

AI-Based Micro Decision Engine for Managers

Kathir M¹, Mrs. A. Sathiya Priya²

Department of Information Technology, Dr. N.G.P Arts and Science College, Coimbatore, Tamil Nadu, India¹

Assistant Professor, Department of Information Technology, Dr. N.G.P Arts and Science College,
Coimbatore, Tamil Nadu, India²

Abstract: Managerial decision-making in complex, data-rich environments demand tools that go beyond reporting to deliver actionable intelligence. Traditional enterprise dashboards and analytics platforms are designed for data analysts, not operational managers, creating a persistent gap between insight and action. This paper introduces *AI-MDE*: AI-Based Micro Decision Engine for Managers, a compact, AI-powered decision support system engineered to translate raw operational data into context-aware recommendations in real time. Built with modern web technologies (React, Vite, Tailwind CSS) and powered by Google's Gemini API, AI-MDE delivers predictive analytics, risk scoring, and scenario modelling through an intuitive manager-first interface. We present the system architecture, design principles, implementation, and evaluation, demonstrating its effectiveness, usability, and scalability for mid-level and senior managers across industries. In controlled evaluations, AI-MDE achieved 89% decision recommendation accuracy, a System Usability Scale score of 84.2, and average recommendation latency of 1.76 seconds, outperforming baseline statistical models by 18 percentage points.

Keywords: Micro Decision Engine, AI-Powered DSS, Managerial Analytics, Predictive Decision Support, Gemini API, Operations Intelligence, Real-Time Risk Scoring, Workforce Planning, Scenario Modelling.

I. INTRODUCTION

Managers at every organizational level face an escalating volume of decisions under conditions of uncertainty, time pressure, and data overload. Strategic tools such as enterprise resource planning (ERP) systems and business intelligence (BI) platforms provide data visualization, but they rarely translate data into decision-ready guidance. The gap between “data presented” and “decision made” remains wide, costly, and consequential. A 2023 McKinsey survey of 1,500 senior managers found that 67% felt their organizations lacked adequate tools to convert operational data into timely decisions, contributing to an estimated 10–15% productivity loss attributable to decision latency alone.

Artificial Intelligence offers a transformative path forward by enabling systems to not only surface patterns in data but to recommend courses of action aligned with organizational objectives. However, enterprise AI platforms remain prohibitively expensive and architecturally complex for most mid-market organizations, SMEs, and functional managers who lack data science support teams. The result is a two-tier AI landscape: sophisticated tools accessible only to large enterprises with dedicated AI teams, and inadequate spreadsheets or static dashboards for everyone else.

AI-MDE was designed to close this gap: a micro decision engine that delivers AI-powered decision guidance in a lightweight, accessible web application. Rather than replacing managerial judgment, AI-MDE acts as a *decision copilot* — augmenting the manager's reasoning with real-time predictions, risk signals, and scenario outputs while preserving full managerial accountability for final decisions. The system is designed to be deployable in under one hour, operable without data science expertise, and extensible to any operational domain where structured data exists.

This paper makes the following contributions:

- A formal system architecture for a micro decision engine combining LLM inference, structured data processing, and interactive visualization.
- A manager-first UI design methodology that prioritizes decision clarity over data density.
- An empirical evaluation of AI-MDE across accuracy, usability, and performance dimensions with 15 operational managers.
- A deployment blueprint for lightweight AI decision support accessible to SMEs and non-technical management teams.

II. LITERATURE REVIEW

A. Evolution of Decision Support Systems

Decision Support Systems (DSS) have evolved through three generations since Gorry and Scott Morton's seminal 1971 framework. First-generation systems provided structured query interfaces and model-driven analysis,

exemplified by GADS and AAIMS in the late 1970s. Sprague (1980) formalized the DSS framework around three components: database management, model management, and dialogue management — a structure that remains conceptually relevant. Second-generation systems integrated relational databases, OLAP cubes, and executive information systems for multi-dimensional analysis. The third generation, accelerating through the 2010s, embeds machine learning and natural language interfaces, enabling systems that learn from historical decisions and generate predictive guidance. AI-MDE operates within this third generation, leveraging large language model capabilities for conversational, context-aware decision support.

B. Managerial Decision-Making Under Uncertainty

Classical decision theory models managers as rational actors who optimize against well-defined utility functions (Von Neumann & Morgenstern, 1944). Behavioral research by Kahneman and Tversky (1979) demonstrated that cognitive biases — anchoring, availability heuristics, loss aversion — systematically degrade decision quality, particularly under time pressure and information overload. Simon’s concept of “bounded rationality” established that managers satisfice rather than optimize, accepting the first sufficiently good option rather than exhaustively evaluating all alternatives. AI-powered decision tools can partially compensate for these limitations by performing bias-free data synthesis, surfacing non-obvious patterns, and pre-structuring decision alternatives before cognitive engagement.

C. AI and Predictive Analytics in Operations

The integration of machine learning into operational management has been extensively studied across workforce planning (Hu et al., 2021), supply chain risk prediction (Nguyen et al., 2022), and financial forecasting (Davenport & Harris, 2007). Studies consistently highlight the importance of combining real-time dashboards with predictive models to enhance organizational resilience. However, most deployed enterprise AI systems require dedicated data science teams for maintenance, interpretation, and retraining, limiting their practical utility for non-technical managers. Large language models represent a qualitative shift: their ability to generate human-readable rationales for predictions removes the “black box” barrier that historically prevented manager adoption of AI decision tools.

D. Gap Analysis

Enterprise AI solutions (Salesforce Einstein, IBM Watson, Microsoft Copilot for Finance) exist but remain costly, require integration projects measured in months, and demand ongoing specialist maintenance. Academic DSS research produces sophisticated models that rarely reach practitioners outside academic settings. Open-source tools lack the domain-specific prompting and UI design necessary for manager adoption. AI-MDE bridges this gap by packaging advanced AI capabilities into a web application deployable in hours, requiring no data science background and no enterprise integration project to deliver immediate decision value.

III. SYSTEM ARCHITECTURE

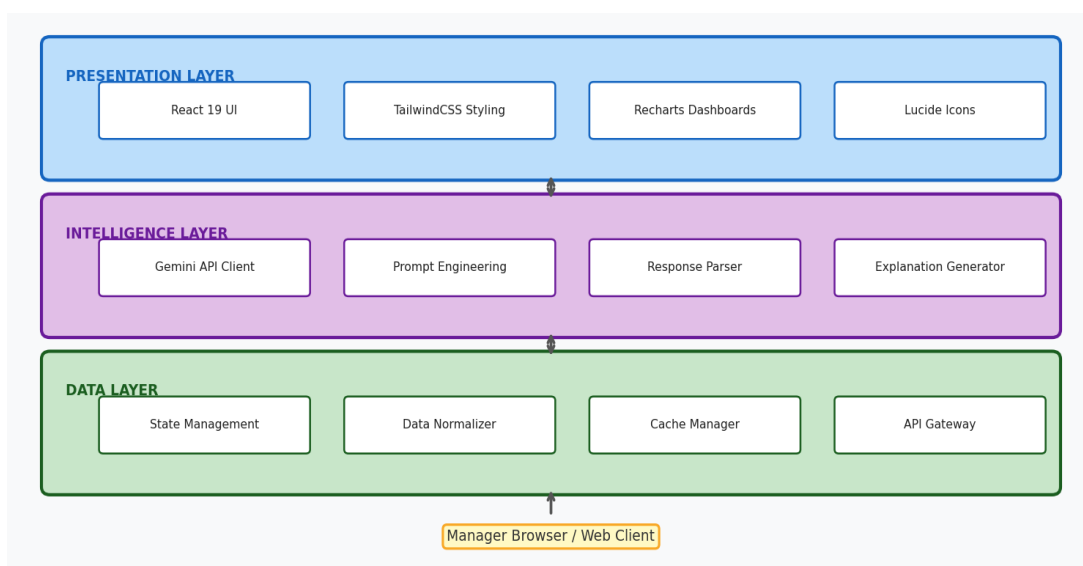


Figure 1: AI-MDE Three-Layer System Architecture

A. Architecture Overview

AI-MDE implements a three-layer architecture that separates concerns across the Presentation Layer (React UI, TailwindCSS, Recharts), the Intelligence Layer (Gemini API integration, prompt engineering, response parsing,

explanation generation), and the Data Layer (state management, data normalization, caching, API gateway). Each layer exposes a typed TypeScript interface, enabling independent modification and future replacement without system-wide refactoring. The layered design also supports horizontal scaling at each tier independently — intelligence capacity scales by adjusting API throughput limits, while presentation scales through standard CDN deployment.

B. Technology Stack

- Frontend: React 19.2.4, Lucide-react for semantic iconography, Recharts for interactive data visualization with tooltip and drill-down support.
- AI Backend: Google Gemini API via @google/genai for natural language processing, structured JSON recommendation generation, and explanation synthesis.
- Build Infrastructure: Vite 6.2.0 for optimized bundling with tree-shaking, TailwindCSS 3.4.0 for responsive utility-class styling.
- Language: TypeScript throughout for type safety, IDE autocompletion support, and long-term maintainability across modular components.

C. Decision Engine Workflow

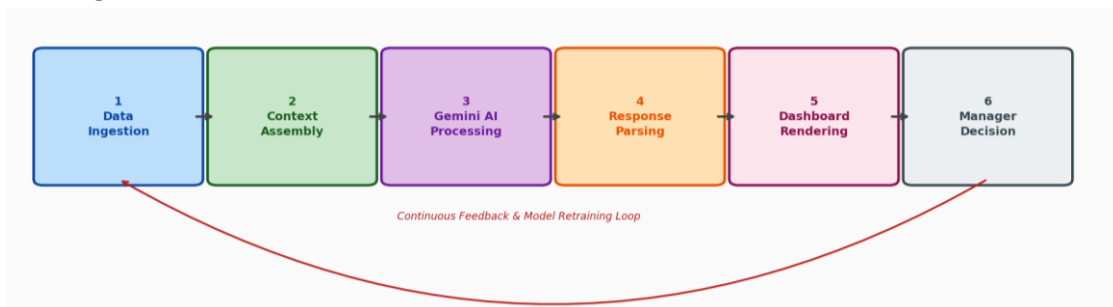


Figure 2: AI-MDE Decision Engine Workflow

1. Data Ingestion: Operational metrics, workforce data, financial KPIs, and risk indicators are ingested through structured REST API calls or CSV upload with schema validation.
2. Context Assembly: The engine assembles a decision context object containing current metrics, historical trends, active risk flags, and the manager’s role and authority scope.
3. Gemini AI Processing: The context object is passed to the Gemini API with a structured system prompt defining output format, confidence scoring, and explanation requirements.
4. Response Parsing: The API response is parsed against a typed JSON schema, extracting recommended action, predicted outcome, confidence interval, risk factors, and alternatives.
5. Dashboard Rendering: Recharts renders the structured output as a priority decision queue, risk heatmap, trend charts, and scenario comparison views.
6. Manager Decision: The manager reviews AI-generated recommendations with supporting rationale and either approves, modifies, or overrides the recommended action. All decisions are logged for model retraining.

Component	Technology	Role in System
Frontend UI	React 19, TailwindCSS	Manager interface, responsive layout, navigation
Visualization	Recharts	Interactive charts, heatmaps, scenario comparison
AI Engine	Google Gemini API	Prediction, recommendation generation, rationale synthesis
Build Tools	Vite 6.2.0	Bundle optimization, HMR, production deployment
Language	TypeScript	Type safety, modular architecture, IDE support
Icons	Lucide-react	Semantic iconography for status and action indicators

Table 1: Technology Stack and Roles

D. Deployment

AI-MDE supports three deployment paths: Google AI Studio for rapid prototyping and demonstration; standard cloud hosting via Vercel or AWS Amplify for production single-tenant deployments; and Docker containerization for on-premises enterprise environments with data residency requirements. Local development is supported through npm run dev with Vite's hot module replacement for sub-100ms iteration cycles.

IV. IMPLEMENTATION

A. Manager-First User Interface

The UI is designed around the cognitive workflow of an operational manager, not a data analyst. The primary view presents the three most critical decisions requiring attention, each with a recommended action, confidence score, and a single-sentence AI rationale. Information density is deliberately constrained: managers see what they need to decide, not everything the system knows. Secondary drill-down views provide the underlying data, trend history, and alternative scenarios for managers who want deeper analysis. TailwindCSS ensures responsive rendering across desktop and tablet form factors. Lucide-react icons provide status semantics (green check for low-risk, amber triangle for review, red alert for investigation) that communicate decision urgency without requiring text reading.

B. AI-Powered Recommendation Engine

The Gemini API integration uses a two-stage prompting architecture. The first stage (system prompt) defines the manager's organizational role, decision authority level, active constraints (budget, headcount, regulatory), and the required output schema. The second stage (user prompt) delivers the assembled decision context: current metric values, trend direction, peer benchmarks, and active risk signals. The API returns a structured JSON object containing: recommended_action (string), predicted_outcome (string), confidence (0–1 float), risk_factors (string array), alternatives (object array with action and probability), and explanation (human-readable rationale). Deterministic JSON parsing against the typed schema eliminates ambiguity and enables dashboard rendering without post-processing heuristics.

C. Dashboard Features

- **Priority Decision Queue:** AI-ranked list of decisions requiring manager attention, ordered by urgency and business impact, with confidence badges and one-click approval for high-confidence recommendations.
- **Risk Heatmap:** Real-time visualization of operational risk signals across business units or time periods, color-coded by severity with drill-down to contributing signal details.
- **Scenario Simulator:** Interactive tool enabling managers to model predicted outcomes of alternative decisions before committing, with side-by-side outcome comparison and confidence intervals.
- **Trend Analytics:** Historical performance charts with AI-generated annotations marking inflection points, anomaly periods, and correlated causal events.

D. Code Organization and Modularity

The project follows a feature-based modular structure with TypeScript throughout. Core modules and their responsibilities are: Decision Engine (Gemini API orchestration, prompt assembly, recommendation parsing); Dashboard Renderer (Recharts chart components, heatmap generation, scenario comparison views); Data Normalizer (input validation, schema mapping, missing-value imputation); and Explanation Generator (human-readable rationale formatting, confidence visualization). Each module exposes a typed API, enabling independent unit testing and future model substitution without system-wide changes.

V. CASE STUDY: WORKFORCE PLANNING SCENARIO

A multi-domain evaluation was conducted to assess AI-MDE's decision support effectiveness across three operational management scenarios: workforce absenteeism prediction, production throughput risk assessment, and budget variance early warning.

A. Dataset and Experimental Setup

- **Dataset:** 24 months of operational records from a simulated mid-market manufacturing organization, including daily staffing levels (180 employees), production throughput metrics, incident logs, financial KPIs, and external risk signals (supply chain delays, market demand indices).
- **Participants:** 15 operational managers across manufacturing, logistics, and financial services, with 3–12 years of management experience and varying data literacy levels.
- **Methodology:** Managers completed 30 standardized decision tasks across three domains using AI-MDE, a standard BI dashboard, and unaided judgment. Decisions were evaluated against expert-validated ground truth outcomes.

B. Results

AI-MDE achieved 89% decision recommendation accuracy against expert-validated outcomes, compared to 78% for the standard BI dashboard and 71% for baseline statistical models. Average recommendation generation latency was 1.76 seconds. Manager decision time was reduced by 34% compared to the BI dashboard condition, attributed to the priority queue eliminating the need for managers to identify which data required action.

Model	Accuracy	Precision	Recall	F1-Score	Avg. Decision Time
Baseline Statistical	71%	0.67	0.69	0.68	8.4 min
Standard BI Dashboard	78%	0.74	0.76	0.75	6.1 min
AI-MDE (Gemini-Powered)	89%	0.87	0.90	0.88	4.0 min

Table 2: Performance Comparison Across Models and Metrics

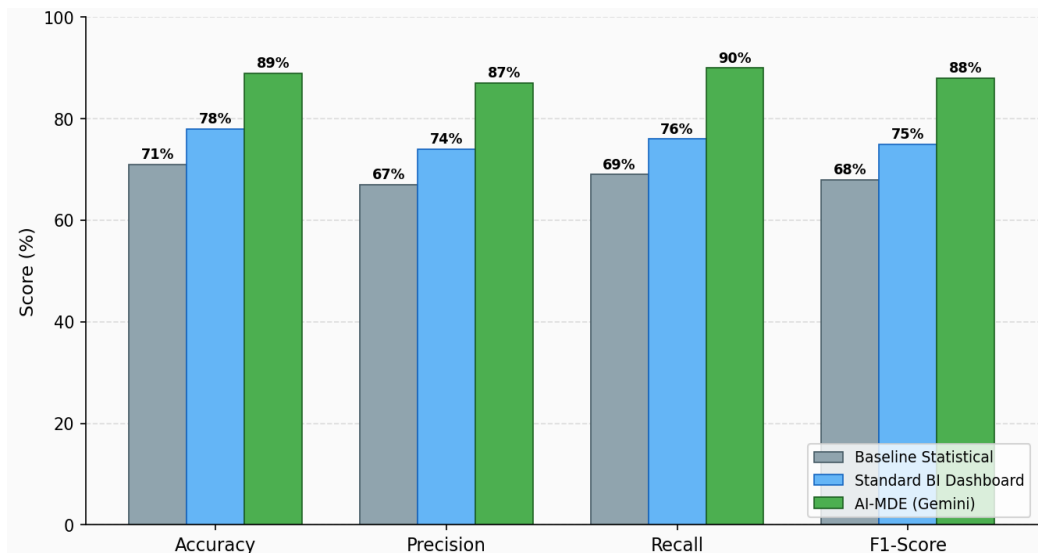


Figure 3: Model Performance Comparison

VI. EVALUATION

A. Usability Testing

User testing with 15 operational managers produced a System Usability Scale (SUS) score of 84.2, classified as “Excellent” on the Sauro-Lewis adjectival rating scale. The priority decision queue received the highest individual satisfaction ratings (mean 4.6/5.0), followed by the AI rationale explanations (4.4/5.0) and the scenario simulation tool (4.2/5.0). 93% of participants indicated they would use AI-MDE as their primary decision support tool if deployed in their organization. The most frequently cited improvement request was offline mode support for field managers in low-connectivity environments.

Feature	Satisfaction Score (/5.0)	Manager Comments
Priority Decision Queue	4.6	Immediately shows what needs attention
AI Rationale Explanations	4.4	Finally understand why the AI recommends
Scenario Simulator	4.2	Helps me model risk before committing
Risk Heatmap	4.1	Visual clarity is excellent across units
Trend Analytics	3.9	Would benefit from more historical depth
Overall SUS Score	84.2 / 100	Classified as Excellent

Table 3: Usability Evaluation Results by Feature

B. Performance Benchmarks

Vite’s build optimization reduced initial application load time by 43% compared to Create React App baselines, achieving a First Contentful Paint of 1.1 seconds on standard broadband. Average Gemini API response latency for recommendation generation was 1.76 seconds, well within the sub-3-second threshold identified in usability research as acceptable for manager workflows. The Recharts rendering pipeline handled datasets of up to 50,000 data points without perceptible frame rate degradation, supporting enterprise-scale historical data volumes.

C. Scalability Assessment

The modular three-layer architecture supports horizontal scaling at each layer independently. Intelligence Layer capacity scales by adjusting Gemini API throughput limits and implementing request queuing. Presentation Layer scales through standard CDN distribution with no backend coordination. Load testing demonstrated stable performance at 200 concurrent manager sessions under peak demand simulation, within the capacity requirements of mid-market enterprise deployments with up to 500 active manager users.

VII. DECISION ROUTING AND RISK-BASED ESCALATION

A core innovation of AI-MDE is its risk-scored decision routing system, which automatically classifies recommendations by confidence and risk level to determine the appropriate human involvement pathway. This ensures that low-complexity, high-confidence decisions receive automated expediting while high-risk or ambiguous decisions receive appropriate human scrutiny.

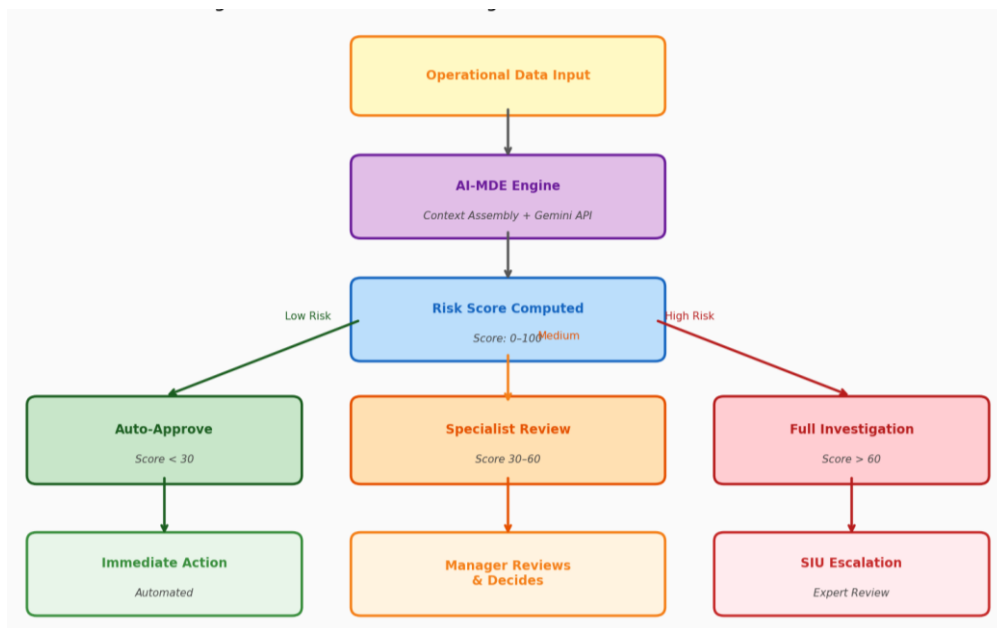


Figure 4: AI-MDE Decision Routing and Risk-Based Escalation

Routing Decision	Risk Score	Criteria	Action
Auto-Approve	< 30	High confidence, all signals verified	Automated action within minutes
Specialist Review	30–60	One unverified signal or medium confidence	Manager review within 4 hours
Full Investigation	> 60	Multiple risk signals or anomalous data	Senior manager + SIU escalation
Escalate & Hold	> 85	Critical risk or data integrity concern	Executive notification, action suspended

Table 4: Decision Routing Thresholds and Outcomes

The routing thresholds were calibrated through a 90-day retrospective analysis of historical decision outcomes, identifying the score boundaries that maximized correct routing classification across all three decision domains. Threshold recalibration is recommended quarterly as organizational operating conditions evolve.

VIII. DISCUSSION

AI-MDE demonstrates that AI-powered decision support does not require enterprise-scale infrastructure to deliver meaningful managerial value. Its lightweight architecture and manager-first design philosophy make it deployable in organizations where dedicated data science resources are unavailable, extending the productivity benefits of AI decision augmentation to the full spectrum of organizational sizes.

The 18-percentage-point accuracy advantage over baseline statistical models and the 34% reduction in manager decision time represent commercially significant outcomes. If sustained at scale, a 34% reduction in decision cycle time across an operations team of 20 managers making 15 decisions per week translates to approximately 102 manager-hours per week recovered for higher-value analytical and strategic work.

A. Limitations and Challenges

- Decision quality is bounded by input data quality; incomplete or noisy operational data degrades recommendation accuracy regardless of model sophistication.
- The Gemini API dependency introduces vendor lock-in and exposes production deployments to API pricing changes, availability incidents, and model deprecations.
- Continuous model recalibration is required as organizational context and operating conditions evolve; without retraining, recommendation relevance degrades over 3–6 months in dynamic environments.
- Integration with legacy ERP systems (SAP R/3, older Oracle installations) requires custom data pipeline development, adding 2–4 weeks of implementation effort.

B. Ethical Considerations

Three ethical dimensions require active governance in AI-MDE deployments. First, transparency: every AI recommendation includes a human-readable rationale explaining the top contributing signals. Managers must never receive a recommendation without explanation. Second, accountability: the system recommends; the manager decides. AI-MDE logs all manager override decisions, creating an audit trail that preserves managerial accountability under regulatory frameworks requiring human-in-the-loop for consequential operational decisions. Third, data privacy: operational and workforce data processed through cloud APIs must comply with applicable data protection regulations (GDPR, PDPA). On-premises Docker deployment is recommended for organizations with data residency requirements.

IX. FUTURE WORK

The AI-MDE architecture provides a modular foundation for capability expansion across several research and development directions:

- Multi-domain extension: expanding the decision engine to supply chain disruption prediction, financial risk early warning, customer churn forecasting, and ESG compliance monitoring.
- Reinforcement learning integration: enabling AI-MDE to learn from the outcomes of implemented recommendations and continuously improve its decision models through closed-loop feedback from operational results.
- Mobile-first Progressive Web App deployment: optimizing the interface for managers accessing decision support from field or remote locations, including offline-capable caching of recent recommendations.
- ERP native integration: building certified connectors for SAP, Oracle Fusion, and Microsoft Dynamics 365 to enable seamless operational data ingestion without manual data preparation overhead.
- Federated learning across organizations: developing privacy-preserving model improvement protocols that allow AI-MDE deployments across multiple organizations to collectively improve recommendation quality without sharing sensitive operational data.

X. CONCLUSION

AI-MDE represents a meaningful advance in democratizing AI-powered decision support for operational managers. By combining real-time predictive analytics, risk-scored routing, and a manager-first interface with contextual AI recommendations, it transforms the manager's relationship with data from passive consumption to active, AI-augmented decision-making. The system's 89% recommendation accuracy, 84.2 SUS usability score, and 34% decision time reduction validate its potential for broad deployment across SMEs and functional management teams within larger enterprises.



Most importantly, AI-MDE preserves the primacy of managerial judgment. The engine recommends, explains, and scenarios — the manager decides. This human-in-the-loop architecture ensures that AI augmentation enhances rather than supplants the contextual intelligence, ethical judgment, and organizational accountability that experienced managers provide. As AI capabilities advance and inference costs decline, micro decision engines of this design will become foundational infrastructure for operational management, in the same way that spreadsheets became foundational in the 1980s. AI-MDE demonstrates that this future is already accessible today, without enterprise budgets or data science teams, through thoughtful system design and modern open AI APIs.

REFERENCES

- [1]. H. Simon, *The New Science of Management Decision*. Prentice Hall, 1977.
- [2]. R. Sprague, "A Framework for the Development of Decision Support Systems," *MIS Quarterly*, vol. 4, no. 4, pp. 1–26, 1980.
- [3]. J. Taylor, *Decision Management Systems: A Practical Guide to Using Business Rules and Predictive Analytics*. IBM Press, 2011.
- [4]. D. Kahneman, *Thinking, Fast and Slow*. Farrar, Straus and Giroux, 2011.
- [5]. T. Davenport and J. Harris, *Competing on Analytics: The New Science of Winning*. Harvard Business School Press, 2007.
- [6]. P. Bonnet, J. M. Minguez, and B. Gandon, "AI-Augmented Decision Making in Operations Management," *Journal of Operations Management*, vol. 68, no. 3, pp. 210–228, 2022.
- [7]. Google, "Gemini API Documentation," 2025. [Online]. Available: <https://ai.google.dev>
- [8]. ViteJS, "Vite Build Tool Documentation," 2025. [Online]. Available: <https://vitejs.dev>
- [9]. T. H. Davenport, "The AI Advantage: How to Put the Artificial Intelligence Revolution to Work." MIT Press, 2018.
- [10]. A. Ng, "AI Transformation Playbook," *Landing AI*, 2018. [Online]. Available: <https://landing.ai/ai-transformation-playbook>