

The Impact of Artificial Intelligence, Machine Learning, and Deep Learning on Business Processes: A Socio-Technical Framework for Intelligent Process Automation, Decision Support, and Governance

Laraib Tabish¹, Falak Ali^{2*}, Maitikuerban Abuduaini³

¹Department of Business Management, Government College University, Faisalabad, Pakistan.

²Department of Computer Science, Government College University, Faisalabad, Pakistan.

³Sakarya University, Sakarya, Turkey.

*Corresponding author: falakali20001@gmail.com

Abstract:

Artificial intelligence (AI), machine learning (ML), and deep learning (DL) are moving from experimental projects into the operational core of organizations, where they increasingly execute, inform, and redesign business processes. Their effect is frequently described in terms of automation and efficiency, yet the actual consequences for a business process depend on where the technology is placed, how its outputs are routed into decisions, and whether the surrounding data, security, and governance arrangements can support it. This paper examines how modern AI, ML, and DL tools affect business processes across the full process lifecycle of design, execution, decision-making, monitoring, and improvement. Drawing on applied studies of ML-based face recognition, image-based electronic toll collection, peer-to-peer energy trading simulation, robotic process automation (RPA) and AI-bot identity management, low-code/no-code intelligent automation, real-time AI decision processing, human–bot discrimination, honeypot-based threat intelligence, and federated learning, the paper distinguishes three effects: task automation, decision augmentation, and process transformation. It then proposes a six-layer socio-technical framework comprising process data and integration, AI/ML/DL intelligence, automation and execution, security, identity and privacy, governance and human oversight, and business value. A cost-sensitive, three-zone decision-routing model is introduced to show how confidence thresholds should be derived from business costs rather than from model accuracy alone, and a process-level evaluation framework is defined that combines quality, time, cost, risk, and human-oversight measures. The analysis concludes that AI improves business processes most reliably when it is applied to well-understood processes with standardized data, when automated decisions are bounded by explicit thresholds and human review, and when software bots are governed as accountable identities.

Keywords: Artificial intelligence, machine learning, deep learning, business process management, intelligent process automation, robotic process automation, low-code/no-code, decision support, AI governance, human-in-the-loop.

I. INTRODUCTION

A business process is a structured set of activities that transforms inputs into an outcome of value for a customer, whether that customer is external, such as a buyer placing an order, or internal, such as a department requesting a purchase. For several decades organizations automated such processes through enterprise resource planning systems, workflow engines, and hand-coded business rules. These systems were effective for stable, well-specified tasks, but they could not interpret unstructured inputs such as images, documents, or behavior, and they could not adapt when the conditions of a process changed.

Modern AI tools change this situation. ML models learn decision rules from historical data, DL models extract features directly from images, text, and signals, computer-vision systems capture transactional information without manual entry, and software bots execute repetitive actions across applications. As a result, AI no longer affects only the speed of individual tasks; it changes where decisions are made within a process, who or what makes them, how exceptions are handled, and how the process itself is redesigned over time.

Applied studies illustrate this breadth. Ahmed et al. [1] compared K-nearest neighbor (KNN), linear discriminant analysis (LDA), support vector machine (SVM), and principal component analysis (PCA) approaches for face recognition and reported that the combined PCA+LDA+INN approach achieved the highest average accuracy among the evaluated methods, a result relevant to

identity-verification steps in customer onboarding and access processes. Mujeeb et al. [2] proposed an electronic toll collector framework based on image classification, automating a revenue-collection process that traditionally depended on manual handling. Ahmed et al. [3] developed a software simulator for peer-to-peer energy trading that integrates generation, consumption, storage, and trading functions, and Ahmed et al. [4] examined how low-code/no-code platforms combined with AI/ML can enhance intelligent automation. Together these works show that AI can be embedded in very different processes, but also that the benefit depends on data representation, integration, and control.

The central research question of this paper is therefore: How do modern AI, ML, and DL tools affect the design, execution, control, and governance of business processes, and what architectural and organizational safeguards are required for these effects to be beneficial?

The contributions of this paper are: (1) a distinction between three types of AI effect on business processes, namely task automation, decision augmentation, and process transformation; (2) a mapping of AI, ML, and DL techniques onto the stages of the business process lifecycle; (3) a six-layer socio-technical framework for AI-enabled business processes; (4) a cost-sensitive, three-zone decision-routing model for human-in-the-loop automation; (5) an analysis of organizational, security, and privacy risks with corresponding controls; and (6) a process-level evaluation framework that extends beyond model accuracy.

The remainder of the paper is organized as follows. Section II reviews related work, Section III states the problem, Section IV analyses the effect of AI across the process lifecycle, Section V presents the proposed framework, Section VI introduces the decision-routing model, Sections VII to IX discuss data, organizational, and risk considerations, Section X presents application scenarios, Section XI defines the evaluation framework, and Sections XII to XVI provide discussion, design principles, limitations, future directions, and conclusions.

II. BACKGROUND AND RELATED WORK

Generations of Process Automation: Process automation can be viewed as three overlapping generations. The first generation consists of rule-based workflow and transaction systems in which every path is specified in advance. The second generation, robotic process automation (RPA), uses software bots to imitate human interaction with existing applications, automating repetitive tasks without changing the underlying systems. The third generation, often called intelligent process automation, combines execution bots with ML and DL models that classify inputs, predict outcomes, and recommend or take decisions. Large language models are now extending this generation further, with emerging applications in cybersecurity and engineering automation [5]. Each generation extends the range of automatable work, but each also moves more judgement from people into software.

Machine Learning for Classification and Verification: Many process steps are classification decisions: whether a document is an invoice or a contract, whether a claim is routine or suspicious, or whether a presented face matches an enrolled identity. A supervised model can be written as $\hat{y} = f(X; \theta)$, where X represents input features, θ the learned parameters, and $\hat{y}$ the predicted class. The face-recognition comparison in [1] shows that the choice of feature representation and classifier materially affects accuracy, which in a business process translates directly into the rate of correctly verified customers and the rate of erroneous acceptances or rejections.

Deep Learning and Computer Vision in Transactional Processes: Computer vision allows a process to capture information from the physical world without manual data entry. The electronic toll collector framework in [2] uses image classification to automate toll-related operations, replacing manual collection with automated recognition linked to payment. The same pattern applies to goods receipt in warehouses, damage assessment in insurance, and quality inspection in manufacturing, where an image becomes the trigger for a downstream financial or logistical transaction.

Spatio-Temporal Deep Learning in Operations: Logistics and service operations depend on knowing where vehicles and people will be. Ul Abideen et al. [6] proposed a deep-learning approach that plans spatial-temporal trajectories by predicting future locations, and a related study [7] models interactions among pedestrians in three dimensions to improve trajectory prediction. In business

terms, such models can inform dispatch, delivery-time estimation, and route planning for fleets, and pedestrian-flow forecasting can support staffing and layout decisions in stations, venues, and retail spaces.

AI in Market and Trading Processes: The peer-to-peer energy trading simulator in [3] illustrates a business process in which many participants generate, consume, store, and trade a resource. Such processes are natural candidates for predictive models that forecast demand and supply, optimize prices, and detect anomalous transactions, because the volume and speed of decisions exceed what manual coordination can handle.

RPA, AI Bots, and Identity: When software bots execute business processes, they require credentials, permissions, and audit trails in the same way that human employees do. Hussain et al. [8] proposed an identity and access management framework for RPA and AI bots, emphasizing that each bot's permissions must be defined, enforced, and audited. This is a business-process concern as much as a security concern, because an over-privileged bot can execute an entire process, including payments or record changes, without the segregation of duties that applies to human staff.

Low-Code/No-Code and Intelligent Automation: Ahmed et al. [4] discuss the integration of low-code/no-code technologies with AI/ML for intelligent automation. These platforms allow business users, rather than only professional developers, to model processes and embed AI services, shortening the time between identifying an improvement and deploying it. The same accessibility, however, can produce inconsistent data models, undocumented logic, and uncontrolled deployment of AI components unless platform use is governed.

Real-Time AI Decision Processing: Some business processes, such as payment authorization, fraud screening, and dynamic pricing, require decisions within strict time limits. Ali [9] surveys the scheduling, inference, and security foundations of real-time AI decision processing, and a companion survey [10] examines divide-and-conquer strategies that split large decision tasks into smaller subproblems that can be processed in parallel. Established real-time operating-system concepts [11], such as deadline-driven scheduling and bounded response times, provide the vocabulary for specifying these requirements. For business processes, these works highlight that a model is only useful if its answer arrives before the process step that depends on it.

Behavioral ML, Fraud, and Security Monitoring: Digital sales and service channels must distinguish legitimate users from automated abuse. Khalil et al. [12] investigated mouse-pointing endpoint prediction using kinematic template matching to distinguish humans from bots, an approach that can be integrated into account creation, checkout, and booking processes. At the infrastructure level, Shahbaz et al. [13] describe the deployment of honeypot systems for threat detection and intelligence, which can provide early warning of attacks against the systems on which automated processes depend. Because bots and inference services run on shared hosts, protecting operating-system kernel code and data [14] is a

precondition for trusting their outputs. Where process inputs come from connected vehicles, as in fleet logistics or usage-based charging, trust-based management of vehicles communicating over wireless sensor networks [15] can score the reliability of each data source before its data triggers a transaction.

Privacy-Preserving Collaborative Learning:

Organizations often benefit from training models on data held by several parties, for example across branches or partner firms, but cannot share raw records. Federated learning lets each participant train locally and share only model updates; a simplified aggregation is $\theta_{t+1} = \sum_k \frac{n_k}{n} \theta_k$. Bounceur et al. [16] demonstrate this approach for collaborative intrusion detection, and the same principle applies to cross-organization fraud models in business processes.

Classification, Standards, and Process Data:

Business processes depend on shared categories such as customer types, product codes, and transaction statuses. Bowker and Star [17] show that classification systems shape what an organization can see and act upon, and Edwards [18] describes the effort, or data friction, required to move and reconcile data across systems. The European Interoperability Framework [19] similarly treats legal, organizational, semantic, and technical alignment as preconditions for seamless services. These works imply that an AI model inherits the categories, gaps, and inconsistencies of the process data on which it is trained.

III. PROBLEM STATEMENT

AI tools are frequently adopted as point solutions: a classifier is added to one step, a bot to another, and a forecasting model to a third. Each component may be evaluated by its own accuracy, yet the value of a business process is determined by its end-to-end outcome. This value can be expressed conceptually as

$$V = f(Q, T, C, R) \quad (1)$$

where Q denotes outcome quality, T end-to-end time, C cost, and R risk, including compliance, security, and reputational risk. Model accuracy influences Q but does not by itself determine T , C , or R . A highly accurate model placed in a step that is not a bottleneck may leave cycle time unchanged, while an automated decision without appropriate controls may reduce cost while increasing risk.

The problem addressed in this paper is therefore the absence of a structured, process-level view of how AI, ML, and DL affect business processes, combined with the absence of architectural and governance mechanisms that ensure improvements in one dimension of (1) are not achieved at the expense of another.

IV. AI ACROSS THE BUSINESS PROCESS LIFECYCLE

The effect of AI differs according to the stage of the process lifecycle at which it is applied. Fig. 1 summarizes five stages and the AI capabilities most relevant to each.




Fig. 1. AI, ML, and DL capabilities across the business process lifecycle.

Design: During process design, AI can support the discovery of how a process actually runs by analyzing event logs and transaction histories, revealing loops, rework, and deviations from the intended model. Because process models are graphs of activities and transitions, detecting duplicate or near-identical models in a large process repository is a graph-matching problem, for which target-cell selection strategies developed for graph isomorphism [20] are relevant. Low-code/no-code platforms [4] shorten the path from a redesigned process model to a working application. The principal effect at this stage is process transformation: the process can be redesigned around capabilities that did not previously exist, rather than merely accelerated.

Execute: During execution, RPA bots and recognition models perform tasks previously carried out manually, such as reading documents, capturing images, and entering data into multiple systems. Image-based capture in [2] is an example of this task-automation effect. Because bots act on behalf of the organization, their identities and permissions must be managed explicitly [8].

Decide: Decision steps are where ML and DL have the largest qualitative effect, because they replace or support human judgement. Classification, scoring, and forecasting models can approve, reject, prioritize, or price cases. Real-time decisions further require that inference fits within the time budget of the process step [9], [10].

Monitor: Monitoring uses anomaly detection, fraud and bot detection [12], and security intelligence [13] to identify cases and events that deviate from expected behaviors. Monitoring is also where the performance of the AI components themselves must be tracked, because the data on which a model was trained may cease to represent current conditions.

Improve: Outcomes, overrides, and exceptions recorded during execution become training data for model retraining and evidence for process redesign. This feedback loop is what distinguishes intelligent process automation from static automation: the process can improve continuously, provided that the feedback is captured systematically and reviewed before models are updated.

The time effect of AI can be made explicit by decomposing the cycle time of a case into waiting and service components:

$$T_{case} = \sum_i (w_i + s_i) \quad (2)$$

where w_i is the waiting time before activity i and s_i is its service time. Task automation primarily reduces s_i . In many processes, however, the larger share of cycle time

lies in w_i , caused by queues, handoffs, and approvals. AI therefore reduces end-to-end time substantially only when it removes handoffs or approval queues, for example by deciding routine cases automatically, rather than when it merely speeds up an activity that was not a bottleneck. A related measure is the straight-through processing rate,

$$STP = N_{auto} / N_{total} \quad (3)$$

where N_{auto} is the number of cases completed without human intervention and N_{total} the total number of cases. Table I maps representative techniques onto process functions and their principal risks.

TABLE I. AI, ML, AND DL TECHNIQUES MAPPED TO BUSINESS PROCESS FUNCTIONS

Technique	Process function	Cited example	Principal risk
Supervised ML classification	Identity verification, case classification	Face recognition [1]	False match / false rejection
DL computer vision	Automated capture and billing	Toll collection [2]	Misrecognition, fraud
Spatio-temporal DL	Dispatch, ETA, routing, staffing	Trajectory prediction [6], [7]	Location error
Prediction and optimization	Pricing, trading, scheduling	P2P energy trading [3]	Data integrity
RPA and AI bots	Rule-based back-office execution	Bot IAM [8]	Over-privileged bots
Low-code/no-code + AI	Rapid process (re)design	Intelligent automation [4]	Ungoverned deployment
Real-time inference	Time-critical decisions	Real-time AI [9], [10]	Deadline misses
Behavioural ML	Fraud and bot screening	Human-bot detection [12]	Legitimate users rejected
Federated learning	Cross-party model training	Federated IDS [16]	Model poisoning

V. PROPOSED FRAMEWORK

The analysis above indicates that the benefits of AI depend on conditions outside the model itself. The proposed socio-technical framework, shown in Fig. 2, therefore organizes an AI-enabled business process into six layers.




Fig. 2. Proposed six-layer socio-technical framework for AI-enabled business processes.

Process Data and Integration Layer: This layer captures event logs, documents, images, sensor readings, and transaction records, and exposes them through application programming interfaces with shared vocabularies and metadata describing source, time, quality, ownership, and permitted use [19]. A socio-technical framework for secure and privacy-preserving interoperability between heterogeneous intelligent systems, on which this layer builds, is presented in [21].

AI/ML/DL Intelligence Layer: This layer hosts classification, prediction, computer vision, optimization, and anomaly-detection models. In practice these services and the process data they consume are frequently hosted on cloud platforms, whose architecture, intelligent applications, and security considerations are reviewed in [22]. Each model should be registered with its purpose, training data, performance on validation data, and the process steps that consume its outputs.

Automation and Execution Layer: This layer contains RPA bots, workflow engines, and low-code/no-code applications [4] that act on model outputs. It also implements the decision-routing logic described in Section VI, which determines whether a case is completed automatically or sent to a person.

Security, Identity, and Privacy Layer: Every bot and service operates under its own identity with least-privilege permissions and audit logging [8]. The layer includes fraud and bot detection on customer-facing channels [12], threat intelligence such as honeypots [13], and privacy-preserving learning where data cannot be pooled [16].

Governance and Human Oversight Layer: This layer defines who owns each process and model, which decisions may be automated, what thresholds apply, how long data may be retained, and how automated decisions are reviewed, explained, and corrected.

Business Value and Application Layer: This layer represents the processes that deliver value, such as order-to-cash, procure-to-pay, customer onboarding, trading, and service. Its outcomes are measured with the process-level metrics defined in Section XI and fed back to the lower layers.

VI. DECISION ROUTING AND HUMAN-IN-THE-LOOP

The most consequential design choice in an AI-enabled process is how model outputs are converted into actions. A model typically produces a score or confidence p for each case. Rather than a single cut-off, a three-zone routing rule is proposed:

$$\begin{aligned} &\text{auto-accept if } p \geq \tau_h; \text{ human review if } \tau_l < p < \tau_h; \\ &\text{auto-reject if } p \leq \tau_l \end{aligned} \quad (4)$$

where τ_l and τ_h are lower and upper thresholds. Cases in the middle zone, where the model is least certain, are reviewed by a person. The expected cost of the decision step for a given pair of thresholds is

$$E[C] = c_{FP} \cdot FP + c_{FN} \cdot FN + c_H \cdot N_H \quad (5)$$

where FP and FN are the numbers of false positives and false negatives among automated decisions, N_H is the

number of cases routed to human review, and c_{FP} , c_{FN} , and c_H are the corresponding unit costs. Widening the review zone reduces automated errors but increases review cost and lowers the straight-through rate in (3).

This formulation has an important managerial implication: thresholds should be chosen from business costs, not from model accuracy alone. In a customer-onboarding process based on face verification [1], a false acceptance may allow an impostor to open an account, whereas a false rejection delays a legitimate customer. Where the cost of a false acceptance is much higher, τ_h should be set conservatively, even if this lowers apparent automation. The same model may therefore be deployed with different thresholds in different processes, and the thresholds should be reviewed by the process owner rather than fixed by the development team.

VII. DATA, CLASSIFICATION, AND PROCESS STANDARDIZATION

AI models learn from process data, and process data reflect the categories and practices of the organization that produced them [17]. If two branches record the same transaction status with different codes, or if a field named CustomerID refers to a billing account in one system and to a person in another, a model trained on the combined data will learn inconsistent patterns. The effort required to reconcile such differences can be expressed as a data-friction cost [18]:

$$F = C_{\text{extract}} + C_{\text{clean}} + C_{\text{map}} + C_{\text{validate}} \quad (6)$$

where the terms represent extraction, cleaning, semantic mapping, and validation costs. Conversion between representations often trades fidelity for size: for three-dimensional product models used in online catalogues and configurators, quadric-error-metric simplification [23] reduces model complexity while bounding geometric error, and the same trade-off arises whenever detailed records are mapped onto simpler shared schemas. AI can reduce some of these costs through automated classification, entity matching, and anomaly detection, but it can also increase them if models are built on undocumented assumptions. Standardized vocabularies, documented data dictionaries, and explicit mapping rules, as recommended in public-sector interoperability practice [19] and in AI-enabled interoperability frameworks [21], are therefore prerequisites rather than optional refinements. A practical corollary is that process standardization should normally precede AI deployment: automating a poorly defined or wasteful process tends to reproduce its weaknesses at greater speed.

VIII. ORGANIZATIONAL AND WORKFORCE EFFECTS

AI changes the composition of work within a process more often than it removes the process owner's responsibility. Routine data entry and rule-based checking move to bots and models, while people increasingly handle exceptions, review uncertain cases, and resolve disputes. This shift creates new roles, including model owners who are accountable for model behaviour in a specific process, exception handlers who deal with cases that automation

cannot resolve, and process analysts who interpret monitoring data and redesign workflows.

Low-code/no-code platforms [4] extend development capacity to business staff, which can accelerate improvement but also blurs the boundary between business and IT responsibilities. Organizations therefore need clear rules on who may deploy which AI components, how such components are tested, and who maintains them. Accountability must also remain traceable: when a bot or model executes a decision, the organization should be able to identify which version of which model acted, under which identity, and under whose authority.

Finally, staff who review AI recommendations can be influenced by them. If reviewers routinely accept model suggestions without examination, the nominal human oversight in (4) provides little protection. Review interfaces should therefore present the evidence behind a recommendation, and the rate at which reviewers override the model should be monitored as an indicator of whether oversight is meaningful.

IX. RISKS AND CONTROLS

Introducing AI into business processes creates risks that are partly technical and partly organizational. Software bots expand the set of identities that can act within enterprise systems, and over-privileged bots can bypass segregation of duties [8]. Customer-facing processes are exposed to automated abuse that behavioural detection can mitigate [12], although such detection may also reject legitimate users. Systems that host automated processes are targets for attack, and threat intelligence from honeypots can provide early warning [13]. Automated decisions in processes such as hiring, credit assessment, and biometric verification can also reproduce historical inequities through proxy variables even when sensitive attributes are removed, which calls for fairness testing and bias-mitigation measures across the model lifecycle [24]. Collaborative models trained through federated learning reduce raw-data sharing but remain vulnerable to poisoned updates [16]. Linking datasets for AI purposes can enable secondary uses that were not part of the original purpose of collection, a concern documented in the context of large-scale information-system interoperability [25]. Table II maps these risks to controls and to the layer of the framework in which each control is implemented.

TABLE II. RISK-CONTROL MAPPING FOR AI-ENABLED BUSINESS PROCESSES

Risk	Process consequence	Control	Layer
Over-privileged bots	Unauthorized payments or record changes	Bot identities, least privilege, audit [8]	Security
Erroneous automated decisions	Wrong approvals or rejections	Three-zone routing, cost-based thresholds	Execution, Governance
Automated abuse of channels	Fake accounts, fraud	Behavioural bot detection [12]	Security
Attacks on process systems	Disruption, data loss	Threat intelligence, honeypots [13]	Security

Risk	Process consequence	Control	Layer
Compromised bot hosts	Manipulated bot actions	Kernel code and data protection [14]	Security
Untrusted vehicle or sensor data	Wrong billing or dispatch	Trust-based source scoring [15]	Security
Model poisoning	Degraded shared models	Robust aggregation, update validation [16]	Security
Secondary use of linked data	Privacy and compliance breach	Purpose limitation, access rules [25]	Governance
Inconsistent data semantics	Biased or unreliable models	Shared vocabularies, data dictionaries [19]	Data
Algorithmic bias	Discriminatory approvals or rejections	Fairness testing, bias mitigation [24]	Intelligence, Governance
Model drift	Gradual loss of quality	Monitoring, scheduled revalidation	Intelligence
Ungoverned citizen development	Undocumented logic	Platform governance, testing [4]	Governance

X. APPLICATION SCENARIOS

Customer Onboarding and Identity Verification: A bank or telecommunications provider can combine document capture, face verification [1], and a risk score to open accounts. Routine cases with high verification confidence are completed automatically, uncertain cases are routed to an agent under (4), and the onboarding bot operates under its own audited identity [8].

Automated Revenue Collection: In toll collection [2] and similar usage-based charging, image recognition triggers billing without manual handling. The process benefit is largest where manual collection created queues, consistent with the waiting-time argument in (2), while misrecognition and fraud require exception handling and audit.

Trading and Market Processes: In energy trading platforms [3] and analogous marketplaces, forecasting and optimization models can support pricing and matching of supply and demand. Because many automated decisions occur in short intervals, decision latency must fit within the market's time constraints [9].

Back-Office Automation: Invoice processing, reconciliation, and report generation can be automated with RPA bots and low-code/no-code applications that call ML services [4]. Governance of bot permissions [8] and documentation of citizen-developed logic determine whether these gains remain controllable as the number of bots grows.

Digital Channel Protection: E-commerce checkout, ticket booking, and account registration can use behavioural signals [12] to distinguish genuine users from bots, while infrastructure-level threat intelligence [13] protects the systems on which these processes run.

XI. PROCESS-LEVEL EVALUATION FRAMEWORK

Evaluating an AI-enabled process only by classification accuracy is insufficient. Model-level metrics remain necessary, namely Accuracy = $(TP + TN)/(TP + TN + FP + FN)$, Precision = $TP/(TP + FP)$, Recall = $TP/(TP + FN)$, and F1 = $2(Precision \times Recall)/(Precision + Recall)$, but they must be complemented by process-level measures. Table III defines the recommended measures.

TABLE III. RECOMMENDED PROCESS-LEVEL EVALUATION MEASURES

Measure	Definition	Dimension of (1)
Cycle time	Mean and 90th-percentile end-to-end time per case, per (2)	Time (T)
Straight-through rate	Share of cases completed without human intervention, per (3)	Time / Cost
Exception rate	Share of cases requiring manual handling or rework	Quality / Cost
First-pass yield	Share of cases completed correctly without correction	Quality (Q)
Cost per case	Total processing cost divided by number of cases, incl. review cost c_H	Cost (C)
Decision error cost	Expected cost of automated errors, per (5)	Risk (R)
Override rate	Share of model recommendations changed by reviewers	Oversight (H)
Model F1 and drift	F1 on recent labelled cases compared with validation F1	Quality / Risk
Security and privacy incidents	Incidents involving bots, models, or linked data	Risk (R)

These measures can be combined into an evaluation vector

$$E = (Q, T, C, R, H) \quad (7)$$

where H represents the effectiveness of human oversight. For empirical validation, the measures should be collected before and after AI deployment and, where possible, compared with a control group of cases or sites processed without the intervention, so that improvements can be attributed to the AI component rather than to seasonal or organizational changes.

XII. DISCUSSION

The analysis indicates that AI, ML, and DL affect business processes through three distinct mechanisms. Task automation reduces service time and manual effort; decision augmentation changes the quality, consistency, and speed of judgements; and process transformation enables redesigned processes built around new capabilities. These effects are complementary but not automatic. Automation of a non-bottleneck task leaves cycle time largely unchanged, decision automation without cost-based thresholds can lower cost while raising risk, and transformation without governance can produce a collection of uncoordinated bots and models.

A second observation is that the most visible metric, model accuracy, is often the least decisive for business value. The value of an AI-enabled process depends on where the model is placed in the flow of work, how its

outputs are routed, how its errors are priced, and whether the data it uses are consistent. The proposed framework and the decision-routing model are intended to make these dependencies explicit, so that managers and engineers evaluate AI investments in terms of process outcomes rather than component performance.

Finally, the reviewed applications show that the same underlying techniques recur across domains: classification in identity verification, vision in revenue collection, forecasting in trading, and behavioural analysis in channel protection. This suggests that organizations can build reusable AI capabilities, but that each reuse requires process-specific thresholds, controls, and ownership.

XIII. RECOMMENDED DESIGN PRINCIPLES

- 1. Process Before Technology:** Standardize and simplify a process before automating it, so that AI does not accelerate existing waste.
- 2. Target Bottlenecks:** Apply AI where it removes waiting time, handoffs, or approval queues rather than only where tasks are easy to automate.
- 3. Cost-Based Thresholds:** Set decision thresholds from the business costs of errors and review, as in (5).
- 4. Human Review of Uncertain Cases:** Route low-confidence cases to people and monitor override rates to keep oversight meaningful.
- 5. Bots as Accountable Identities:** Give every bot its own least-privilege identity and audit trail.
- 6. Shared Data Semantics:** Maintain vocabularies, data dictionaries, and mapping rules for process data.
- 7. Governed Citizen Development:** Allow low-code/no-code innovation within defined testing and approval rules.
- 8. Privacy by Design:** Limit linked data to authorized purposes and use federated learning where raw data cannot be pooled.
- 9. Continuous Monitoring:** Track model drift, exception rates, and security events alongside process KPIs.
- 10. Outcome-Based Evaluation:** Judge AI investments with the process-level vector in (7), not with model accuracy alone.

XIV. LIMITATIONS

This study is a conceptual and literature-based analysis. It does not report a new implementation, dataset, or experiment, and no performance values are claimed for the proposed framework. The application studies cited are drawn from a limited set of domains, including identity verification, transportation, logistics, energy, cybersecurity, and automation platforms, and the framework's applicability to other sectors such as healthcare or public administration requires further examination. The formulas presented are conceptual models intended to structure evaluation, not calibrated estimates.

XV. FUTURE RESEARCH DIRECTIONS

Future work should implement the proposed framework in one or more organizations and measure the evaluation vector in (7) before and after deployment; estimate the cost parameters in (5) for specific processes and study how threshold choices affect straight-through rates and error costs; investigate how reviewer reliance on AI recommendations changes over time; examine governance models for large populations of RPA and AI bots; and evaluate federated learning for cross-organization process models such as shared fraud detection[26], [27], [28], [29], [30].

XVI. CONCLUSION

Modern AI, ML, and DL tools are changing business processes from rule-driven sequences into adaptive systems that capture information automatically, make or support decisions, and learn from their outcomes. Their effect, however, is not determined by algorithmic capability alone. It depends on whether AI is placed at genuine bottlenecks, whether automated decisions are bounded by cost-based thresholds and human review, whether bots are governed as accountable identities, and whether process data share consistent meaning.

The six-layer framework, the three-zone decision-routing model, and the process-level evaluation framework proposed in this paper provide a structured basis for designing and assessing AI-enabled business processes. Organizations that combine AI capability with process standardization, security, privacy, and governance are more likely to realize sustainable improvements in quality, time, cost, and risk than those that adopt AI as isolated point solutions.

REFERENCES

- [1] N. Ahmed, F. A. Khan, Z. Ullah, H. Ahmed, T. Shahzad, and N. Ali, "Face recognition comparative analysis using different machine learning approaches," *Advances in Science and Technology Research Journal*, vol. 15, no. 1, pp. 265–272, 2021, doi: 10.12913/22998624/132611.
- [2] A. Mujeeb, N. Ahmed, H. Arshed, and F. A. Khan, "Electronic toll collector framework," *Advances in Science and Technology Research Journal*, vol. 16, no. 1, pp. 294–302, 2022, doi: 10.12913/22998624/144537.
- [3] N. Ahmed, F. Ajmal, M. Usman, Z. Javed, M. H. Abid, M. Yousaf, and M. Z. Noohani, "A software capable of simulating peer-to-peer energy trading," *International Journal of Applied and Physical Sciences*, vol. 7, pp. 18–26, 2021, doi: 10.20469/ijaps.7.50003.
- [4] N. Ahmed, M. Shahbaz, H. S. U. R. Khalil, Z. Javed, M. A. Khalid, F. A. Khan, and Z. Maroof, "The synergy of low-code/no-code and AI/ML: Enhancing intelligent automation," *World Journal of Engineering and Technology*, vol. 13, no. 4, pp. 754–763, 2025, doi: 10.4236/wjet.2025.134048.
- [5] M. Shahbaz, M. Emran, R. H. S. Kannavalli, F. A. Khan, and M. Z. Noohani, "Toward AI-augmented applied systems: Emerging AI and LLM directions across trajectory prediction, cybersecurity, and engineering automation," *International Journal of Computer & Organization Trends*, vol. 3, pp. 9–15, 2026, doi: 10.14445/22492593/IJCOT-V16I3P302.
- [6] Z. Ul Abideen, X. Sun, C. Sun, and H. S. U. R. Khalil, "Improved deep learning-based approach for spatial-temporal trajectory planning via predictive modeling of future location," *KSII Trans. Internet Inf. Syst.*, vol. 18, no. 7, pp. 1726–1748, 2024.
- [7] Z. Ul Abideen, N. Ahmed, H. S. U. R. Khalil, and M. Shahbaz, "Advancing pedestrian trajectory prediction with interaction-aware

- 3D-dual contextualized modeling,” *Egyptian Informatics Journal*, vol. 31, Art. no. 100742, 2025, doi: 10.1016/j.eij.2025.100742.
- [8] S. T. Hussain, M. Shahbaz, M. Emran, and M. H. Junejo, “Securing the bot eco system: IAM framework for RPA and AI BOTs,” in *Proc. 2026 IEEE 5th Int. Conf. Computing Mach. Intell. (ICMI)*, 2026, pp. 1–9, doi: 10.1109/ICMI68585.2026.11539844.
- [9] F. Ali, “Real-time decision processing in artificial intelligence systems: A survey of scheduling, inference, and security foundations,” *Int. J. Emerg. Trends Eng. Dev.*, vol. 16, no. 5, pp. Aj1–Aj4, 2026, doi: 10.5281/zenodo.22728746.
- [10] F. Ali, “Divide-and-conquer decomposition strategies for real-time artificial intelligence decision systems: A comparative survey,” *Int. J. Emerg. Trends Eng. Dev.*, vol. 16, no. 5, pp. Ak1–Ak6, 2026, doi: 10.5281/zenodo.22728743.
- [11] N. Ahmed, *Understanding Real Time OS Concepts*, vol. 6, no. 9. Zenodo, 2015, doi: 10.5281/zenodo.22148015.
- [12] H. S. U. R. Khalil, N. Ahmed, and Y. Yang, “Mouse pointing endpoint prediction to distinguish between human and bot using kinematic template matching,” in *HCI International 2025—Late Breaking Papers, Lecture Notes in Computer Science*, vol. 16337. Cham, Switzerland: Springer, 2026, pp. 343–364, doi: 10.1007/978-3-032-12801-0_23.
- [13] M. Shahbaz, R. Imtiaz, M. H. Junejo, M. Emran, N. Ahmed, and H. S. U. R. Khalil, “Deployment of honeypot system for threat detection and intelligence in modern cybersecurity architectures,” in *Proc. 2026 1st Int. Conf. Innovations in Information and Communication Technology (IICT)*, 2026, pp. 1–5, doi: 10.1109/IICT68374.2026.11468479.
- [14] N. Ahmed, H. Y. Rana, M. T. Qureshi, and B. Shakoor, “How to protecting kernal code and data,” *International Journal of Computer Organization Trends*, vol. 5, no. 4, pp. 20–23, 2015, doi: 10.14445/22492593/IJCOT-V23P301.
- [15] M. H. Junejo, M. Shahbaz, and M. Emran, “Trust-based management for Internet of Vehicles via wireless sensor networks: Safeguarding trust during journeys,” in *Proc. 2026 IEEE 5th Int. Conf. Computing Mach. Intell. (ICMI)*, 2026, pp. 1–10, doi: 10.1109/ICMI68585.2026.11539803.
- [16] A. Bounceur, M. Ibrahim, M. Saqib et al., “Privacy-preserving multiclass intrusion detection in in-vehicle CAN networks using federated learning,” *Sci. Rep.*, 2026, doi: 10.1038/s41598-026-64490-z.
- [17] G. C. Bowker and S. L. Star, *Sorting Things Out: Classification and Its Consequences*. Cambridge, MA, USA: MIT Press, 1999.
- [18] P. N. Edwards, *A Vast Machine: Computer Models, Climate Data, and the Politics of Global Warming*. Cambridge, MA, USA: MIT Press, 2010.
- [19] European Commission, *New European Interoperability Framework: Promoting Seamless Services and Data Flows for European Public Administrations*. Brussels, Belgium: Publications Office of the European Union, 2017.
- [20] U. C. Sankal, N. Ahmed, F. A. Khan, and R. Hanni, “Target cell selection algorithms in graph isomorphism problem,” *International Journal of Scientific and Technical Research in Engineering*, vol. 6, no. 4, pp. 35–47, Jul.–Aug. 2021, doi: 10.5281/zenodo.22149607.
- [21] M. Qasim, L. Tabish, and F. Ali, “AI-enabled interoperability in intelligent digital systems: A socio-technical framework for secure, scalable, and privacy-preserving integration,” *International Journal of Advanced Multidisciplinary Research and Educational Development*, vol. 2, no. 5, pp. 538–546, Sep.–Oct. 2026. [Online]. Available: <https://ijamred.com/volume2/issue5/paper-details/IJAMRED-V2I5P63.html>
- [22] M. Shahbaz, M. Emran, R. H. S. Kannavalli, F. A. Khan, and M. Z. Noohani, “Cloud computing: Architecture, intelligent applications, security, and open-source case studies,” *International Journal of Computer & Organization Trends*, vol. 16, no. 3, pp. 1–8, 2026, doi: 10.14445/22492593/IJCOT-V16I3P301.
- [23] H. Chang, D. Zhang, Y. Dong, Y. Yang, H. S. U. R. Khalil, and X. Wang, “Two-round optimization algorithm based on quadric error metrics,” *IEEE Access*, vol. 13, pp. 30021–30035, 2025, doi: 10.1109/ACCESS.2025.3541436.
- [24] M. Qasim, L. Tabish, and F. Ali, “Mitigating algorithmic bias in information systems: A comprehensive framework for fair artificial intelligence,” *International Journal of Advanced Multidisciplinary Research and Educational Development*, vol. 2, no. 5, pp. 547–552, Sep.–Oct. 2026. [Online]. Available: <https://ijamred.com/volume2/issue5/paper-details/IJAMRED-V2I5P64.html>
- [25] European Union Agency for Fundamental Rights, *Interoperability and Fundamental Rights Implications*. Vienna, Austria, 2018.
- [26] A. Bounceur, M. Ibrahim, M. Saqib et al., “Privacy-preserving multiclass intrusion detection in in-vehicle CAN networks using federated learning,” *Sci. Rep.*, 2026, doi: 10.1038/s41598-026-64490-z.
- [27] N. Ahmed, F. A. Khan, Z. Ullah, H. Ahmed, T. Shahzad, and N. Ali, “Face recognition comparative analysis using different machine learning approaches,” *Advances in Science and Technology Research Journal*, vol. 15, no. 1, pp. 265–272, 2021, doi: 10.12913/22998624/132611.
- [28] A. Mujeeb, N. Ahmed, H. Arshed, and F. A. Khan, “Electronic toll collector framework,” *Advances in Science and Technology Research Journal*, vol. 16, no. 1, pp. 294–302, 2022, doi: 10.12913/22998624/144537.
- [29] Z. Ul Abideen, N. Ahmed, H. S. U. R. Khalil, and M. Shahbaz, “Advancing pedestrian trajectory prediction with interaction-aware 3D-dual contextualized modeling,” *Egyptian Informatics Journal*, vol. 31, Art. no. 100742, 2025, doi: 10.1016/j.eij.2025.100742.
- [30] H. S. U. R. Khalil, N. Ahmed, and Y. Yang, “Mouse pointing endpoint prediction to distinguish between human and bot using kinematic template matching,” in *HCI International 2025—Late Breaking Papers, Lecture Notes in Computer Science*, vol. 16337. Cham, Switzerland: Springer, 2026, pp. 343–364, doi: 10.1007/978-3-032-12801-0_23.