IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 7, November 2026
Impact Factor: 8.2

An Integrated Artificial Intelligence and Machine Learning Framework for Intelligent Decision-Making

Prof. Amit Narayan Shinde[1] and Prof. Sudhakar Babulal Baviskar[2]
Assistant Professor, SSPM's, Rani Laxmibai Mahavidyalaya, Parola, Jalgaon, KBCNMU, Jalgaon[1]
Lecturer, R. C. Patel College of Engineering and Polytechnic, Shirpur, DBATU, Lonere and MSBTE Mumbai[2]
amitshindednu@gmail.com and baviskar.sudhakar276@gmail.com

Abstract: *Artificial intelligence (AI) and machine learning (ML) are increasingly used to support decisions that involve large datasets, complex relationships, uncertainty, and time-sensitive choices. Yet, the value of these technologies depends on more than predictive accuracy. Effective decision support also requires reliable data, appropriate model selection, interpretable outputs, meaningful human oversight, and governance throughout the system lifecycle. This paper develops an integrated AI and ML framework for intelligent decision-making by synthesizing research on organizational decision structures, human–AI collaboration, explainable AI, and trustworthy AI governance. The study adopts a qualitative, secondary-research-based conceptual design. It critically reviews established peer-reviewed studies and institutional guidance, then organizes their contributions into a lifecycle model comprising problem definition, data preparation, analytical modelling, decision intelligence, human review, implementation, and continuous monitoring. The proposed framework distinguishes prediction from decision-making and emphasizes that algorithmic outputs should be combined with domain knowledge, organizational objectives, constraints, and accountability mechanisms. The analysis identifies recurring implementation challenges, including data and model drift, bias, opacity, automation bias, fragmented accountability, and weak evaluation of real-world outcomes. It argues that a human-centred, risk-sensitive architecture is more appropriate than a universal automation model. The framework contributes an integrated structure for designing and evaluating AI-enabled decisions and offers practical implications for organizations seeking to improve decision quality while preserving transparency, fairness, and human agency.*

Keywords: artificial intelligence; machine learning; intelligent decision-making; decision support systems; explainable AI; human–AI collaboration; AI governance

I. INTRODUCTION

Decision-making is a core organizational and societal activity. Managers allocate resources, clinicians assess treatment options, financial institutions evaluate risk, public agencies prioritize services, and industrial operators respond to changing conditions. These decisions are often made under uncertainty and involve multiple objectives, incomplete information, and consequences that extend beyond the immediate task. Digitalization has increased both the volume and velocity of available information, making conventional approaches to analysis difficult to scale. Artificial intelligence (AI), particularly machine learning (ML), offers methods for detecting patterns, estimating outcomes, classifying cases, recommending actions, and supporting decisions across complex environments.

AI is a broad field concerned with systems that perform tasks associated with human intelligence, whereas ML is a family of computational methods through which systems learn patterns from data. The distinction matters because intelligent decision-making is not equivalent to applying a predictive model. A model may estimate the probability of

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 7, November 2026 Impact Factor: 8.2

customer default, equipment failure, disease progression, or supply disruption; a decision-maker must still determine what action is justified, given costs, benefits, constraints, uncertainty, and ethical considerations. A high-performing prediction can therefore lead to a poor decision when the target variable is misaligned with the real objective, the data are unrepresentative, or the recommendation is applied without appropriate context.

Research on organizational decision-making suggests that humans and algorithms have different strengths. Algorithms can process large alternative sets consistently and rapidly, while human decision-makers contribute contextual understanding, tacit knowledge, moral reasoning, and the ability to interpret ambiguous situations (Shrestha et al., 2019). Jarrahi (2018) similarly emphasizes human–AI symbiosis in environments characterized by uncertainty and equivocality. These perspectives challenge the assumption that organizations must choose between full automation and exclusively human judgment. Instead, decision processes can be structured to allocate tasks according to the relative capabilities and limitations of people and computational systems.

However, integration is not automatically beneficial. Complex models may be difficult to interpret, historical data may reproduce discriminatory patterns, and users may either over-rely on algorithmic recommendations or reject them after visible errors. Raisch and Krakowski (2021) describe an automation–augmentation paradox: automation and human augmentation are interdependent rather than mutually exclusive, and overemphasis on either can create undesirable organizational consequences. The design question is therefore not simply whether to use AI, but how to configure the relationship among data, models, people, processes, and governance.

This paper addresses that question by proposing an integrated framework for intelligent decision-making. Its central argument is that decision quality emerges from the coordinated operation of technical and socio-organizational components. These include problem formulation, data governance, analytical modelling, explanation, decision rules, human review, implementation, and monitoring. The framework is intended to be adaptable across sectors rather than optimized for one application.

1.1 Aim and Objectives

The main aim is to develop a conceptually grounded framework that integrates AI and ML capabilities with human judgment and organizational governance to support intelligent decision-making. The specific objectives are: (a) to examine major research perspectives on AI-enabled decision support; (b) to identify technical, organizational, and ethical factors influencing decision quality; (c) to develop an integrated lifecycle framework; and (d) to propose evaluation dimensions and implementation principles.

1.2 Significance of the Study

The study is significant because many AI adoption initiatives focus heavily on model accuracy while giving less attention to the full decision process. A framework that links prediction to action can help organizations avoid treating algorithmic outputs as self-justifying decisions. It also offers a common structure for researchers comparing AI-enabled decision systems and for practitioners planning responsible deployment.

II. REVIEW OF LITERATURE

2.1 AI and the Changing Structure of Organizational Decisions

Shrestha et al. (2019) examine how AI changes organizational decision structures by comparing human and algorithmic decision-making across factors such as search-space specificity, interpretability, the number of alternatives, speed, and replicability. They identify three broad arrangements: delegation of decisions to AI, sequential human–AI decision-making, and aggregation of human and AI judgments. Their contribution is important because it treats AI adoption as a question of decision architecture rather than merely a technology purchase. A limitation is that the framework is primarily conceptual and does not prescribe one universally optimal arrangement. The appropriate structure depends on task characteristics, error costs, regulation, and the consequences of a decision.

Jarrahi (2018) argues that human and AI capabilities can be complementary. AI can support analytical tasks involving large volumes of data and complex pattern recognition, while humans contribute intuition, contextual interpretation, and flexible reasoning. This perspective is especially relevant when decision problems are uncertain or ambiguous.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 7, November 2026 Impact Factor: 8.2

Nevertheless, complementarity should not be assumed to occur naturally. It requires interfaces, role definitions, training, escalation procedures, and organizational incentives that encourage critical engagement rather than passive acceptance.

Raisch and Krakowski (2021) add a paradox perspective. Their analysis cautions that automation and augmentation can reinforce one another and that treating them as separate strategic choices can produce cycles of over-automation or excessive human intervention. This insight supports a dynamic framework in which organizations periodically reassess which tasks should be automated, which should remain human-led, and where joint judgment is most appropriate.

2.2 Machine Learning and Decision Support

ML methods provide a range of analytical capabilities. Supervised learning uses labelled examples to estimate outcomes such as risk categories or demand levels. Unsupervised learning identifies structure in data without predefined outcome labels, supporting segmentation and anomaly detection. Reinforcement learning studies sequential decisions in which actions influence future states and rewards. Deep learning can represent complex patterns in high-dimensional data, but its computational and interpretive demands may be substantial. The choice of method should be driven by the decision problem, available data, required latency, error costs, and need for explanation—not by novelty alone.

Decision support systems traditionally combine data, analytical models, and user interfaces to assist decision-makers. AI-enabled systems extend this model through adaptive prediction, automated feature extraction, natural-language interaction, and recommendations. Yet prediction and prescription remain distinct. A predictive model estimates what may happen; a prescriptive system evaluates possible actions under objectives and constraints. For example, predicting high demand does not determine whether an organization should increase inventory, change prices, or redirect capacity. The action depends on costs, service targets, risk tolerance, and operational feasibility.

2.3 Explainability, Interpretability, and Trust

The increasing use of complex ML models has intensified concern about opacity. Doshi-Velez and Kim (2017) argue that interpretability requires rigorous conceptualization and evaluation rather than being treated as a vague desirable property. They highlight the need to determine when explanations are needed and how interpretability should be assessed. Their work is a methodological contribution: it shifts attention from the mere production of explanations to whether those explanations are appropriate for the intended users and tasks.

Ribeiro et al. (2016) introduce LIME, a method that approximates a model locally with an interpretable representation to explain individual predictions. Lundberg and Lee (2017) propose SHAP, a unified approach that assigns feature-importance values to predictions using a framework grounded in Shapley values. These techniques can help users inspect model behaviour, identify influential features, and question unexpected outputs. However, explanation methods do not automatically make a model fair, causally valid, or safe. A feature-attribution explanation may describe how a model generated a prediction without establishing that the prediction is correct or that acting on it is appropriate.

Kostopoulos et al. (2024) review explainable AI-based decision support systems and emphasize the role of transparency and interpretability in AI-supported decisions. Their review reinforces the view that explanation is a design dimension of decision support, not simply a technical add-on. At the same time, the usefulness of explanations depends on audience, task complexity, and decision stakes. Excessive technical detail can overwhelm users, while simplified explanations can conceal uncertainty or limitations. Explanation design should therefore be evaluated for fidelity, comprehensibility, relevance, and its effect on user behaviour.

2.4 Human–AI Interaction and Oversight

Amershi et al. (2019) provide guidelines for human–AI interaction, emphasizing that systems should communicate capabilities and limitations, support correction, manage expectations, and enable users to respond when AI behaviour is uncertain or wrong. These principles connect interface design to decision quality. A user interface that presents a single confident recommendation without uncertainty, evidence, or a route to challenge the output can encourage inappropriate reliance.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 7, November 2026 Impact Factor: 8.2

Human oversight is not equivalent to placing a person at the end of an automated pipeline. Meaningful oversight requires sufficient information, time, competence, authority, and organizational support to question or override system outputs. If a human reviewer is accountable for a decision but lacks access to the underlying evidence or is penalized for disagreeing with the model, oversight may become ceremonial. A well-designed framework must therefore consider institutional power, workload, and accountability in addition to technical controls.

2.5 Responsible AI and Governance

The National Institute of Standards and Technology's AI Risk Management Framework organizes risk management around the functions Govern, Map, Measure, and Manage (Tabassi, 2023). It emphasizes that trustworthiness considerations should be incorporated throughout the AI lifecycle. This approach is valuable because it treats risk as a continuous organizational responsibility rather than a one-time compliance check. The OECD AI Principles likewise promote trustworthy AI grounded in human-centred values, fairness, transparency, robustness, security, and accountability (OECD, 2019/2024).

These governance perspectives complement technical research by identifying conditions that predictive performance alone cannot address. For example, an accurate model may still be unsuitable if its training data were collected without adequate authority, if its use creates unacceptable disparate impacts, or if no mechanism exists to correct errors. Governance should specify ownership, documentation, access controls, audit trails, incident response, and procedures for withdrawing or modifying a system.

2.6 Synthesis and Research Gap

The literature offers substantial insights but is distributed across several streams. Organizational research explains how decision authority can be divided between people and algorithms; ML research focuses on prediction and optimization; explainable AI research examines model transparency; human-computer interaction studies user-system relationships; and governance frameworks address risk and accountability. These streams are often discussed separately. As a result, organizations may implement technically capable models without connecting them to decision objectives, human review, or ongoing governance.

The gap addressed in this paper is the absence, within the selected conceptual synthesis, of a single practical lifecycle structure that connects the technical pipeline to the human and organizational decision process. The proposed framework does not claim to replace specialized domain frameworks. Rather, it integrates their complementary insights into a coherent model that can be adapted and empirically tested in different contexts.

III. RESEARCH METHODOLOGY

3.1 Research Design and Approach

The paper adopts a qualitative, conceptual-analytical research design based on secondary sources. It uses a structured narrative synthesis to compare established academic contributions and institutional guidance relevant to AI-enabled decision-making. This approach is suitable because the primary output is an integrative framework, not an estimate of causal effects or a report of new field observations. The analysis is interpretive: it identifies recurring concepts, compares assumptions, and organizes them into a lifecycle architecture.

3.2 Data Sources

The evidence base consists of published academic articles, conference proceedings, and institutional publications whose bibliographic details can be verified. The selected literature includes work on organizational decision structures, human–AI symbiosis, automation and augmentation, explainable ML, human–AI interaction, and trustworthy AI governance. Institutional material from NIST and the OECD is included to represent formal governance principles. The paper does not use private organizational records, personal data, interviews, or survey responses.

3.3 Source Selection and Analytical Procedure

Sources were selected purposively for conceptual relevance, recognized scholarly or institutional provenance, and direct contribution to one or more framework dimensions. Priority was given to foundational or influential contributions and recent review work that synthesizes a defined research area. The process involved: (1) identifying

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 7, November 2026
Impact Factor: 8.2

recurring themes in the literature; (2) extracting each source's central contribution and stated limitations; (3) comparing the themes across technical, human, organizational, and governance dimensions; and (4) synthesizing the themes into a staged framework. This is a narrative conceptual review, not a registered systematic review. It therefore does not claim exhaustive database coverage, PRISMA screening, or a quantitative meta-analysis.

3.4 Proposed Integrated Framework

The framework is organized into seven connected stages. The stages are presented sequentially for clarity, but implementation is iterative: monitoring may require changes to data, models, decision rules, or even the original problem definition.

Table 1: Proposed Integrated Framework

Stage	Core activities	Principal output
1. Problem framing	Define decision, stakeholders, objectives, constraints, risk level, and success criteria.	Decision specification and accountability owner
2. Data and governance	Acquire lawful and relevant data; assess quality, representativeness, privacy, lineage, and security.	Documented and governed dataset
3. AI/ML analytics	Select baseline and candidate models; train, validate, calibrate, and compare against alternatives.	Validated predictions, scores, or detected patterns
4. Decision intelligence	Combine model outputs with thresholds, costs, constraints, scenarios, and domain rules.	Ranked options or decision recommendations
5. Human review and explanation	Present rationale, uncertainty, limitations, and alternatives; allow challenge, override, and escalation.	Reviewed and contextualized decision
6. Implementation	Execute approved action; record decision, responsible party, and rationale.	Operational action and audit record
7. Monitoring and learning	Track outcomes, drift, bias, incidents, user feedback, and changing conditions.	Updated model, policy, or process

3.5 Technical Components

Data preparation includes validation of missing values, duplicates, inconsistent units, label quality, sampling bias, and potential leakage between training and evaluation data. Data provenance and access rights should be recorded. Model development should begin with a transparent baseline and proceed to more complex models only when the added performance or capability justifies the additional cost and governance burden. Validation should use data splits that reflect the intended deployment setting; temporal or external validation is important when future conditions may differ from historical training data.

The analytical layer may combine statistical learning, tree-based ensembles, neural networks, clustering, anomaly detection, natural-language processing, or optimization methods. Model selection should consider predictive performance alongside calibration, robustness, computational cost, interpretability, and maintainability. In decision settings, probability estimates should be calibrated where they inform thresholds or expected-loss calculations. The system should also communicate uncertainty and abstain or escalate when the input is outside the model's reliable operating range.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 7, November 2026
Impact Factor: 8.2

3.6 Human and Organizational Components

The human layer defines who may receive recommendations, approve actions, override outputs, and handle exceptions. Responsibilities should be matched to competence and authority. Decision-makers need training not only in system operation but also in model limitations, uncertainty, bias, and appropriate reliance. Human review intensity should increase with potential harm, irreversibility, uncertainty, and the difficulty of contesting an outcome. Low-risk, repetitive tasks may permit bounded automation; high-impact decisions generally require stronger review, traceability, and appeal mechanisms.

3.7 Evaluation Dimensions

Evaluation should be multi-dimensional. Technical measures may include accuracy, precision, recall, F1-score, area under the receiver operating characteristic curve, calibration error, latency, and robustness, selected according to the task. Decision-level measures should assess whether recommendations improve the intended outcome, reduce avoidable losses, or support timely action. Human-centred measures include comprehension, appropriate reliance, override quality, workload, and user satisfaction. Governance measures include auditability, privacy incidents, fairness assessments, policy compliance, and clarity of responsibility. No single metric is sufficient to establish overall system quality.

Table 2: Evaluation Dimensions

Evaluation level	Illustrative questions	Possible evidence
Model	Does the model generalize and remain calibrated?	Held-out and external validation; drift tests
Decision	Does use of the model improve the target decision outcome?	Decision audits; cost-sensitive evaluation; outcome comparison
Human interaction	Do users understand when to rely, verify, or override?	Usability studies; structured observation; override analysis
Fairness and rights	Are errors and benefits distributed acceptably?	Subgroup performance; impact assessment; complaints
Governance	Can the organization explain who decided and why?	Logs; documentation; audit and incident records

IV. DATA ANALYSIS AND DISCUSSION

4.1 From Prediction to Intelligent Action

The synthesis indicates that intelligent decision-making requires a clear separation between prediction and action. A prediction is an estimate conditioned on data and a model; a decision is a choice among alternatives under objectives and constraints. Converting a probability into an action requires a decision rule. For instance, a risk score may trigger additional review only when the expected cost of an untreated event exceeds the cost of intervention. The threshold should reflect the relative consequences of false positives and false negatives, resource availability, legal obligations, and stakeholder preferences.

This distinction also explains why optimizing a model against a convenient proxy can produce undesirable outcomes. A model trained to maximize short-term conversion may recommend actions that damage long-term customer trust. A model optimized for average processing time may disadvantage complex cases. The problem-framing stage should therefore define the actual decision objective and identify distributional or long-term consequences. Where objectives conflict, multi-criteria decision analysis or explicit stakeholder prioritization may be required.

4.2 Selecting the Appropriate Human–AI Arrangement

The literature supports a contingent rather than universal approach to automation. Full delegation may be suitable for bounded, repetitive, low-impact decisions where the input space is well understood and errors are reversible. Sequential

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 7, November 2026 Impact Factor: 8.2

arrangements may be useful when AI screens or prioritizes cases for human assessment, or when a human defines context before an algorithm produces a recommendation. Aggregated arrangements may combine independent human and model assessments, particularly when both sources contribute distinct information (Shrestha et al., 2019).

The arrangement should be selected using at least five considerations: task structure, consequence severity, interpretability requirements, time sensitivity, and availability of competent reviewers. A highly structured task with clear labels and low error costs may support automation. A novel, ambiguous, or high-impact decision should retain stronger human authority. This approach is consistent with Jarrahi's (2018) account of complementarity and Raisch and Krakowski's (2021) warning that automation and augmentation interact dynamically.

4.3 The Role and Limits of Explainability

Explainability can support inspection, communication, and contestability, but it should not be treated as a guarantee of trustworthiness. LIME and SHAP offer useful ways to examine local or feature-level contributions (Lundberg & Lee, 2017; Ribeiro et al., 2016). Their outputs can help a decision-maker identify whether a recommendation is driven by plausible factors or suspicious proxies. Yet explanations may be unstable, misleading when features are correlated, or too technical for the intended audience. They may also create an illusion of certainty by presenting a coherent rationale for a model that is poorly calibrated or trained on biased data.

A responsible explanation strategy should be matched to the decision. A frontline employee may need a concise reason code, confidence range, and next step. A model risk specialist may require feature-attribution diagnostics, validation reports, and sensitivity analysis. A person affected by a consequential decision may need understandable reasons and a meaningful route to correction or appeal. Explanation quality should therefore be judged by whether it enables appropriate understanding and action, not simply by whether an explanation graphic is displayed.

4.4 Data Quality, Bias, and Model Risk

The framework places data governance before model development because data limitations propagate through every later stage. Historical records may reflect unequal access, inconsistent measurement, selective labelling, or past institutional decisions. Even when sensitive attributes are excluded, correlated variables can preserve proxy relationships. A model can consequently reproduce existing disparities while appearing statistically efficient. Subgroup evaluation and contextual review are necessary, but fairness cannot be reduced to one metric: different fairness criteria may conflict, and acceptable trade-offs depend on the application and applicable rights.

Model risk also persists after deployment. Changes in customer behaviour, economic conditions, equipment, policies, or data collection can create distribution shift and performance drift. A model validated at development time may no longer be reliable in operation. Monitoring should therefore include input-data quality, prediction distributions, calibration, subgroup error patterns, operational outcomes, and user-reported anomalies. Material drift should trigger investigation, recalibration, retraining, restricted use, or withdrawal.

4.5 Governance, Accountability, and Organizational Readiness

Governance is most effective when embedded in ordinary management processes. The NIST AI RMF's Govern, Map, Measure, and Manage functions provide a lifecycle-oriented structure for identifying context, measuring risks, and managing them over time (Tabassi, 2023). The OECD principles reinforce the importance of human-centred values, transparency, robustness, and accountability (OECD, 2019/2024). Together, these approaches suggest that governance should begin before procurement and continue through retirement.

Accountability requires more than a statement that a human remains responsible. Organizations should identify the business owner, technical owner, data steward, model validator, reviewer, and escalation authority. They should maintain records of model versions, data sources, validation results, explanation methods, policy thresholds, approvals, overrides, and incidents. Such records enable internal review and help determine whether a failure arose from data, model design, decision policy, interface design, or organizational practice.

4.6 Cross-Sector Applicability

The proposed framework can be adapted across sectors, but its operational controls should reflect domain-specific risks. In healthcare, decision support may assist triage or clinical prioritization, but clinical judgment, patient safety,

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 7, November 2026
Impact Factor: 8.2

privacy, and external validation are central. In finance, credit or fraud models require attention to disparate impact, explainability, security, and regulatory expectations. In manufacturing, predictive maintenance can reduce downtime, but sensor reliability, false alarms, and safety-critical escalation matter. In public administration, allocation or eligibility tools raise questions of due process, transparency, contestability, and unequal effects.

Table 3: Cross-Sector Applicability

Context	Potential AI/ML contribution	Critical safeguards
Healthcare	Risk stratification, diagnostic support, resource planning	Clinical validation, clinician oversight, privacy, patient safety
Finance	Credit-risk estimation, fraud detection, portfolio analysis	Fair lending assessment, model validation, explainability, security
Manufacturing	Predictive maintenance, quality inspection, process optimization	Sensor validation, safety thresholds, drift monitoring, fallback procedures
Supply chain	Demand forecasting, disruption prediction, inventory optimization	Scenario analysis, uncertainty communication, human exception handling
Public services	Case prioritization, service planning, anomaly detection	Due process, transparency, appeal routes, impact assessment

4.7 Integrated Interpretation

Across the reviewed sources, five connected conclusions emerge. First, the unit of analysis should be the complete decision process rather than the isolated model. Second, the value of AI depends on fit between task characteristics and the human–AI arrangement. Third, explainability is valuable when it enables informed scrutiny, but it cannot substitute for validity, fairness, or governance. Fourth, oversight must be operationally meaningful and supported by authority, training, and time. Fifth, reliability is a lifecycle property requiring monitoring and adaptation. These conclusions form the conceptual logic of the proposed framework.

The framework also implies a feedback relationship. Decisions produce outcomes that can be compared with expected results; discrepancies provide evidence about model performance, policy assumptions, and implementation quality. However, feedback must be interpreted carefully. If outcomes are affected by the decision itself, naïvely retraining on observed results can reinforce prior interventions or selection effects. Organizations should distinguish model error from changes caused by policy, user behaviour, or external conditions. Where feasible, evaluation designs should use suitable comparison groups, prospective validation, or carefully justified quasi-experimental methods.

4.8 Limitations of the Analysis

The analysis is limited by its narrative and purposive approach. It does not claim exhaustive coverage of all AI decision-support literature, and the proposed framework has not been validated through a field deployment or controlled experiment. The sources span different sectors and methodological traditions, which limits direct comparison. Moreover, the rapid evolution of AI means that technical practices and regulatory expectations may change. These limitations make empirical testing and periodic updating necessary before the framework is adopted as an operational standard.

V. MAJOR FINDINGS

1. Intelligent decision-making is a socio-technical process. Model performance alone cannot establish decision quality; objectives, constraints, human roles, and organizational context must also be considered.
2. Prediction and decision are distinct. Predictive scores require explicit decision rules that account for costs, consequences, uncertainty, and stakeholder priorities.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 7, November 2026 Impact Factor: 8.2

3. No single human–AI arrangement is appropriate for every task. Delegation, sequential collaboration, and aggregated judgment should be selected according to task structure and risk.
4. Explainability supports scrutiny and communication but does not independently guarantee correctness, fairness, or trustworthiness.
5. Meaningful oversight requires competence, access to relevant information, sufficient time, authority to intervene, and clear escalation procedures.
6. Data governance, subgroup evaluation, model validation, and continuous monitoring are essential because bias and performance degradation can emerge throughout the lifecycle.
7. Accountability should be operationalized through named owners, documentation, audit trails, incident management, and reviewable decision records.
8. The proposed seven-stage framework integrates problem framing, data governance, modelling, decision intelligence, human review, implementation, and continuous monitoring.

VI. CONCLUSION

This paper developed an integrated AI and ML framework for intelligent decision-making by synthesizing research on organizational decision structures, human–AI collaboration, explainability, and trustworthy AI governance. The analysis demonstrates that the benefits of AI cannot be evaluated solely through predictive accuracy or automation rates. Intelligent decisions require alignment among the problem being addressed, the quality and legitimacy of data, the suitability of analytical methods, the interpretation of model outputs, the authority of human decision-makers, and the governance arrangements surrounding deployment.

The proposed framework offers seven connected stages: problem framing; data and governance; AI/ML analytics; decision intelligence; human review and explanation; implementation; and monitoring and learning. Its central design principle is that AI should be used in a manner proportionate to task characteristics and decision risk. Bounded, reversible tasks may permit greater automation, while high-impact or ambiguous decisions require stronger human participation, explanation, and contestability. In all cases, the organization must preserve a clear record of how recommendations were converted into actions.

The study contributes conceptually by connecting technical, human, and organizational dimensions that are often examined separately. For practitioners, it provides a planning structure for evaluating AI initiatives before deployment and reviewing them after implementation. For researchers, it offers propositions and evaluation dimensions that can be tested in specific sectors. The framework should not be interpreted as proof that integrated systems always outperform human-only or model-only processes; rather, it specifies conditions and mechanisms through which responsible integration may improve decision quality.

REFERENCES

1. Amershi, S., Weld, D., Vorvoreanu, M., Fourney, A., Nushi, B., Collisson, P., Suh, J., Iqbal, S., Bennett, P. N., Inkpen, K., Teevan, J., Kikin-Gil, R., & Horvitz, E. (2019). Guidelines for human-AI interaction. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (Paper 3, pp. 1–13). Association for Computing Machinery. https://doi.org/10.1145/3290605.3300233
2. Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. arXiv. https://arxiv.org/abs/1702.08608
3. Jarrahi, M. H. (2018). Artificial intelligence and the future of work: Human-AI symbiosis in organizational decision making. Business Horizons, 61(4), 577–586. https://doi.org/10.1016/j.bushor.2018.03.007
4. Kostopoulos, G., Davrazos, G., & Kotsiantis, S. (2024). Explainable artificial intelligence-based decision support systems: A recent review. Electronics, 13(14), Article 2842. https://doi.org/10.3390/electronics13142842
5. Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. In Advances in

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 7, November 2026
Impact Factor: 8.2

Neural Information Processing Systems (Vol. 30). Curran Associates, Inc. https://proceedings.neurips.cc/paper_files/paper/2017/hash/8a20a8621978632d76c43dfd28b67767-Abstract.html

6. Organisation for Economic Co-operation and Development. (2019, updated 2024). OECD AI principles. https://oecd.ai/en/ai-principles
7. Raisch, S., & Krakowski, S. (2021). Artificial intelligence and management: The automation–augmentation paradox. Academy of Management Review, 46(1), 192–210. https://doi.org/10.5465/amr.2018.0072
8. Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?": Explaining the predictions of any classifier. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 1135–1144). Association for Computing Machinery. https://doi.org/10.1145/2939672.2939778
9. Shrestha, Y. R., Ben-Menahem, S. M., & von Krogh, G. (2019). Organizational decision-making structures in the age of artificial intelligence. California Management Review, 61(4), 66–83. https://doi.org/10.1177/0008125619862257
10. Tabassi, E. (2023). Artificial intelligence risk management framework (AI RMF 1.0) (NIST AI 100-1). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.AI.100-1