International Journal of Scientific Research & Engineering Trends
Volume 12, Issue 2, Mar-Apr-2026, ISSN (Online): 2395-566X

Hybrid AI-Agent Driven Process Optimization Framework for Enterprise Decision Systems

Dr. Nilesh Jain, Kishan Vyas, Monil Lalwani, Kushal Shah
CSE-AI / PIET, Parul University, Vadodara, Gujarat, India.

Abstract- — The scope of enterprise reporting systems has grown significantly as far as data size, dashboard interactiveness, and visual analytics are concerned. However, in most companies, these systems operate as a static support for making decisions instead of an active decision system. The issue is that even though the dashboards detect patterns, anomalies, and trends, the user should understand these results and trigger the next steps. This implies some time lag, inconsistencies, and dependency on human decision-making. Thus, in this article, we propose a mixed AI-agent framework of enterprise decision systems where the intelligence ability of artificial intelligence is complemented by the execution capabilities of intelligent agents through staged development. The framework involves data collecting, cleaning, transformation, creation of dashboards, analysis with the help of machine learning algorithms for predictions, classifications, clustering, and detecting anomalies, and finally, the use of agents that translate model outputs into workflow activities such as approval, escalating, prioritizing, and notifying. The primary idea of this research is to split the process of intelligence generation and actions taking at different process stages, thereby increasing the modularity, interpretability, and effectiveness. However, it retains its practical applicability and relevance of dashboard analysis. The paper discusses the conceptual framework, technical design, model types, logic of process mapping, criteria of assessment, governance issues, and the possibility of enterprise implementation of the proposed approach.

Index Terms— Enterprise reporting systems, AI-agent framework, decision automation, intelligent agents, machine learning, predictive analytics, anomaly detection, data transformation, interactive dashboards, workflow automation, decision support systems, modular architecture, enterprise analytics, process mapping, data-driven decision making, governance, scalability, real-time insights.

I. INTRODUCTION

Efforts toward digital transformation have rendered dashboards an indispensable element of monitoring in organizations. Regardless of whether the use of dashboards is in operational review, approval workflows, monitoring systems, or cases-based environments, the use of dashboards represents the graphical level that is needed to summarize and analyze data. But dashboards come with their limitations since they answer the question "what is happening," but they fail to answer the other question, "what should happen next." While this can be solved by human intervention, human intervention presents a challenge in practical enterprise scenarios. For instance, when category-level values exhibit signs of concentration or when records exhibit signs of high risks from past historical trends, an experienced human analyst can identify the issue and solve it accordingly. However, this will rely on the availability of such experts and their consistency. When data volume increases, the process might turn out to be very slow. This is why AI is brought help a system to understand what is happening by ranking, scoring, classifying or segmenting the data in ways that cannot be accomplished through visualization. Even with the use of AI in analyzing data, there comes another challenge since even with the application of machine learning in a certain task, some action must be taken regarding the results obtained. Agents in this case are useful because they will help in deciding what to do based on rules and logic and then performing actions. Agents' role in the system is therefore not to replace ML models but to make intelligence act as workflow logic. Justification for this research is thus to design a pragmatic architecture such that: • Dashboards will remain functional as a monitoring layer; • AI will extend dashboards through analysis reasoning; • Agents will extend AI through actionable workflow execution.

II. BACKGROUND AND MOTIVATION

Efforts toward digital transformation have rendered dashboards an indispensable element of monitoring in organizations. Regardless of whether the use of dashboards is in operational review, approval workflows, monitoring systems, or cases-based environments, the use of dashboards represents the graphical level that is needed to summarize and analyze data. But dashboards come with their limitations since they answer the question "what is happening," but they fail to answer the other question, "what should happen next." While this can be solved by human intervention, human intervention presents a

International Journal of Scientific Research & Engineering Trends
Volume 12, Issue 2, Mar-Apr-2026, ISSN (Online): 2395-566X

challenge in practical enterprise scenarios. For instance, when category-level values exhibit signs of concentration or when records exhibit signs of high risks from past historical trends, an experienced human analyst can identify the issue and solve it accordingly. However, this will rely on the availability of such experts and their consistency. When data volume increases, the process might turn out to be very slow. This is why AI is brought into play as a tool to make sense of data. In this regard, AI can help a system to understand what is happening by ranking, scoring, classifying or segmenting the data in ways that cannot be accomplished through visualization. Even with the use of AI in analyzing data, there comes another challenge since even with the application of machine learning in a certain task, some action must be taken regarding the results obtained. Agents in this case are useful because they will help in deciding what to do based on rules and logic and then performing actions. Agents' role in the system is therefore not to replace ML models but to make intelligence act as workflow logic. Justification for this research is thus to design a pragmatic architecture such that: • Dashboards will remain functional as a monitoring layer; • AI will extend dashboards through analysis reasoning; • Agents will extend AI through actionable workflow execution.

III. RESEARCH OBJECTIVES

The primary objectives of this research include: • the creation of a process-oriented architecture that integrates data analysis, artificial intelligence, and intelligent agents; • defining a stage-by-stage differentiation between intelligence production and implementation; • selecting adequate machine learning techniques for structured corporate data; • aligning intelligent agent behavior with advanced stages of corporate workflows; • developing an architecture that is modular, understandable, extensible, and adaptable for actual enterprise settings; • determining metrics for the assessment of AI outcomes and intelligent agent performance; • illustrating the role of dashboards as a link rather than a destination in enterprise decision-making systems.

V. RESEARCH CONTRIBUTION AND NOVELTY

In addition to the combined application of AI and agents, what constitutes the innovation aspect is a step-by-step distribution of the roles between AI and agents. Discussions of enterprise AI either talk about forecasting tools or automation systems, rarely making clear how these tools and systems interact in the process.

In this respect, this study introduces a conceptual framework that:

1. places AI in the initial analysis stage;
2. places agents in the subsequent implementation stage;
3. does not get rid of dashboards altogether;
4. is based on the process itself, not on tools or technologies;
5. is logically relevant for structured data environments.

VI. RELATED WORK

AI in enterprise systems has long been applied to tasks like predictive modeling, pattern recognition, and aiding decision-making based on data [1], [2]. Business Intelligence systems have grown into robust visualization environments for KPI tracking, trends assessment, and comparative analytics [3]. However, closing the loop between insight and process remains a current challenge. Machine learning techniques like logistic regression, tree-based ensemble models, and clustering have proven effective in tabular domains due to their predictive capabilities and relative interpretability [4]–[6]. In domains where explainability and decision responsibility are crucial, interpretable or semi-interpretable methods remain very relevant. Intelligent agent research has concentrated on topics like autonomous operation, sensor integration, rules management, reaction to events, collaboration, and goal-oriented performance [7]. Agent-based systems excel in scenarios where repeatable actions can be defined in terms of policies or workflows. Still, the amount of scientific work that distinguishes between AI and agents by process stage in a dashboard-oriented enterprise framework is limited. Although many approaches mention "intelligent automation," few papers elaborate on how intelligence is generated, propagated, and implemented in end-to-end pipelines. This work builds on that research path by developing a process-oriented approach rather than simply providing more algorithms.

VII. VENTERPRISE DECISION PIPELINE

The entire enterprise decision framework follows the sequence illustrated in Fig. 1. The pipeline starts with data acquisition and pre-processing, moves on to modeling and dashboard creation, and then proceeds to two specific layers, namely, the:
• AI layer, where interpretation and scoring take place; and the
• agent layer, where the interpretation results lead to business actions.

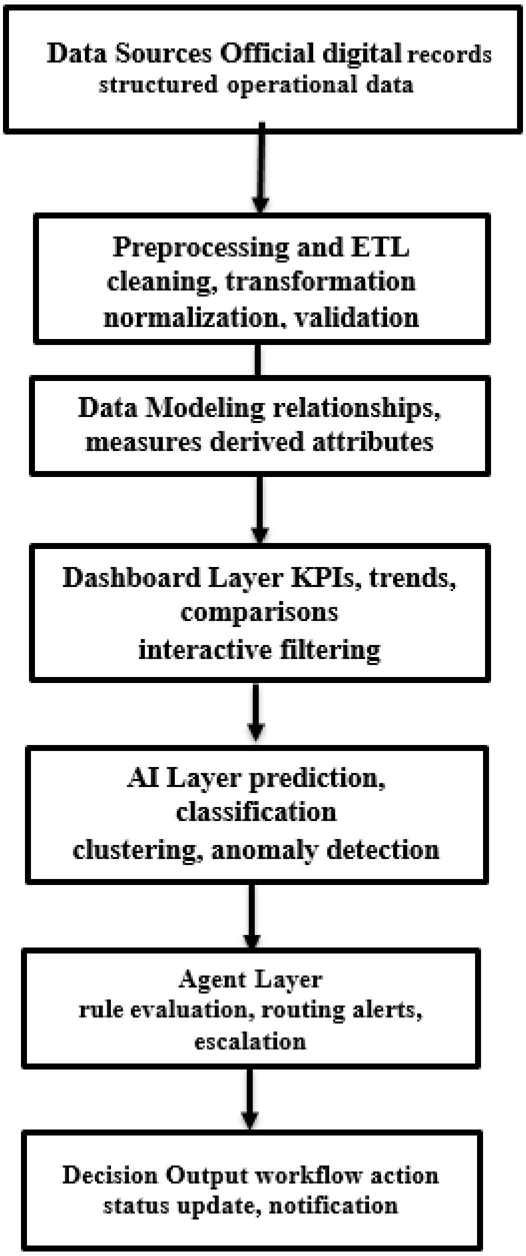

Fig. 1. Proposed end-to-end hybrid AI-agent enterprise decision pipeline

VIII. DATA AND DASHBOARD FOUNDATION

Though the research considers AI and agents, data collection pipeline is indispensable here since the quality of both elements hinges on data quality, thus requiring an equally robust foundation.

A. Data Collection

Source data for such systems is usually extracted from official digital databases, portals, transactions, records, or organizational reporting interfaces. The data extracted can include: numerical columns; categorical variables; entity metadata; timeline-based columns; process statuses.

B. Data Preparation

Raw operational data is normally to undergo:

- handling missing values;
- fixing inconsistent labels;
- standardization;
- adding additional analytical columns;
- eliminating duplicate columns;
- grouping certain data to enhance visibility on the dashboard.

C. Model Design & Visual Representation

The processed data is then designed into a model that will allow for:

- KPI calculation;
- analysis of trends;
- comparing different segments;
- calculating category-level sums;
- making decisions based on filters.

IX. THE NEED FOR A STAGE-WISE HYBRID APPROACH

One of the key points in this paper is that both AI and agent systems cannot be used randomly. They have their own specific tasks to perform. AI is more appropriate where there is uncertainty and inferential reasoning needs to be done. On the other hand, agent systems are ideal when it comes to responding to something and executing policies after identifying a situation. Without proper differentiation, it becomes difficult to design an architecture for both AI and agent systems.

X. DESIGN PRINCIPLES OF THE PROPOSED FRAMEWORK

In order to make the framework implementable and scientifically sound, it is designed based on the following design principles.

A. Modularity
Any layer in the framework could be easily substituted for another without requiring the whole system to be redesigned. For example, the logistic regression could be replaced by gradient boosting or a rule-based agent could be substituted with a more advanced agent type while still using the dashboard and data layers.

B. Explainability
Enterprise-level decisions often have to be supported. The framework should maintain outputs of the models used as well as the rules of agents used transparent and understandable to a human. This is why probabilistic inference is separated from deterministicactions in the architecture.

C. Auditability
All significant output should have audit records. In other words, if an alert was raised or a workflow was rerouted, the system could identify whether it was triggered based on the classification, threshold or anomaly.

D. Scalability
The framework should allow the increase in the volume of the input data, variety of processes and rule complexity over time. Otherwise, the solution will look both weak from the academic perspective and prone to fail in practice.

E. Human Centric Control
The framework suggests automation; however, it assumes that any process could be automated only to some extent, thus, supporting a continuum from mere recommendations to assisted automation and full rule-based automation where applicable.

F. Process Alignment
The system has to align with the process logic. It means that thresholds, routing rules, and response options have to be aligned with the process stages and responsibilities.

XI. AI LAYER DESIGN

The AI layer of the proposed architecture represents the module responsible for generating intelligence. It receives cleaned and modeled enterprise data as input and generates corresponding outputs.

A. Input Feature Space
Given the fact that the framework operates on structured data, the feature space might include:

- value-related fields;
- category labels and grouped classes;
- source-specific features;
- status-related fields;
- time duration or sequence-related fields;
- counts/frequencies/participation rates;
- features generated in the course of preprocessing operations.

Depending on model requirements, some feature engineering may be needed such as encoding, scaling, discretization, and aggregation.

B. Logistic Regression for Binary Decisions
The most interpretable approach to deal with binary decisions is the application of logistic regression. The latter is helpful whenever the system has to predict the occurrence of an event, namely whether a particular record is considered regular or exceptional, approved or reviewed, low- or high-risk.

The logistic regression model looks as follows:

$$P(Y = 1) = 1 / (1 + e^{-(\beta 0+\beta 1x1+\beta 2x2+\cdots+\beta nxn)}) \quad (1)$$

The key advantage of the model is its clarity of interpretation.

Coefficients give the researcher insights on the direction of impact, making the decision easier to explain and justify.

Table I
Functional Distinction Between Ai And Agent Layers

Aspect	AI Layer	Agent Layer	Why It Matters
Primary function	Interprets data and generates intelligence	Executes actions based on conditions and rules	Preve nts role confu - sion
Nature of output	Probabilistic, ranked, scored, or classified re- sults	Deterministic workflow actions	Support s explaina bil- ity
Data depende	Requires historical and feature-rich	Requires processed signals	Impr oves

ncy	data	or triggers	modular de-sign
Typical methods	Regression, classification, clustering, anomaly detection	Rule engine, event handling, routing, escala-tion	Matches tool to task
Best process stage	Early analytical stages	Final response and execution stages	Enables process-aware architecture
Business value	Better insight and prioritization	Faster and more consistent action	Connects thinking to doing

C. Random Forest for Classification
In case of classification problems with several levels, random forest algorithm becomes relevant due to its reliability and ability to work well with tabular structured data, outperforming individual trees on the test data. It may be employed for the following purposes:

- risk tiering;
- priority classification;
- performance ranking;
- categorical sensitivity identification. A possible setting is as follows:
- number of estimators = 100;
- maximum depth = 10;
- criterion = gini.

D. K-Means for Segmentation
Not all enterprise problems are supervised. Whenever the task is to identify naturally forming clusters in the dataset, the k-means algorithm comes handy. It can be formulated by means of the following objective function:
$J = \Sigma(k)\ \Sigma(x \in Ci)\ \|x - \mu i\|^2$ (2)
Clustering can help to identify:

- high- vs low-value entity clusters;
- dense vs sparse structures;
- behavioral clusters valuable for threshold creation.

E. Anomaly Detection
Unusual cases often occur in operational datasets. However, not all deviations constitute errors – at least in terms of being worth special consideration. For interpretable detection of anomalies, the use of z-score is justified.

$Z = (x - \mu) / \sigma$ (3)

Large magnitude values indicate significant departure from normal distribution characteristics. Records that qualify for this criterion may be sent to review or flagged for agent handling.

F. Feature Importance and Interpretability
In addition to identifying predictions, a critical aspect of enterprise AI is interpreting model behavior. Tree-based algorithms generate important feature importance indicators. They may reveal:

- what drives high-risk or high-priority classification;
- whether value-related features prevail over category-related;
- changes in process behavior over time.

This is especially important whenever the output of a particular prediction is likely to trigger further actions.

G. Representative AI Outputs
Among other things, the AI layer is capable of producing the following outputs:

- probability score;
- class label;
- anomaly status;
- cluster affiliation;
- confidence indicator;
- leading features contributing to the prediction.

XII. AGENT LAYER DESIGN

The agent layer has to be responsible for implementing analytical results as behaviors of enterprises. The responsibility of this layer is different from the one of the AI layer, which produces intelligent analytics but not actions or behaviors.

A. Perception-Decision-Action Structure

It is designed that this layer is composed of three modules: Perception Module

This module accepts inputs from the AI layer, which include labels, scores, anomalies or rankings.

Decision Engine

This module employs business rules such as the following:

- if risk score is larger than a certain threshold, escalate;
- if priority is high and there is an anomaly, alarm immediately;
- if classification has low sensitivity, go through standard workflow path;
- if clustering results belong to the exceptional segments, human intervention will be required.

Action Module

This module performs actions such as the following:

- changing statuses;
- sending notifications;
- approving requests;
- escalating to supervisory stages;
- annotating dashboards;
- branching processes.

B. Types of Agents

There are many types of agents in the framework we propose, which include:

- reactive agents for responding to events immediately;
- rule-based agents that act on policies in deterministic way;
- goal-based agents for optimization purpose.

C. Why Agents Have to Go to Later Stages

A major rationale for our proposal is that agents should go to later stages since:

- later stages need actions, not interpretations;
- decisions have to become workflow behavior;
- it is easier to achieve operational consistency if actions are defined;
- event-driven systems avoid manual monitoring delays.

XIII. WORKFLOW MAPPING ACROSS VARIOUS ENTERPRISE APPLICATIONS

The following manner can be used to enhance the pragmatic relevance of the presented framework by demonstrating how similar mapping of architecture can be carried out across various enterprise processes independently of proprietary organizational information. As demonstrated, the presented model is not limited to a particular scenario, but rather operates as an enterprise decision system design pattern.

XIV. PROCESS MAPPING LOGIC

This model does not target any specific application area but focuses on the process aspect. It may be mapped to multiple workflow types in an enterprise context.

A. Approval-Based Processes

For processes that involve a lot of approvals, the AI module could predict the approval likelihood, detect anomalous patterns, and classify cases based on risk and importance. The agent module could determine if the case would be routed directly, escalated for higher-level approval, rejected because of failure to meet thresholds, or sent for special scrutiny.

B. Case Management

For workflow processes centered around the management of cases, the AI module could aid in setting priorities, segmentation, or detecting anomalies, while the agent module could update case ownership and the level of response required.

XV. AN ALGORITHMIC PERSPECTIVE ON THE FRAMEWORK

The framework logic can be expressed in algorithmic form as follows:

If critical condition (as defined in rules and thresholds) is detected for AI outputs then escalate and notify else If review condition is detected then put in review queue else Perform standard workflow end process state update, write decision log This logic makes it clear that intelligence and execution are separate processes.

XVI. EXAMPLE SCENARIO WITH MULTIPLE STEPS

In order to provide a better explanation of the above logic, we can use a more detailed story. Imagine some general business

International Journal of Scientific Research & Engineering Trends
Volume 12, Issue 2, Mar-Apr-2026, ISSN (Online): 2395-566X

process, involving a record entering the monitoring and approval system.

Firstly, the record is extracted from the source system and added to the structured data set. Then, during data preprocessing, its fields are cleaned up and transformed as needed, labels fixed if incomplete, additional derived columns created. At this point, the record becomes a part of the dashboard layer and its metrics/KPIs at a categorical level receive updated values.

Then comes the analysis done by the AI layer. The logistic regression predicts with high probability that the record is high priority. Similarly, the random forest classifier confirms it as a high-priority record. Furthermore, an anomaly detector rule says that one of the values in the record falls outside of the historical distribution significantly. Up to this point, nothing specific has been decided yet, but a significant decision context has been generated.

Finally, the output bundle of decisions is passed to the agent decision engine. It checks whether certain criteria are met. The current situation suggests that the agent will disallow default processing of this record since both conditions (high priority, anomalous) hold true. As a result, the following multi-part response is generated: escalation and notification of the appropriate user/queue, setting the record to review mode,

annotating the dashboard, changing the state of the record.
Clearly, this scenario proves the need for both layers in a monitoring application – without the agent, no operational action would happen, while a dashboard alone could not have triggered that action since it is only used for showing results of analyses.

Table II
Illustrative Workflow Mapping Of Ai And Agent Roles

Workflow Type	AI Layer Role	Agent Layer Role	Expected Benefit
Approval workflow	Predict approval likelihood, classify risk, identify anomalies	Route for direct approval, escalation, or review	Faster decisions with better consis- tency
Case handling workflow	Prioritize cases, detect exceptional patterns, classify urgency	Assign owner, notify authority, update case state	Better triage and reduced backlog
Monitoring workflow	Detect threshold breaches, segment operational patterns, forecast devia- tion	Trigger alerts, log incident, launch in- tervention path	Early response to operational changes
Review workflow	Score records for importance, iden- tify outliers, rank attention areas	Send review tasks, route to specialist queue, mark status	More focused analyst effort
Performance oversight workflow	Compare categories, infer weak or strong segments, identify trends	Notify managers, initiate corrective workflow, schedule follow-up	Better governance and continuous improvement

XVII. CRITERIA FOR TECHNICAL EVALUATION

The evaluation must take into account both the quality of the AI component and that of the agent component. Taking the entire system in isolation will result in concealing critical weaknesses.

A. AI Component Level Indicators
For classification problems, some common indicators include:

- accuracy;
- precision;
- recall;
- F1-score.

RMSE may be an indicator for regression-based forecasting. For clustering, there is the silhouette score. Anomaly detection may need precision at highest alerts.

B. Agent Component Level Indicators
Since agents perform operational functions, the following criteria should be used to evaluate their performance:

- decision latency;
- routing reliability;
- alert accuracy;
- proportion of automation results;
- minimization of manual labor.

International Journal of Scientific Research & Engineering Trends
Volume 12, Issue 2, Mar-Apr-2026, ISSN (Online): 2395-566X

C. System-Level Indicators

Some of the criteria that will enable assessment of the holistic framework from a business point of view include:

- improved speed of response;
- fewer manual repetitions;
- improved consistency in decision implementation;
- enhanced quality of prioritization;
- enhanced traceability from dashboards to actions. See Table III below

Table III
Suggested Evaluation Metrics For The Hybrid Framework

Layer	Metric	Purpose	Interpretation
AI layer	Accuracy, precision, recall, F1-score	Measure classification quality	Higher values indicate Stronger predictive separation
AI layer	RMSE	Evaluate regression-based estimates	Lower value indicates bet- ter error control
AI layer	Silhouette score	Measure cluster separation quality	Higher value indicates bet- ter cluster compactness and separation
AI layer	Feature importance stability	Check interpretability robustness	Stable drivers improve trust
Agent layer	Decision latency	Measure response speed after AI output	Lower delay improves pro- cess responsiveness
Agent layer	Automation rate	Measure proportion of actions completed without manual intervention	Higher rate indicates Stronger operational efficiency
Agent layer	Rule consistency	Measure whether identical signals trigger identical responses	Higher consistency improves governance
System level	Manual effort reduction	Compare effort before and after frame- work adoption	Indicates productivity gain
System level	End-to-end resolution time	Measure complete process performance	Reflects practical enterprise impact

XVIII. GOVERNANCE, TRUST, AND CONTROL CONSIDERATIONS

Enterprise decision systems should not only possess a certain level of technological proficiency, but they should also have governability. This consideration is particularly critical when AI results influence operations.

A. Model Governance

Each predictive model should have clear documentation regarding:

- input variables,
- pre-processing assumptions,
- range of possible outcomes,
- circumstances of re-training,
- potential limitations.

The above documentation prevents the risk of treating predictive models as black-box technologies.

B. Rule Governance

Apart from models, the agent layer also needs to be governed. In other words, rules need to be version-controlled, reviewed, and consistent with enterprise policies. For instance, if the risk threshold is altered from 0.80 to 0.70, the cause for this adjustment needs to be clearly identified, as it can significantly impact the behavior of workflow.

C. Audit Trail

A good implementation preserves the line of reasoning that begins with dashboard state and ends with agent behavior. Such audit trail can be helpful for purposes of:

- audit,
- management explanation,
- troubleshooting,
- continuous improvement.

D. Human Override

It is crucial to implement means of overriding (or deferring) certain actions triggered by the system in specific cases. An appropriately designed hybrid solution helps humans work effectively; it does not tie the organization to excessive automation.

XIX. SCENARIOS FOR ILLUSTRATING THE PROCESS

First, three exemplary process scenarios are described below.

A. Scenario 1: Sensitivity-driven Approvals

The record arrives in the system and is presented via the dashboard layer. The AI layer calculates the probability of risk at 0.86 and tags the record as "high risk." Based on this information, the agent layer triggers a rule which states that if probability > 0.80 and label = high risk, then trigger advanced review and alert authority. The system thus goes beyond mere awareness to actual action.

B. Scenario 2: Detection of Outliers as Triggers for Advanced Review

An input value stands out because it seems far removed from any historical trends or averages. A flag is raised based on the application of Z-score rule in the AI layer. However, rather than immediately rejecting the value, the agent layer prevents standard flow and sets up a review checkpoint in an example of balanced human intervention.

C. Scenario 3: Segmentation and Cluster-based Process Routing

A clustering model detects records falling within a sensitive historical cluster. Then, based on the results obtained, the agent layer routes the records for stricter processing procedures.

XX. FACTORS IN IMPLEMENTATION

Some key factors must be considered in implementing the suggested approach in an actual enterprise setting.

A. Data Governance

The performance of the artificial intelligence layer depends on the quality of data received by the system. Proper data governance should include:

- ensuring source integrity;
- standardizing data fields;
- controlling data transformation processes;
- managing versions of calculated columns and measures.

B. Rule Governance

Since the agent layer drives operational activities, it is crucial to keep the definitions of rules documented and updatable. Thresholds should not be implemented without careful consideration. Proper rule governance will involve:

- scheduling regular reviews of thresholds;
- specifying exception-handling procedures;
- recording audit trails for actions taken;
- determining override mechanisms if necessary.

C. Human-in-the-Loop Approach

While some activities can proceed through complete automation, other situations call for semiautomation. This means that although the framework will suggest and route activities, final approval must come from a human decision-maker.

D. Scalability

The architecture should also accommodate:

- model retraining;
- rule refinement;
- new types of workflows;
- enterprise-level integration with ERP, CRM, or case management systems.

XXI. DISCUSSION

It might be assumed that dashboards, AI and automation can be considered separate technology stacks which can be just deployed sequentially, but the truth is that their performance is based on integration logic. This solution takes care of it in a very straightforward manner – building up from data foundation to dashboard visualization, AI analysis and agent execution.Interpretability is one of the biggest advantages of the suggested architecture. First, the dashboard shows the visible state of the system. Secondly, the AI analysis provides explanations of this visible state through scores, classification or segmentation of the data. Thirdly, the agent layer gives an understanding of what actions should follow this analysis. Such approach greatly improves the transparency that may be missing from the "intelligent automation" solutions.Second advantage is practical adaptability. The presented architecture does not rely on any specific data set in any domain and can be applied across different enterprise workflows like approvals, reviews, monitoring, or case-based decision making. The focus is on the process rather than on a particular organization or tooling. This makes the architecture highly educational, suitable not only for practical implementation but also for academic discussion. Each layer of the architecture can be explained separately, yet they all form the full coherent system.Another important aspect of the proposed architecture is step-by-step adoption. Organizational units that have already implemented dashboards don't have to change the whole stack at once but rather gradually adopt new technology solutions: first, enhance the dashboard with AI insights in the form of score or anomaly flags and, finally, introduce the agent layer for selected processes.

XXII. LIMITATIONS

This research is theoretical and architectural in nature, rather than being an implementation case study focused on benchmarks. The limitations of the paper are:

- lack of a public enterprise data set which completely captures all possible process situations;
- little attention to unstructured data and deep learning approaches;
- the possibility that the domain dependence may affect the use of threshold-based rules;
- necessity of iterative tuning in practice;
- organizational reluctance due to excessive automation.

XXIII. FUTURE SCOPE

Future scope of research includes but not limited to:

- decision pipeline based on streaming data processing in real time;
- using reinforcement learning for adaptively changing agents' behavior;
- generating explanations in the natural language format;
- use of large language models for contextual support;
- use of graphs for reasoning about complex relationships in enterprise;
- evaluation against publicly available benchmark datasets.

One interesting area of research is hybrid governance, whereby humans, AIs, and agents have well-defined responsibilities depending on how sensitive the decision is. Another possible line of research involves comparing different architectures for

Integrating AIs and agents. For example, a researcher can investigate whether putting light-weight agents closer to the beginning of the process results in greater responsiveness of the system or certain classes of workflows require feedback loops between the AI layer and dashboard layer.

XXIV. CONCLUSION

In this paper, an innovative framework was presented for optimizing processes within enterprise decision systems using AI agents. The framework starts with enterprise data collection and continues with data cleaning, transformation, modeling, and dashboarding. Then two types of intelligence layers are added. First, there is an intelligence layer that covers the prediction, classification, clustering, anomaly detection, and explainability functions. Second, there is an agent layer that covers the rule validation, routing, escalation, alerting, and action functions.

The key innovation of this research lies in the clear division of intelligence creation and action execution stages. Instead of fusing AI and automation into a single system, this research assigns these tasks to different stages according to their technical nature. The result is a more straightforward, understandable, manageable, scalable, and evaluative framework. While enterprises continue to adopt dashboards and analytics, it is natural that they consider going beyond these tools. This research offered such an approach. The proposed framework can help turn dashboards from mere information displays into decision-making platforms.

Acknowledgment

The author expresses heartfelt appreciation to those who guided him academically and professionally through internships where he gained practical experience in enterprise analytics, dashboarding, business process analysis, and intelligent systems. These experiences were integral in developing the thoughts discussed herein.

REFERENCES

1. S. Russell and P. Norvig, Artificial Intelligence: A Modern Approach, 4th ed. Hoboken, NJ, USA: Pearson, 2021.
2. C. M. Bishop, Pattern Recognition and Machine Learning. New York, NY, USA: Springer, 2006.
3. S. Few, Information Dashboard Design: Displaying Data for At-a-Glance Monitoring, 2nd ed. Burlingame, CA, USA: Analytics Press, 2013.
4. T. Hastie, R. Tibshirani, and J. Friedman, The Elements of Statistical Learning, 2nd ed. New York, NY, USA: Springer, 2009.
5. L. Breiman, "Random forests," Machine Learning, vol. 45, no. 1, pp. 5–
32, 2001.
6. J. MacQueen, "Some methods for classification and analysis of multi- variate observations," in Proc. 5th Berkeley Symp. Math. Statist. Prob., Berkeley, CA, USA, 1967, pp. 281–297.
7. M. Wooldridge, An Introduction to MultiAgent Systems, 2nd ed. Hobo-ken, NJ, USA: Wiley, 2009.
8. E. Turban, R. Sharda, and D. Delen, Decision Support and Business Intelligence Systems, 10th ed. Boston, MA, USA: Pearson, 2011.
9. I. Goodfellow, Y. Bengio, and A. Courville, Deep Learning. Cambridge, MA, USA: MIT Press, 2016.
10. T. M. Mitchell, Machine Learning. New York, NY, USA: McGraw-Hill, 1997.