International Journal of Web of Multidisciplinary Studies
(Peer-Reviewed, Open Access, Fully Refereed International Journal)
website: *https://ijwos.com*
Volume 03, Issue 11,November 2026
OPEN ACCESS
E-ISSN : 3049-2424
DOI: https://doi.org/10.71366/ijwos

AIBDO-X: AI Powered Business Decision Optimizer Executive Tool

Mukesh Kolhe[1], Amrapali Natkar[2], Gautam Chinde[3], Tanuja Rajput[4]
[1,2,3,4] *Department of Artificial Intelligence & Data Science, Pune, India*

Article Info

Article History:
Published: 02 Nov 2026

Publication Issue:
Volume 03, Issue 11
November 2026

Page Number:
1-10

Corresponding Author:
Mukesh Kolhe

Abstract:
The traditional business intelligence (BI) tools and products, like Tableau, Microsoft Power BI, Qlik Sense, Google Looker Studio, SAP BusinessObjects or IBM Cognos, are very much geared towards descriptive charting and historical KPI tracking. These tools indicate the events that occurred and when they occurred but they do not provide the "why" or the suggestions for management action that could correct a metric's decline. Root cause identification, inter-departmental bottlenecks identification, and planning corrective interventions remain completely manual. In real life situations, enterprise teams also deal with separated workflows where they jump between separate data processing, predictive modelling, time-series forecasting and executive reporting solution software. To close this, we are introducing the AI Powered Business Decision Optimizer Executive Tool (AIBDO-X). This zero cost Decision Intelligence platform, developed with a FastAPI back-end, orchestrates 22 special purpose AI agents via over 30 modular Python files. The sites include user uploaded raw CSV or Excel files from 7 functional areas: Marketing, Finance, Sales, Human Resources (HR), Supply Chain, Customer Operations, and Risk. At Anasys, AIBDO-X implements a ten-step analytical workflow for each department running in nearly 50 seconds. This automated pipeline performs schema validation, missing data imputation, exploratory data analysis (EDA), domain-specific feature engineering, and competitive benchmarking over over a dozen machine learning estimators and multi-horizon forecasting with models such as ARIMA, Prophet and Holt-Winters. The system does not generate black box predictions, but instead it combines the capabilities of a Root Cause Agent (Root Cause Analysis, feature importance using SHAP, LIME and statistical correlation analysis) and an Explainability Agent that transforms the understandability of feature attributions into plain language business insights. Seven domain-specific Strategic Agents then evaluate whether each of these departments' operations is feasible, then a central Decision Agent integrates them all into one executive action plan. All past choices are recorded in a persistent Memory Agent to track events in the institution. Last but not least, a fault-tolerant LLM Manager ensures requests are sent to multiple API keys for Google's Gemini 3.5 Flash and Groq GPT-OSS-120B APIs - even if rate limits, it proves to be deterministic LLM rules act which also enables uninterrupted, zero-cost operations.

Keywords: Automated Machine Learning (AutoML), Business Decision Support Systems, Explainable AI (XAI), Large Language Models, Multi-Agent Architecture, Predictive Business Intelligence, Root Cause Attribution, Time-Series Forecasting.

1. Introduction

Companies nowadays generate enormous amounts of operational data through their sales interactions, marketing endeavors, financial transactions, staff lists, procurement processes, customer support queues, and risk management logs. Despite the ubiquitous use of commercial Business Intelligence (BI) platforms such as Microsoft Power BI, Tableau, Google Looker Studio, Qlik Sense, SAP BusinessObjects, IBM Cognos, ThoughtSpot, and Microsoft Fabric with Copilot, decision-making is still frequently stuck on the stage of manual interpretation. The backbone of standard

Copyright © 2026 The Author(s) : This is an open access article under the CC BY license (http://creativecommons.org/licenses/by-nc/4.0/)

BI software is used for descriptive reporting. It produces interactive visuals, KPI scorecards and drill down filters indicating what has happened (such as an 18% revenue decline in a month) but does not explain why it has occurred or recommend a specific next course for action for leadership (next steps would include raising awareness about the breach, a discussion of the drivers behind it, what actions to take next, or what changes to make in the cause) [1], [2].

Each time you use descriptive dashboards you experience one of the following operational challenges: In the simplest capacity, visual reports will let the manager know that a metric is failing, without letting them know the root cause of the failure, giving them no opportunity for automated decision support. The collection of charts and departmental spreadsheets requires cross-referencing of the data to determine if stock levels need to be topped up, decision makers are required to adjust unit prices, run region promotions or reallocate ad spend, which is a manual process. This is complicated even more in a small or medium scale organisation, as you are unlikely to find a dedicated data scientist to pinpoint which KPI is really impacting profitability or to establish whether a regional movement over time is outside the normal range.

Even if organizations implement predictive modules, black-box forecasts that predict a decline in revenue or profits do not explain the reasons behind the numbers, and can reduce faith of the manager [4], [9]. Forecasting engines also are used separately, without an explanation of the diagnostic process and concrete steps [7]. In the meantime, static dashboards require analysts to spend hours manually building slide decks and adding narrative summaries. In addition to reporting delays, the entire analytics lifecycle is still siloed, with a multitude of different applications for data scrubbing, visualization, forecasting, AI-enabled interrogation, and report generation. What's most important, though, is that conversational query is still very limited: Most platforms depend on navigating through a clunky multi-tab interface or manual SQL styling; and therefore, natural-language query Copilots are still locked behind premium enterprise-licensing interfaces [11] and [12].

To fill the gap between shady business decisions in departmental spreadsheet and clear decisions by executives, we created AIBDO-X (AI Powered Business Decision Optimizer Executive Tool). AIBDO-X provides a new framework that leverages a FastAPI and Uvicorn server that is based on Python 3.11+. With a multi-agent structure, more than 30 modular Python scripts and 22 specialized AI agents replace the analyst-dependent BI workflows. Raw CSV or Excel files can be uploaded by users to up to seven functional divisions, which start a ten-stage automated pipeline linked to a central Orchestrator Agent that is like a virtual CEO. Early stages deal with schema and value validation, missing data imputation with outlier removal, automatic exploratory data analysis (EDA) and automatic generation of charts, and domain-specific feature engineering. It then conducts competitive learning of over 10 machine learning estimators as well as multi-method time-series forecasting pipelines. It then uses SHAP and LIME root-cause diagnostics, makes tangible business recommendations, creates text-based explanations with Google's Gemini and performs strategic analyses at the department level within Groq. Once teams in all the departments are done, a central Decision Agent balances, reviews and harmonizes the departmental team priorities and generates a unified executive action plan backed by a continuous presence of the Memory Agent and a natural-language Conversational BI helper.

This paper makes four primary technical contributions:

Zero-Preprocessing Automated Data Pipeline: We create a basic layer of data (file_handler.py, validation_agent.py, cleaning_agent.py, eda_agent.py, feature_agent.py) which can accept raw files uploaded from CSV and Excel, clean data anomalies, and automatically compute domain information fields such as Profit Margin, Customer Lifetime Value (CLV), Growth Rate, Sales Ratio, Inventory Turnover and Recency-Frequency-Monetary (RFM) segments [2], [5], [6].

Unlike traditional forecasts that are treated as a single product, AIBDO-X connects benchmarks contained in analytics_agent.py (for 10+ classification/regression models) directly to Direct to root cause and recommendation models in root cause_agent.py and recommendation_agent.py, directly to the forecast in Prophet, ARIMA and Holt-Winters in forecast_agent.py. This coupling is done by pairing attributions of SHAP and LIME with quantified operational interventions [4], [7]–[9].

Many business departments are simulated and conflicts are resolved cross-departmentally, we model seven main business departments (Marketing, Finance, Sales, HR, Supply Chain, Customer, Risk) as independent strategic agents. decision_agent.py evaluates and ranks the competing proposals from the departments and arbitrates the final decisions, while memory_agent.py keeps a log of the decisions from the previous sessions [1–3, 13].

We are writing the Heavy Executive Synthesis block with three google Gemini 3.5 API keys, with priority failover, and the lesser Departmental Reasoning and Chat API calls with two Groq (openai/gpt-oss-120b) API keys, with round-robin load balancing. Risk-free deterministic rule-based fallback ensures continuous operation for no cost [10, 14].

2. Literature Review

AIBDO-X is inspired by recent advances in four closely related fields; autonomous agent orchestration, automated tabular data science, explainable forecasting and machine learning, and conversational business intelligence.

A. Beginnings of autonomous multi-agent architectures and orchestration moves away from large language model (LLM) prompt-based workflows and step through the beginnings of structured coordination. Ferrag et al. [1] describe this transition and make a formalization of such "Perception – Memory – Action" cycle in the context of core agentic patterns including task planning, tool invocation, self-reflection and multi-agent collaboration. They show that in multi-domain tasks, single monolithic LLM prompts often fail because of context limits or hallucinations unless assisted with explicit intent routing or a persistent episodic memory [1]. Following these structural ideas, a recent survey of orchestration [14] classified multi-agent topologies into three different types: decentralized, centralized, and hierarchical. This research also provided results of the hierarchical manager-worker topologies in which a managerial agent decomposes complicated workflows into specialized sub-agents and recomposes the structured outputs of these sub-agents to produce the output of the top agent, providing better runtime control and reduced tokens.

A number of agent collaborations in financial and recommendation domains are also emerging, but are role specialized. Xiao et al. [13] crafted TradingAgents as a framework simulating a trading company that integrates financial market analysis, homogenizes the decision-making process through its division into seven personas—such as analyst, researcher, trader, and risk management—and separates the analytical and executive functions into distinct tiers of the LLM. While TradingAgents [13] demonstrates the use of multi-role debate and Iterated Model Routing it is restricted to stock market tickers and not multi-departmental company spreadsheets. At the same period, Yang and Yin [3] proposed the innovative CG-MAR system to improve multi-agent recommendations deals with the problem of candidate retrieval, coarse ranking, pairwise comparative refinement, verification and a dedicated memory module. The pairwise ranking and persistent memory approach from [3] and the multi-role governance model from [13] are adapted to create AIBDO-X for optimizing enterprise decisions from multiple departments.

B. AI-Assisted Data SystemsGiving a footnote to the enterprise tables and to the executive KPIs is still a major engineering challenge. As Pang pointed out in the 2017 report [2] on traditional BI pipelines, they require substantial manual ETL scripting and advanced SQL knowledge to ensure that raw data matches the business metrics. To remove this manual burden, two business-semantic data blueprints were proposed in [2] which would pair a pair of Validation and Generation Agents with consensus checks to automate the data preparation task. Automated tabular data science involves a series of analytical tasks that usually include data exploration, anomaly fixing, feature generation, and model fitting, all of which can be dynamically assigned to data in automated end-to-end analytical workflows as demonstrated by Hong et al. [5] using a tool called Data Interpreter, which is implemented as a hierarchical task graph. On the contrary, Li et al. [6] take a phase-driven approach in which they develop a finite machine learning toolkit, consisting of five specific agent phases which correspond to data reading, planning, coding, reviewing and summarizing, with AutoKaggle achieving an 85% valid submission rate on the machine learning competitive benchmarks for tabular data. AIBDO-X is bringing together key concepts of [2], [5] and [6] to make a deterministic four-phase business data preparation pipeline.

C. Standard BI platforms focus their charting efforts on forecasting.Forecasting is usually treated as a standalone charting exercise in standard BI platforms. But an exhaustive survey by Chen et al. [7] of 47 time-series architectures using LLM as their foundational building blocks found that agentic forecasting works best when the LLM is used to help control existing time-series models, such as ARIMA, and machine learning regressors, with validation and decision layers in between. Yeh et al. [8] argued in favor of this hybrid approach by proposing Data-Centric Agent for Time Series (DCATS) and showing that the AutoML agents studied are able to reduce forecasting errors by an average of 6% over a range of time horizons without hyperparameter tuning by humans.

But when companies try to get predictive accuracy with no mathematical transparency, they are not fully content with the companies they're in charge of. The authors in [4] proposed a soft computing pipeline using SHapley Additive exPlanations (SHAP) and Local Interpretable Model-agnostic Explanations (LIME) to explain the reasons behind why a customer churns. Continuing the discussion on how these metrics are delivered to end-users, Bilal et al. [9]

explored using LLMs for Explainable AI (XAI) and presented a two-stage hybrid model – extracts feature attributions using a math-based approach (e.g., SHAP scores, LIME scores) and converts these numerical values into business-readable text by the LLM. Roumeliotis et al. [10] demonstrated also in another work on business sentiment categorization that LLMs can be used as explainable ensemble aggregation over a set of base classifiers, providing structured and human-readable JSON decisions with comprehensive rationales, and fall-backing to deterministic confidence-based rules when parsing LLM output fails. Root_cause_agent.py, explainability_agent.py, decision_agent.py and llm_manager.py in our architecture implement this two-stage approach of attribution-to-narrative design.

D. Conversational Business Intelligence over Tabular Data which is what eliminates the need to apply complex SQL queries and click on filters. In Baidu, Lian et al. [11] built CoeusBI, an interactive LLM-based analytics platform, which is based on dedicated routing and view-generation agents to make it possible to conduct multi-turn conversational analytics without any manual semantic configuration. Similarly, Al Awad et al. [12] proposed Talk2Data, a tabular question answering system. By connecting an intent router (that splits in- and out-bound quantitative calculation requests into distinct flow of questions) to a light multi-turn memory to take out schema information and chat history and embody within the prompt, Talk2Data recouped to an accuracy of 95.8% to 97.9% on the number of queries that were answered correctly [12]. Combined with these principles, AIBDO-X applies automated intent classification (query_classifier.py) directly on top of single-department and consolidated executive reports and session-aware natural-language querying (conversational_bi.py).

3. Existing BI Limitations and Problem Formulation

When analyzed with Copilot, using Copilot-like capabilities on mainstream enterprise analytics applications (including Microsoft Power BI, Tableau, Google Looker Studio, Qlik Sense, SAP BusinessObjects, IBM Cognos, ThoughtSpot, and Microsoft Fabric) shows a distinct architectural pattern. These tools work great at ingesting data (whether spreadsheets, SQL databases, cloud warehouses or REST APIs) and the visual presentation component (interactive charts, KPI scorecards and slice-and-dice filters), but the real decision making power remains with human analysts. The following eight structural weaknesses are currently present in commercial tools:

1. Problem 1 (Descriptive-Only Reporting): Dashboards normally show past changes, e.g., an older user indicates a revenue slide of 18% on the month, but provide no insight into why the activity is declining, what further (operational) data element led to the change, or who in a leadership team should take responsibility for the resulting behavior.
2. Problem 2 (Absence of Prescriptive Decision Support): The analyst must manually cross-validate visual charts, compare individual departmental charts and make plans for recovery individually. Commercial BI tools cannot automatically make decisions on specific actions such as increasing safety levels, changing unit pricing, introducing localized pricing discounts, focusing on particular sales territories, etc.
3. Problem 3 (Steep Analytical Skill Barrier): One of the difficulties for smaller and medium-sized entities is the lack of their own data science team (Problem 3 – Steep Analytical Skill Barrier). Non-technical business owners find it hard to define what's most important to visualize, which KPI is controling net profitability or when a change in this metric indicates a statistically significant anomaly.
4. Problem 4 (Opaque Predictive Add-Ons): Commercial suites with predictive plugins that alert to shrinking revenues and/or margins are still opaque. Failure to reveal the essential feature attributions, weakens confidence in the forecast by an executive [4, 9].
5. Problem 5 (Siloed Time-Series Projections): Native forecasting tools produce standalone trends that are fully isolated from root cause analyses, prescriptive action planning processes and constraints on the supply side [7].
6. Problem 6 (Manual Reporting Overhead): The traditional dashboard just shows figures and simply static visuals which then require manual effort from managers to prepare narrative text summaries, executive briefs, and slide decks, which takes hours to perform.
7. Problem 7 (Fragmented Toolchains): Rarely, if ever, does end-to-end business intelligence exist in a single place, meaning that teams need to assemble separate pieces of various tools to build a dashboard, create time-series models, use machine learning, and develop narrative reports.

8. Problem 8 (Restricted Conversational Querying): Traditional BI platforms, whether they are literally arranged as one or two pages of filters, require end-users to click through multiple pages of filter panels, or to learn SQL and manually craft queries to extract insights from these platforms; though natural-language copilots are functional in limited applications or restricted to expensive enterprise licensing tiers [11, 12].

As listed in Table 1, our AIBDO-X architecture can be compared to typical commercial BI platforms in terms of its functional capabilities with respect to ten operational dimensions.

Table 1 compares the different functional capabilities of mainstream BI suites with the proposed AIBDO-X architecture.

Feature / Capability	**Mainstream BI Suites**	**Proposed AIBDO-X Platform**
Direct ingestion of uncleaned CSV/Excel files	No (Requires manual ETL)	Yes (Automated ingestion pipeline)
Automated anomaly repair & data cleaning	No	Yes (Validation + Cleaning Agents)
Domain-specific KPI & feature engineering	No	Yes (Profit Margin, CLV, RFM)
Competitive AutoML benchmarking (10+ models)	No	Yes (Selected via R^2 / Accuracy)
Multi-horizon time-series forecasting	Basic / Standalone	Yes (Prophet, ARIMA, Holt-Winters)
Explainable AI (SHAP/LIME) root-cause attribution	No / Opaque	Yes (Dedicated Root Cause Agent)
Cross-departmental multi-agent simulation	No	Yes (7 Dept. Agents + CEO Orchestrator)
Conflict-resolved executive action planning	No	Yes (Central Decision Agent)
Natural-language Conversational BI	Restricted / Paid Add-On	Yes (Groq LLM + Report Context)
End-to-end platform licensing cost	Paid Enterprise Tiers	Zero Cost (₹0 Operational Cost)

4. Proposed System Architecture and Multi-Agent Methodology

AIBDO-X brings the whole analytical life cycle into an end-to-end automated decision pipeline, instead of piecemeal BI tools and capabilities by analysts. This codebase is organised in 10 functional blocks in over 30 python scripts There are 11 RESTful services exposed via backend services using FastAPI and Uvicorn (api/app.py, api/routes.py, api/schemas.py, api/services.py, run_api.py). At the user end, the users interface with a 5-page HTML5, CSS3, Bootstrap 5 and Vanilla JavaScript (index.html, dashboard.html, reports.html, chat.html, usage.html) web app). Fig. 1 illustrates the complete system architecture and multi-agent data flow of AIBDO-X.

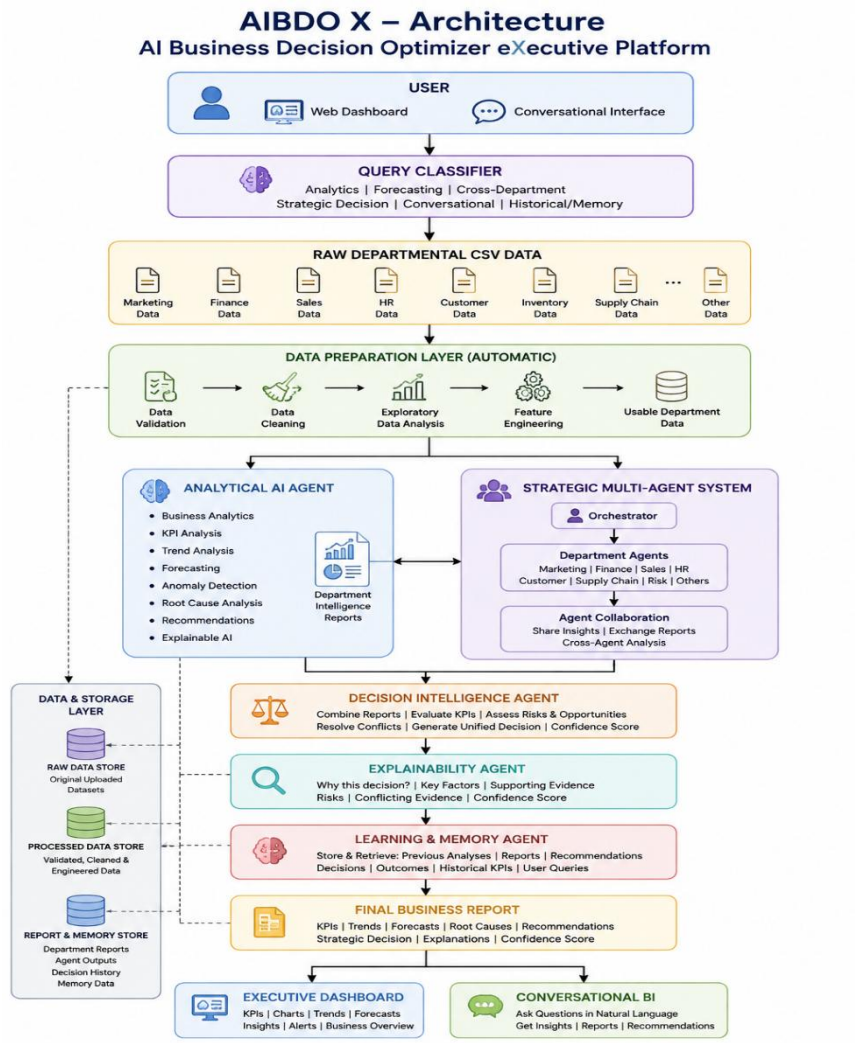


Fig. 1. Complete system architecture and multi-agent data flow of the AIBDO-X Executive Platform.

A. Part A: For every user who uploads up to 7 raw CSV or excel spreadsheets that have sales, marketing, finances, HR, supply chain, customer or risk data, the central Orchestrator Agent (agents/orchestrator_agent.py -- which inherits from the common LLM interface agents/base_agent.py) initiates a deterministic sequence of data preparation [2,5]:

- File Ingestion (data_pipeline/file_handler.py): Raw file uploads that include data in tabular format are first checked by the File Handler which checks the file extension, character encoding, and memory requirements before loading data into memory.
- Schema and Quality Auditing (data_pipeline/validation_agent.py): Then the Validation Agent conducts an audit over each table to identify tables lacking headers, mismatched column types, duplicate records, null percentages, malformed entries, entries of infinity (+ or – ∞) and extreme statistical deviations.
- Automated Anomaly Repair (data_pipeline/cleaning_agent.py): Any quality issues found are automatically fixed by the Cleaning Agent. It removes duplicate rows, normalizes the time stamps, truncates the extreme values based on interquartile range (IQR) thresholds and fills missing values dynamically. If there is a gap in an observer's

response, the imputation method compensates for that sample skewness, and if there is a gap in a categorical field then the mode is imputed:

$$\hat{x}_{\text{imputed}} = \begin{cases} \text{mean}(x), & \text{if } |\text{skew}(x)| \text{ is less than } 1 \\ \text{median}(x), & \text{if } |\text{skew}(x)| \geq 1 \\ \text{mode}(x), & \text{if } x \text{ is categorical} \end{cases} \quad (1)$$

- Exploratory Profiling (data_pipeline/eda_agent.py): The other multiperiod butts/dates features, like top/bottom performing product lines, bottom/coherent trend regions, seasonal cycles, and multiperiod pairwise correlations, are computed automatically by the Exploratory Profiling (data_pipeline/eda_agent.py) with the dataset cleaned.
- Domain KPI Construction (data_pipeline/feature_agent.py): Finally, a Feature Engineering Agent (data_pipeline/feature_agent.py) is able to extract high-value business metrics directly from the activated columns. The ones available are such as Profit Margin, Customer Lifetime Value (CLV), Growth Rate, Sales Ratio, Inventory Turnover, or Recency-Frequency-Monetary (RFM) customer tier [2, 6]:

$$\text{Profit Margin} = \frac{\text{Revenue} - \text{Cost}}{\text{Revenue}}, \quad \text{Inventory Turnover} = \frac{\text{COGS}}{\text{Avg_Inventory}} \quad (2)$$

B. Part B: Once the domain-specific KPIs have been added to the data, the execution moves to the predictive modeling and forecasting layer [6]–[8].

The Estimators that integrate with the Analytics Agent engine are available, for example, in the ml_engine/analytics_agent.py module, and include: Linear Regression, Ridge, Lasso, Decision Trees, Random Forests, Gradient Boosting, XGBoost, LightGBM, K-Nearest Neighbors (KNN), and Support Vector Regression (SVR). In regression problems: The best estimator, m^, is chosen so as to maximize R^2 (coefficient of determination); In classification problems: The best estimator, m^ is chosen so as to maximize predictive accuracy:

$$m^* = \arg\max_m \left[1 - \frac{\sum_i \left(y_i - \hat{y}_i^{(m)}\right)^2}{\sum_i (y_i - \bar{y})^2} \right] \quad (3)$$

The best model is saved as a .pkl artifact by using utils/model_manager.py. This manager tests the model freshness, feature length of input and baseline accuracy thresholds before serving live inferences on the POST /api/predict endpoint.

The Forecast Agent (ml_engine/forecast_agent.py) is used to perform multi-horizon time-series forecasting for all core metrics like Revenue, Sales Volume, Net Profit, Customer Demand and Stock Levels. It includes Facebook Prophet, statsmodels ARIMA, Holt-Winters Exponential Smoothing, and lag-feature machine learning regressors [7] and [8] to address a variety of seasonal and non-stationary behaviours.

C. Part C & Part F: AIBDO-X uses a three-step reasoning pipeline [4], [9] to overcome black-box prediction and siloed trend lines (Problem 1, 2, 4, 5), including all outputs of the models in this proposal and all time-series trajectories.

The first step is to determine the Root Cause Agent (reasoning/root_cause_agent.py), which determines the reason for historical metric change, or projection of future trajectory change. Combines statistical correlation, permutation feature importance, LIME and SHapley Additive exPlanations (SHAP) [4] into a single analysis. If a business gets a trained model f and some index j, then the marginal contribution ϕ_j of that business variable with respect to the remaining feature subset S⊆F\ {j} is actually calculated as:

$$\phi_j(f, x) = \sum_{S \subseteq F \setminus \{j\}} \frac{|S|!\,(|F| - |S| - 1)!}{|F|!} [f(S \cup \{j\}) - f(S)] \quad (4)$$

Second, the Recommendation Agent (reasoning/recommendation_agent.py) translates these prioritised feature drivers and forecast trajectories into numerical business interventions, such as transferring ad spends or changing unitisable prices, scheduling promotional discounts, reordering stock into the warehouse, etc., or initiating churn-prevention campaigns [3]. Third, the Explainability Agent (agents/explainability_agent.py) helps to convert the raw mathematical

weights, supporting metrics, operational risks and confidence intervals, provided by Google Gemini 3.5 Flash, into clear executive prose [9], [10].

D. Part E & Part F: A decision that impacts on one team typically impacts on many others, particularly in the case of D. Part E & Part F, Seven Strategic Department Agents, Decision Agent and Memory Enterprise decisions. To emulate the same process of cross-functional executive deliberation that takes place in [1] and [13], AIBDO-X sends seven Strategic Department Agents that are based on Groq (openai/gpt-oss-120b) as shown in Table 2.

Table 2: Functional Roles and Outputs of the Seven Strategic Department Agents

Strategic Agent (Module)	**Core Domain Scope**	**Primary Evaluation Output**
Marketing (marketing_agent.py)	Campaign performance, demand shifts, competitor activity	Demand Score & Marketing Plan
Finance (finance_agent.py)	Capital allocation, ROI projections, cash flow, margins	Financial Feasibility Rating
Sales (sales_agent.py)	Territory performance, pipeline velocity, target attainment	Sales Growth Potential
HR (hr_agent.py)	Workforce capacity, hiring needs, training readiness	Resource Readiness Index
Supply Chain (supply_chain_agent.py)	Warehouse stock, supplier lead times, fulfillment logistics	Operational Feasibility Score
Customer (customer_agent.py)	Retention rates, satisfaction scores, feedback sentiment	Customer Acceptance Score
Risk (risk_agent.py)	Regulatory compliance, market volatility, competitive threats	Composite Risk Level

In the real world of business, the major functions of the departments often compete for attention and resources; for example, while Marketing is keen to practice hard discounting to ensure required volume and Sales asks to do so to cover inventory levels, Supply Chain may be keeping an eye on inventory levels. The Decision Agent (agents/decision_agent.py) considers all seven departmental assessments, evaluates, weights and ranks them to resolve these trade-offs. It will punish stresses in the case of conflicts between departments or risks that are exposed, and it will choose one unified plan of executive action using a composite utility function U(a) [10]([13]:

$$U(a) = \sum_{d=1}^{7} w_d \, S_d(a) - \lambda_{\text{risk}} R_{\text{risk}}(a) - \lambda_{\text{conf}} \text{ConflictPenalty}(a) \quad (5)$$

After the optimal strategy is determined the Memory Agent (agents/memory_agent.py) stores the whole execution trace, departmental scores and final executive verdict in the directory: data/processed/reports/consolidated/. This

longitudinal ledger permits leadership to conduct follow-up audits on an along-term basis that can include monitoring done for either the implementation or fulfillment of previous recommendations, for instance, monitoring for implementation of previous recommendations [1] [3].

E. Part D: Query Classifier (nlp/query_classifier.py) takes incoming natural-language queries and sends them to the appropriate BI to get the results required for interactive exploration through /api/chat and chat.html. Within this module user questions are classified according to their intent, such as Analytics, Strategic or Reporting, and analytics information is specified in question temporal bounds, KPI names, and departmental targets are extracted from the user's question. The generated departmental summaries are then used to pull grounded context from the conversational BI engine (nlp/conversational_bi.py) and the historical memory archives. Executives can then get data-driven answers to their questions, like "Why did net profit go down?", "Did I need to restock on some SKUs quickly?" or "What should be our top priority as a leader?," all without writing any SQL or navigating through dashboard filters [11], [12].

5. Conclusion

Enterprise analytics is not just about individual dashboards, though they are a key component, but successfully moving to real business action. In this work, we propose an end-to-end multi-agent framework, called AI Powered Business Decision Optimizer Executive Tool (AIBDO-X), to automate the process from data found in department-level spreadsheets to executive-level recommendation report.This work introduces the AI Powered Business Decision Optimizer Executive Tool (AIBDO-X), which is an end-to-end multi-agent framework that can automate the journey from an input of raw data stored in departmental spreadsheets to an executive-level recommendation report. The platform combines all the tasks of data preparation, predictive modeling and strategic planning within a single automated workflow of ~50 seconds, rather than these activities being disjointed manual work. Unprocessed files in CSV format and averaged Excel files go into an automated schema validation, skewness-aware imputation, and domain KPI engineering before participating in a competitive AutoML tournament alongside more than ten estimators, including multi-horizon forecasting using ARIMA, Prophet and Holt-Winters. AIBDO-X guarantees each of its recommendations comes from a mathematical basis and factoring in the operational feasibility, by integrating the SHAP and LIME feature attribution with seven domain-specific Strategic Agents and a central conflict-resolving Decision Agent. Moreover, with the integration of a multi-key Gemini and Groq rotation manager, a persistent Memory Agent and a natural-language Conversational BI module, enterprises can obtain enterprise-class decision intelligence for free.

There are some extensions planned for future versions of the platform. The first is to make all of the FastAPI and multi-agent runtime into a collection of microservices that are containerized in Docker for easy-scale, cloud deployments. Second, it currently handles static tabular uploads, but future versions will include support for live database connectors to SQL and enterprise resource planning (ERP) data streams. Lastly, we intend to add autonomous reinforcement learning feedback loops in the Memory Agent, integrating and allowing the system to adapt weights for the different utilities of the departments according to the financial feedback from the real world.

References

1. M. A. Ferrag, N. Tihanyi, and M. Debbah, “From LLM reasoning to autonomous AI agents: A comprehensive review,” IEEE Access, vol. 14, pp. 1–38, Jun. 2026, doi: 10.1109/ACCESS.2026.3698694.
2. C. Pang, “Toward data systems that are business semantic centric and AI agents assisted,” IEEE Access, vol. 13, pp. 113751–113764, Jul. 2025, doi: 10.1109/ACCESS.2025.3583260.
3. X. Yang and J. Yin, “CG-MAR: LLM-based multi-agent recommendation system via community graph,” IEEE Access, vol. 14, pp. 106203–106218, Jul. 2026, doi: 10.1109/ACCESS.2026.3709706.
4. S.-V. Oprea and A. Bâra, “Explaining customer churn through unified perturbation-based XAI soft computing,” IEEE Access, vol. 14, pp. 24460–24476, Feb. 2026, doi: 10.1109/ACCESS.2026.3663648.
5. S. Hong, Y. Lin, B. Liu, B. Wu, C. Zhang, D. Li, J. Chen, J. Zhang, J. Wang, L. Zhang et al., “Data Interpreter: An LLM agent for data science,” in Findings of the Association for Computational Linguistics: ACL 2025, Jul. 2025, pp. 19796–19821, doi: 10.48550/arXiv.2402.18679.
6. Z. Li, Q. Zang, D. Ma, J. Guo, T. Zheng, M. Liu, X. Niu, Y. Wang, J. Yang, J. Liu et al., “AutoKaggle: A multi-agent framework for autonomous data science competitions,” arXiv preprint arXiv:2410.20424, Nov. 2024, doi: 10.48550/arXiv.2410.20424.

7. Y. Chen, X. Qin, C. Liu, L. Wu, N. I. Samia, and K. Ding, “LLM agents for time-series: A survey,” in Findings of the Association for Computational Linguistics: EMNLP 2026, Aug. 2026, arXiv:2608.26226, doi: 10.48550/arXiv.2608.26226.
8. C.-C. M. Yeh, V. Lai, U. S. Saini, X. Fan, Y. Fan, J. Wang, X. Dai, and Y. Zheng, “Empowering time series forecasting with LLM-agents,” arXiv preprint arXiv:2508.04231, Nov. 2025, doi: 10.48550/arXiv.2508.04231.
9. A. Bilal, D. Ebert, and B. Lin, “LLMs for explainable AI: A comprehensive survey,” arXiv preprint arXiv:2504.00125, Mar. 2025, doi: 10.48550/arXiv.2504.00125.
10. K. I. Roumeliotis, D. Margaris, D. Spiliotopoulos, and C. Vassilakis, “Large language models as explainable AI ensemble aggregators for business review sentiment analysis: A comparative study with classical ensembles,” Appl. Sci., vol. 16, no. 13, p. 6479, Jun. 2026, doi: 10.3390/app16136479.
11. J. Lian, C. Li, Y. Shao, M. Wang, Y. Dong, X. Liu, W. Zhang, C. Gui, T. Wan, and M. Dong, “CoeusBI: A comprehensive interactive business intelligence system powered by LLMs at Baidu,” Proc. VLDB Endow., vol. 19, pp. 4076–4089, Jun. 2026, doi: 10.14778/3827998.3828017.
12. M. N. Al Awad, S. Ivanov, O. Tikhonova, and I. Khodnenko, “A multimodal conversational agent for tabular data analysis,” arXiv preprint arXiv:2511.18405, Nov. 2025, doi: 10.48550/arXiv.2511.18405.
13. Y. Xiao, E. Sun, D. Luo, and W. Wang, “TradingAgents: Multi-agents LLM financial trading framework,” arXiv preprint arXiv:2412.20138, Jun. 2025, doi: 10.48550/arXiv.2412.20138.
14. Y. Zhu, L. Liu, J. Yu, and D. Zhang, “LLM-based multi-agent orchestration: A survey of frameworks, communication protocols, and emerging patterns,” Future Internet, vol. 18, no. 6, p. 326, Jun. 2026, doi: 10.3390/fi18060326.