration approach described in Section 5.2. Organisationally, the transition from a fully manual coverage-determination step to a reviewer-confirmation model required retraining adjusters toward a review-and-correct workflow rather than an originate-from-scratch workflow, a change management consideration that affected adoption speed independently of the underlying model's technical accuracy and that a purely technical evaluation of the framework would not capture.

7.3 Limitations

The present study draws on a single insurer and a single claims process over a nine-month window, and the specific magnitudes reported — particularly the confidence-threshold sensitivity of Section 6.4 — should be read as illustrative of the method rather than as universal parameters transferable to other claim types, regulatory jurisdictions, or industries with different tolerance for automated decisioning. The decision-intelligence layer's reliance on retrieval-augmented LLM inference means its accuracy is bounded by the completeness and currency of the underlying policy corpus; a corpus that has drifted out of date relative to an actual policy change would cause the cognitive layer to condition its output on superseded guidance, a failure mode that the framework's confidence scoring does not, by itself, detect, since the model may express high confidence in an answer that is well-supported by an out-of-date source. The continuous learning loop's reliance on human reviewer corrections as training signal also means the framework inherits whatever systematic bias exists in historical reviewer decisions; the framework does not, on its own, correct a bias that its human reviewers have consistently exhibited, and organisations deploying it should audit training-data quality for this risk independently of the framework's own accuracy metrics.

A further limitation concerns generalisability across languages and jurisdictions: the cognitive task-automation layer’s accuracy on unstructured content depends on the underlying LLM’s competence in the specific language and regulatory vocabulary of the deployment context, and organisations operating across multiple jurisdictions with materially different insurance or financial regulation should expect to validate the framework’s confidence calibration separately for each jurisdiction rather than assuming a single validated threshold transfers unchanged.

7.4 Toward Increasingly Autonomous, Accountably Governed Process Management

A natural extension is to allow the confidence threshold itself to vary adaptively per case attribute, in the manner proposed for risk-tiered escalation deadlines in workflow-governance frameworks [8], so that lower-stakes, well-understood claim categories operate at a more permissive threshold than novel or high-value categories without requiring a separately configured task type for every risk tier. A further direction is extending the decision-intelligence layer's uncertainty quantification beyond a single scalar confidence value to a richer representation of which specific aspect of a case the model is uncertain about — the coverage determination itself, or merely the specific settlement amount — allowing the adaptive routing engine to route only the uncertain sub-decision to a human reviewer rather than the entire case, a refinement that could further reduce human-escalation volume without a corresponding increase in residual error.

CONCLUSION

This paper argued that the synergy of predictive process intelligence and language-capable cognitive automation, composed within a single governed BPM architecture, enables a redefinition of business process management from a static specification of deterministic routing into a living system that predicts outcomes, drafts and evaluates judgement-dependent decisions, learns continuously from operator feedback, and escalates to a human precisely when its own calibrated confidence indicates it should. Evaluated over a nine-month deployment in insurance claims processing, the framework raised straight-through processing from 22% under traditional BPM to 79% under the full AI-augmented configuration, reduced median cycle time from 46 to 6.5 hours, and demonstrated that a substantial share of its eventual accuracy was attributable to the continuous learning loop rather than to the initial model alone. The framework's principal limitations are its dependence on a current, well-maintained policy corpus for retrieval-augmented inference and its inheritance of whatever bias exists in the historical reviewer

DOI: 10.46121/pspc.54.3.14

decisions it learns from, both of which require deliberate organisational attention rather than being resolved by the architecture automatically. More broadly, the results suggest that the next stage of business process management is not a further increment of rule-based automation, but the disciplined integration of predictive and generative AI capability directly into the governed process runtime, with human oversight preserved not as a residual manual step but as an explicit, auditable, and continuously informed part of the system's own design.

REFERENCES

1. **Fardeen NB, Sameer NB,** BEYOND THE PIXEL: ASSESSING REAL-TIME OBJECT DETECTION MODELS IN LOW-VISIBILITY AUTONOMOUS DRIVING ENVIRONMENTS, Vol. 53 No. 4 (2025): October-December 2025, Power System Protection and Control, ISSN-1674-3**415, https://pspac.info/index.php/dlbh/article/view/371**
2. **Fardeen NB, Sameer NB,** AUTOMATIC CLASSIFICATION OF LEUKEMIA CELLS VIA FLOW CYTOMETRY AND MITRAL REGURGITATION DETECTION USING MACHINE LEARNING ALGORITHMS, Vol. 49 No. 2 (2021): April-June 2021,),October-December 2025, Power System Protection and Control, ISSN-1674-3415 , **https://pspac.info/index.php/dlbh/article/view/361**
3. **Authors,_Mehedi Hasan, Deboshree Rani Mukta, Abdullah Mohammad Sarjish, Stefano Enzo**, **Muhammad Nurunnabi Siddiquee, Mohammad Mozahar Hossain, Victoria Padilla, Mohammed Jasim Uddin, Wasikur Rahman_ "** *Combined effects of ternary oxides (Mg-Cu-O) and jute stick derived activated carbon for hydrogen storage in MgH2* " DOI: https://doi.org/10.1007/s11356-025-36394-4 , **Apr-2025**
4. **Authors: _Md. Wasikur Rahman Md. Sabbir Ahmed, Selim Reja, Sk Md Ali Zaker Shawon, Farhana Nazmin, Abdullah Mohammad Sarjish, Yousuf Ali, M. Azizur R. Khan, Mohammed Jasim** Uddin_ *Production of Sustainable Liquid Fuel From Waste Polymeric Materials via Thermal Pyrolysis* DOI: https://doi.org/10.1155/er/1449087 ,
5. **Sarjish, Abdullah Mohammad, et al.** " *Nano Polymers for Clean Energy Solutions* ." Nanostructured Polymers: Synthesis and Performance. Singapore: Springer Nature Singapore, 2026. 461-495. https://doi.org/10.1007/978-981-95-3581-1_18 , **Feb-2026**
6. **Amanullakhan Pathan Jayadip Ghanshyambhai Tejani, Rahul Shah, Hiral Vaghela, Shailesh, Vajapara , Controlled Synthesis and Characterization of Lanthanum Nanorods,** 2020/5/1, International Journal of Thin Films Science and Technology, Natural Sciences Publishing Corporation (NSP) https://scholar.google.com/citations?view_op=view_citation&hl=en&user=efbNMkwAAAAJ&citation_for_view=efbNMkwAAAAJ:M3NEmzRMIkIC
7. **Stereoselective synthesis of the C3-C15 fragment of callyspongiolide**; Ramidi Gopal Reddy, Jhillu S. Yadav and Debendra K. Mohapatra. (Tetrahedron Letters. 2018, 59, 3579-3582). https://doi.org/10.1016/j.tetlet.2018.08.041, https://www.sciencedirect.com/science/article/pii/S0040403918310426
8. **Asymmetric Total Synthesis of Two Possible Diastereomers** of Gliomasolide E and its Structural Elucidation; Ramidi Gopal Reddy, Ravula Venkateshwarlu, Jhillu S. Yadav and Debendra K. Mohapatra. (J. Org. Chem. 2017, 82(2), 1053-1063). https://doi.org/10.1021/acs.joc.6b02611, https://pubs.acs.org/doi/abs/10.1021/acs.joc.6b02611
9. **Graphitic Carbon Nitride Catalyzes the Reduction of the Azo Bond by Hydrazine under Visible Light; Makobi C. Okolie**, Glory G. Ollordaa, Gopal R. Ramidi, Xin Yan, Yufeng Quan, Qingsheng Wang and Yingchun Li. (Nanomaterials 2024, 14(17), 1402), https://doi.org/10.3390/nano14171402, https://www.mdpi.com/2079-4991/14/17/1402

DOI: 10.46121/pspc.54.3.14

10. **Lakshmi Rahul Reddy Mareddy, 2026, IJCNIS,**
11. **Lakshmi Rahul Reddy Mareddy, 2026, ICAA, ,https://eudoxuspress.com/index.php/pub/article/view/4780/3569**
12. **Lakshmi Rahul Reddy Mareddy, 2026, IJCMI , https://www.ijcmi.in/index.php/ijcmi/article/view/70**
13. **Rangavinay Teja Royal Jagga** , 2026, https://ieeexplore.ieee.org/document/11448495
14. **Rangavinay Teja Royal Jagga , 2026, https://ieeexplore.ieee.org/document/11448514**
15. **Rangavinay Teja Royal Jagga , 2026, https://ieeexplore.ieee.org/document/11448337**
16. **Rangavinay Teja Royal Jagga , 2026, https://ieeexplore.ieee.org/document/11497560**
17. **Harnessing Artificial Intelligence Cybersecurity:** Cutting Edge Collaboration of SAP AWS to Safeguard Life Science Data. 2026 https://doi.org/10.1109/aiei69164.2026.11497833
18. **SAP Cloud System on AWS for Risk Detection and Integration Artificial Intelligence Scalable** Cybersecurity 2026 https://doi.org/10.1109/aiei69164.2026.11497435
19. **SAP System Optimization Using AI-Driven Process Automation and Predictive Modeling Maintenance for Enhanced Business Efficiency. 2025** https://doi.org/10.1109/computingcon64838.2025.11376613
20. https://ieeexplore.ieee.org/abstract/document/11376613
21. https://ieeexplore.ieee.org/abstract/document/11497833
22. https://ieeexplore.ieee.org/abstract/document/11497435
23. **Kaleshwar Aryasomayajula**, PREVENTING BIAS IN AI-BASED DECISION-MAKING: ANALYZING TECHNIQUES TO REMOVE UNFAIR PREJUDICE IN ALGORITHMIC OUTCOMES, Vol. 50 No. 2 (2022): April-June 2022, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/294
24. **Kaleshwar Aryasomayajula, MODERN** DEVELOPMENT PRACTICES: INVESTIGATIONS INTO SERVERLESS COMPUTING, LOW-CODE/NO-CODE PLATFORMS, AND DEVELOPER PRODUCTIVITY IN REMOTE ENVIRONMENTS, Vol. 50 No. 4 (2022): October-December 2022, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/364
25. **Kaleshwar Aryasomayajula**, Micro services: "A Comparative Analysis of Monolithic vs. Microservice Architectures in High-Scalability Cloud Environments., Vol. 51 No. 2 (2023): April-June 2023 , Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/374
26. **Alam, M. Z., Rahman, R., Sozib, H. M., Ahmed**, H., Hossain, A., Sabeena, A. A., Tasnim, A. F., Ahmed, F., Sarkar, M. I., & Erdei, T. I. (2026). Enhancing Thyroid Disease Diagnosis with Machine Learning and Counterfactual Explainable AI. IEEE Access, 1–1. https://doi.org/10.1109/access.2026.3663497
27. **Yeasmin, S., Semi, M. M. A., Rony, M. K. K., Das, S., Sabeena, A. A., Rahman, R., Biswas, B., Ahmed, F., & Hossain, A. (2025).** Artificial Intelligence for Mental Health Monitoring: A Solution for Digital Behavioral Health Care and Education—An Umbrella Review. Health Science Reports, 9(1), e71703. https://doi.org/10.1002/hsr2.71703
28. **Hasan, S., Rahman, K. A., Ahmed, F., & Hossain, A. (2026). An integrated AI-driven framework for** maternal resource intelligence shortages across U.S. hospitals. Integrative Biomedical Research, 10(1), 1–8. https://doi.org/10.25163/biomedical.10110694
29. **Riipa, M. B., Ahmed, F., Rony, M. K. K., Hossain, A., Islam, A., Utsho, M. R., Kamal, M. B., Sharmin, S., &** Tasnim, A. F. (2026). The role of artificial intelligence in predicting

DOI: 10.46121/pspc.54.3.14

cardiovascular outcomes: a systematic review and meta-analysis. Biostatistics & Epidemiology, 10(1). https://doi.org/10.1080/24709360.2026.2670804

30. **Semi, M. M. A., Das, S., Utsho, M. R., Hossain,** A., Kamal, M. B., Sizan, A. A., Tasnim, A. F., Yeasmin, S., & Parvin, M. R. (2026). Artificial Intelligence in Public Health Education: A Scoping Review of workforce competency development. Health Science Reports, 9(3). https://doi.org/10.1002/hsr2.72066
31. **Ahmed, F., Hasan, S., Hossain, A., & Rahman, K. A. (2026). Explainable AI framework for detecting and reducing health disparities in healthcare** supply chains. Journal of AI, ML and DL, 2(1), 1–8. https://doi.org/10.25163/ai.2110685
32. **Godavari Modalavalasa," SCALABLE CLOUD SOLUTIONS THROUGH ARTIFICIAL INTELLIGENCE GOVERNANCE**: APPLICATIONS IN HEALTHCARE AND FINANCIAL SYSTEMS", Vol. 54 No. 1 (2026): January-March 2026, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/219, , DOI: https://doi.org/10.46121/pspc.54.1.25
33. **Godavari Modalavala**sa, AUTONOMOUS DATA ENGINEERING PIPELINES: A POLICY-DRIVEN ARCHITECTURE FOR SECURE AND SCALABLE CLOUD-NATIVE ANALYTICS, Vol. 53 No. 4 (2025): October-December 2025, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/181, DOI: https://doi.org/10.46121/pspc.53.4.25
34. **Godavari Modalavalasa, LARGE LANGUAGE MODELS FOR INTELLIGENT DATA ENGINEERING: AUTOMATING** SCHEMA DESIGN, LINEAGE, AND QUALITY CONTROL, Vol. 50 No. 2 (2022): April-June 2022, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/183 , DOI: https://doi.org/10.46121/pspc.50.2.4
35. **Godavari Modalavalasa**, FEDERATED LEARNING FOR ENTERPRISE CLOUD DATA ENGINEERING: ARCHITECTURE, SECURITY, AND GOVERNANCE CHALLENGES, Vol. 51 No. 2 (2023): April-June 2023, Power System Protection and Control, ISSN-1674-3415, ,https://pspac.info/index.php/dlbh/article/view/184, DOI: https://doi.org/10.46121/pspc.51.2.5
36. **Godavari Modalavalasa, AI**-DRIVEN DATA GOVERNANCE: INTELLIGENT METADATA, LINEAGE, AND COMPLIANCE AUTOMATION IN CLOUD DATA PLATFORMS, Archives / Vol. 52 No. 1 (2024): January-March 2024 /, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/182, DOI: https://doi.org/10.46121/pspc.52.1.3
37. **Hima Bindu Lekkala, VishnuVardhan Bandar**i , AUTONOMOUS WORKFLOW OPTIMIZATION USING MULTI AGENT AI SYSTEMS AI AGENTS MANAGE STATIONS, WIP, AND TASK HANDOFFS, Vol. 54 No. 2 (202): April-June 2026, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/306 , DOI: https://doi.org/10.46121/pspc.54.2.08
38. **Lekkala, H. B., & Bandari, V. V. (2026). Autonomous workflow optimization using multi-agent AI systems:** AI agents manage stations, WIP, and task handoffs. *Power System Protection and Control*, 54(2), 95–101. http://pspac.info/index.php/dlbh/article/view/306

39. **Bandari, V. V., & Lekkala, H. B. (2024). Physics-informed reinforcement learning for real-time control of complex manufacturing processes**. *Power System Protection and Control*, 52(2), 164–170. https://pspac.info/index.php/dlbh/article/view/343
40. **Lekkala, H. B., & Bandari, V. V. (2025). AI-based energy-aware scheduling and process optimization** in engineer-to-order smart manufacturing systems. *Power System Protection and Control*, 53(1), 30–36. http://pspac.info/index.php/dlbh/article/view/341
41. **Bandari, V. V., & Lekkala, H. B. (2026). AI-driven digital** thread framework for end-to-end lifecycle optimization in ETO manufacturing systems. *Power System Protection and Control*, 54(2), 318–325. http://pspac.info/index.php/dlbh/article/view/340
42. **Lekkala, H. B., & Bandari, V. V. (2026). Deep learning** for weld defect detection using CNN or vision transformers fusion detection. *Power System Protection and Control*, 54(2), 151–158. https://pspac.info/index.php/dlbh/article/view/314
43. **Bandari, V. V., & Lekkala, H. B. (2024). AI-based operator** behavior monitoring and cost optimization using digital traceability in manufacturing systems. *Power System Protection and Control*, 52(1), 38–45. https://pspac.info/index.php/dlbh/article/view/342
44. **Shreyas Jayaprakash,** MACHINE LEARNING-DRIVEN VERIFICATION: USING ML TO OPTIMIZE COVERAGE ANALYSIS, PREDICT SIMULATION RUNTIMES, AND AUTO-GENERATE TESTBENCHES, Vol. 54 No. 2 (2026): April-June 2026, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/391
45. **Mohammed Shafi Kundiladi,** EVENT-DRIVEN IMAGE AND VEHICLE STATUS MANAGEMENT FOR LOW-POWER IOT DIGITAL LICENSE PLATES, Vol. 53 No. 3 (2025): July-September 2025, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/175
DOI: https://doi.org/10.46121/pspc.53.3.17
46. Controlled Synthesis and Characterization of Lanthanum Nanorods, **Author(s), Jayadeep Tejani, Rahul Shah, Hiral Vaghela, Shailesh Vajapara, Amanullakhan Pathan**, doi:10.18576/ijtfst/090205,
URL: https://www.naturalspublishing.com/Article.asp?ArtcID=20705, May-2020
47. **Kunal Kapoor**, **Apr 2026,** Relation Extraction and NER through Fine Tuned BERT model with transfer Learning, https://ieeexplore.ieee.org/document/11472214/metrics#metrics
48. **Harender Bisht**, Human–AI Collaboration in Insurance Fraud Detection: Ethical Cloud-Native Architectures for Fair and Transparent Decision Support, **Volume 2026, Issue 1,** https://doi.org/10.52710/cfs.873
49. **Harender Bisht,** Ethical AI-Augmented Compliance Systems: Architectural Innovations for Cloud-Native Financial and Insurance Data Integrity, **Vol. 11 No. 1s (2026) ,** https://doi.org/10.52783/jisem.v11i1s.14056
50. **Harender Bisht**, Governance-By-Design For AI-Based Insurance Fraud Detection: Auditability, Accountability, And Regulatory Traceability, **Archives / Vol. 9 No. 1 (2026**) ,https://doi.org/10.63278/jicrcr.vi.3620
51. **Harender Bisht,** Operationalizing Explainable AI in Insurance Fraud Detection: From Model Transparency to Decision Justification, **Vol. 35 No. 1 (2026): JOCAAA**, https://eudoxuspress.com/index.php/pub/article/view/4791,
52. **A. MahendraVardhan and S. Sridhar,** "Determining False Positive Analysis of Software Vulnerabilities with Predefined Scan Rules using Random Forest Classifier and Decision Tree

Technique," **2022 4th International Conference on Advances in Computing**, Communication Control and Networking (ICAC3N), Greater Noida, India, 2022, pp. 622-625, **DOI: https://doi.org/10.1109/ICAC3N56670.2022.10074458**

53. **Amilineni, M.V. 2025**. Audit-Safe Agentic RAG Framework for Regulated Enterprise Systems: A Trustworthy AI Architecture for SAP, Finance, Healthcare, and Government Workflows. INTERNATIONAL JOURNAL OF ADVANCES IN SIGNAL AND IMAGE SCIENCES. (Jan. 2025), 34–46. **DOI: https://doi.org/10.29284/7fjzk883**
54. **Islam, Imtiaz (Annotator), Goard, Michael (Annotator), Daves, Daniel (Annotator),** Credit Card Fraud Detection: A Comparative Study of Anomaly Detection and Supervised Learning under Extreme Class Imbalance, Published **May 5, 2026** | Version v1, **DOI: 10.5281/zenodo.20032207**
55. **Deepa Nagalavi; Shankar Balla Microsoft, USA; Pushpalatha. M; Nashwan Adnan Othman; Malik Bader Alazzam; A. Balakumar, Enhancing Real-Time Language Processing via Advanced PET Signal Analysis and Deep Learning,** ***06-07 June 2025***, DOI: 10.1109/ICICKE65317.2025.11136561 , https://ieeexplore.ieee.org/document/11136561
56. **Venkata Sai Aditya Kondru,FMR, Carmel, USA ,; Shankar Balla; Dhavalkumar Thakar; Dheeraj Velaga; Sri Lekha Bandla,** Intelligent Human–Computer Interaction for Navigation Control Through Vision-Based Hand Gesture Recognition, Conferences >2026 IEEE 15th International , **Date of Conference:** ***07-09 April 2026***,**Date Added to IEEE *Xplore*:** 08 May 2026
DOI: 10.1109/CSNT69054.2026.11502317 , https://ieeexplore.ieee.org/abstract/document/11502317
57. **SUDIPKUMAR GHANVAT , AISHWARYA BADLANI, SHREYA SANDEEP JOSHI,** MARKETING EFFECTIVENESS IN THE POST-COOKIE ERA: A COMPARATIVE ANALYSIS OF FIRST PARTY AND ZERO-PARTY DATA STRATEGIES ON CAMPAIGN PERFORMANCE AND CUSTOMER EQUITY, **VOL. 31 NO. 4 (2023):** JOCAAA , HTTPS://WWW.EUDOXUSPRESS.COM/INDEX.PHP/PUB/ARTICLE/VIEW/3940
58. **Navin Chhibber,; Deepak Singh; Anokh Kishore, Capital One, USA; Nikita Chawla; K. Anguraj,** Multi-Granular Attention-Driven Reinforcement Learning Framework for Web Intelligent Enhancement Systems, , **11 June 2026**, https://ieeexplore.ieee.org/document/11483386
59. **Kapoor K, et al. AI-Driven COVID-19 Care**: Analytics, Scheduling & Geo-Access. International Journal of Drug Delivery Technology (IJDDT). **2026;16(65S):Article** 136. Available at: https://impactfactor.org/PDF/IJDDT/16/IJDDT,Vol16,Issue65s,Article136.pdf
60. **Kilama, D. (2026). Mechanisms of Chinese** Aid to Africa: A Comparative Analysis with Western Philanthropic and Development Models. Interdisciplinary Social Research, 2(1).DOI: https://doi.org/10.64220/isr.v2i1.011, https://scholarlysummit.com/journals/isr/articles/mechanisms-of-chinese-aid-to-africa-a-comparative-analysis
61. **Kilama, D. (2026). Beyond Altruism**: Analyzing China's Strategic Motivations Behind Development Aid to Africa. *Interdisciplinary Social Research, Volume 2 (2026)*(Issue 1), DOI: . https://doi.org/10.64220/isr.v2i1.012, Beyond Altruism: Analyzing China's Strategic Motivations Behind Development Aid to Africa

62. Anumolu DPK, Veena B, Afreen SS, Bandla J, Lakshmi KC, Pulusu VS. In vitro models for studying drug metabolites and metabolizing enzymes: A comprehensive review. Int J Pharm Investig. 2026;16(3):885–91. Available from: http://dx.doi.org/10.5530/ijpi.20260116
63. Asgar Z, Kumar G, Khandelwal P, Pratap B, Gurjar V, Baluni A, et al. Utility of contrast-enhanced computed tomography abdomen in Intestinal Obstruction. Int J Drug Deliv Technol. 2026;16(32s). Available from: http://dx.doi.org/10.25258/ijddt.16.32s.4
64. Anumolu DP, Ramatulasi J, Ashok G, Afreen SS, Mathew C, Shakar PV. Synthesis and characterization of MIP ghost approach for selective extraction of empagliflozin and quantification by liquid chromatography. J Chromatogr Sci. 2025;63(2). Available from: http://dx.doi.org/10.1093/chromsci/bmaf008.
65. Anumolu DPK, Naraparaju S, Addada SB, Pagidipally V, Pulusu VS, Afreen SS. Synthesis of silver nanoparticles using Ecbolium ligustrinum leaves: Characterization through advanced analytical techniques. Nanotechnol Percept . 2024;710–8. Available from: http://dx.doi.org/10.62441/nano-ntp.vi.3618
66. Naraparaju S, Mukti A, Anumolu DP, Chaganti S. Development and validation of a spectrofluorimetric method for the quantification of capecitabine in bulk and tablets. Turk J Pharm Sci [Internet]. 2023;20(4):234–9. Available from: http://dx.doi.org/10.4274/tjps.galenos.2022.46364
67. Naraparaju S, Yamjala P, Chaganti S, Anumolu DP. Spectrophotometric quantification of atomoxetine hydrochloride based on nucleophilic substitution reaction with 1,2-naphthoquinone-4-sulfonic acid sodium salt (NQS). Turk J Pharm Sci [Internet]. 2024;20(6):405–11. Available from: http://dx.doi.org/10.4274/tjps.galenos.2022.09147.
68. Anumolu PD, Shakar PV, Tulasi JR, Soujanya C, Afreen SS, Ashok G. Concurrent discriminative emission intensity quantification of Fexofenadine hydrochloride and Montelukast Sodium. Orient J Chem [Internet]. 2022;38(6):1364–8. Available from: http://dx.doi.org/10.13005/ojc/380605
69. Anumolu PD, Gurrala S, Menkana S, Mathew C, Srimantula L. Application of fluorotag in quantification of metformin in tablet dosage form and biological samples. The Thai Journal of Pharmaceutical Sciences [Internet]. 2022;46(1):93–8. Available from: http://dx.doi.org/10.56808/3027-7922.2550
70. Gurrala S, Raj S, Cvs S, Anumolu PD. Quality-by-design approach for chromatographic analysis of metformin, empagliflozin and linagliptin. J Chromatogr Sci [Internet]. 2022;60(1):68–80. Available from: http://dx.doi.org/10.1093/chromsci/bmab030
71. Gurrala S, Raj S, Cvs S, Anumolu DP, Naraparaju S, Nizampet H. Response surface methodology in spectrophotometric estimation of saxagliptin, derivatization with MBTH and ninhydrin. Turk J Pharm Sci [Internet]. 2022;19(1):9–18. Available from: http://dx.doi.org/10.4274/tjps.galenos.2021.93195