Enterprise Explainable Decision Intelligence Framework (EEDIF): A Service-Oriented Architecture for Trustworthy AI Decision Support.

G.Anthony Kavitha. M.Sc, M.Phil, B.Ed. Department of Computer Applications
TDA College of Arts & Science Kannirajapuram,Tamilnadu – 623 135 kavianto16@gmail.com

Abstract— Artificial Intelligence (AI) is increasingly embedded in enterprise decision support systems to improve operational efficiency, automate business processes, and support data-driven decision-making. Despite their predictive capabilities, many enterprise AI solutions operate as opaque "black box" systems, limiting transparency, governance, regulatory compliance, and stakeholder trust.
Existing Explainable Artificial Intelligence (XAI) techniques primarily focus on interpreting individual model predictions and provide limited support for enterprise-wide architectural integration.
This paper proposes the Enterprise Explainable Decision Intelligence Framework (EEDIF), a service-oriented architecture that introduces explainability as a reusable enterprise capability through the Enterprise Explainability Service (EES). The proposed architecture decouples explainability from AI models while providing standardized services for explanation generation, feature attribution, confidence assessment, policy validation, and audit logging.
A loan approval decision support case study demonstrates the practical application of the framework, and a qualitative evaluation assesses its effectiveness against enterprise requirements, including transparency, governance, auditability, scalability, and reusability. The proposed framework provides a structured architectural foundation for integrating trustworthy and explainable AI into enterprise decision support systems.

Keywords — Explainable Artificial Intelligence (XAI), Enterprise Architecture, Trustworthy AI, Decision Support Systems, Service-Oriented Architecture, Explainability, AI Governance.

I. INTRODUCTION

Artificial Intelligence (AI) has become a key enabler of enterprise decision support systems, supporting intelligent automation, predictive analytics, and operational decision-making across domains such as banking, healthcare, insurance, manufacturing, and public administration. As AI increasingly influences business-critical decisions, organizations require not only accurate predictions but also transparent, explainable, and trustworthy decision-making processes.
Although modern AI models achieve high predictive performance, many operate as opaque "black box" systems whose internal reasoning is difficult for business users, auditors, and regulators to understand. This lack of explainability reduces stakeholder trust, complicates governance, and creates challenges in meeting regulatory and compliance requirements. Consequently, explainability has become a fundamental requirement for enterprise AI adoption.
Explainable Artificial Intelligence (XAI) addresses this challenge by providing techniques that interpret AI model predictions. Approaches such as SHAP and LIME improve transparency by identifying the factors influencing individual predictions. However, these methods are primarily model-centric and are typically embedded within individual AI applications. They provide limited support for enterprise-wide integration, reusable explainability services, or standardized governance across heterogeneous business systems.
Enterprise environments require explainability to be implemented as a shared architectural capability rather than as a feature of individual AI models. A reusable enterprise service can provide consistent explanation generation, confidence assessment, policy validation, and audit logging while reducing implementation complexity and improving governance across multiple AI-enabled applications.
To address this gap, this paper proposes the **Enterprise Explainable Decision Intelligence Framework (EEDIF)**, a service-oriented architecture that introduces the **Enterprise Explainability Service (EES)** as a reusable enterprise component for trustworthy AI decision support. By separating explainability from predictive models, the proposed architecture enables standardized integration across enterprise decision support systems without requiring modifications to underlying AI implementations.
The primary contributions of this paper are:

1. A service-oriented architecture, **EEDIF**, for integrating explainability into enterprise AI decision support systems.

2. A reusable **Enterprise Explainability Service (EES)** providing explanation generation, feature attribution, confidence assessment, policy validation, and audit logging.

3. A loan approval decision support case study demonstrating the practical applicability of the proposed architecture.

4. A qualitative evaluation of the framework against enterprise requirements, including transparency, governance, auditability, scalability, and reusability.

The remainder of this paper is organized as follows. Section II reviews related work and identifies the research gap. Section III presents the proposed EEDIF architecture.

Section IV describes the Enterprise Explainability Service. Section V presents the enterprise explainable decision workflow, followed by the loan approval case study in Section VI. Section VII evaluates the proposed framework, and Section VIII concludes the paper.

II. Related Work and Research Gap

Explainable Artificial Intelligence (XAI) has become an important research area for improving the transparency and interpretability of machine learning models. Widely adopted approaches such as SHAP (SHapley Additive exPlanations) and LIME (Local Interpretable Model-Agnostic Explanations) explain individual model predictions by identifying the contribution of input features. These techniques have significantly improved the interpretability of AI models across domains including finance, healthcare, and fraud detection.

Beyond model-level explainability, recent research has focused on trustworthy AI, responsible AI, and AI governance. Standards and guidelines such as the NIST AI Risk Management Framework, IEEE 7001, and the European Commission's Trustworthy AI Guidelines emphasize transparency, accountability, fairness, and auditability. However, these initiatives primarily define governance principles rather than providing an architectural approach for integrating explainability into enterprise decision support systems.

Despite these advances, existing approaches remain largely model-centric. Enterprise organizations typically operate multiple AI services across heterogeneous business applications, making it difficult to maintain consistent explainability, governance, and audit capabilities.

Embedding explainability independently within each application often results in duplicated implementation effort, inconsistent explanation quality, and increased maintenance complexity.

Unlike existing research that primarily focuses on model-level explainability techniques, this work approaches explainability from an enterprise architecture perspective. The proposed **Enterprise Explainable Decision Intelligence Framework (EEDIF)** introduces explainability as a reusable enterprise capability through service-oriented architecture. By decoupling explainability from individual AI models, the framework enables standardized explanation generation, governance, policy validation, and auditability across heterogeneous enterprise decision support systems.

III. Enterprise Explainable Decision Intelligence Framework (EEDIF)

The **Enterprise Explainable Decision Intelligence Framework (EEDIF)** is a service-oriented architecture that integrates explainability into enterprise decision support systems as a reusable organizational capability. Instead of embedding explainability within individual AI models, the framework separates explainability into an independent enterprise service, enabling standardized deployment across heterogeneous applications while promoting transparency, governance, and maintainability.

As illustrated in **Fig. 1**, the proposed architecture is organized into four logical layers: **Enterprise Applications**, **Decision Orchestration Layer**, **Enterprise Explainability Service (EES)**, and **AI Models & Enterprise Data Sources**. These layers are supported by cross-cutting enterprise capabilities, including governance, security, monitoring, auditing, configuration management, and integration services, ensuring consistent operation across the enterprise AI ecosystem.

The architecture is guided by the following design principles:

- **Separation of Concerns:** Explainability is implemented independently from AI prediction models.
- **Service Reusability:** A common explainability service supports multiple enterprise applications.
- **Interoperability:** Standardized interfaces enable integration with heterogeneous AI services.
- **Governance by Design:** Policy validation and audit logging are incorporated into the decision lifecycle.
- **Human-Centric Decision Support:** AI recommendations are accompanied by transparent explanations to support informed human decision-making.

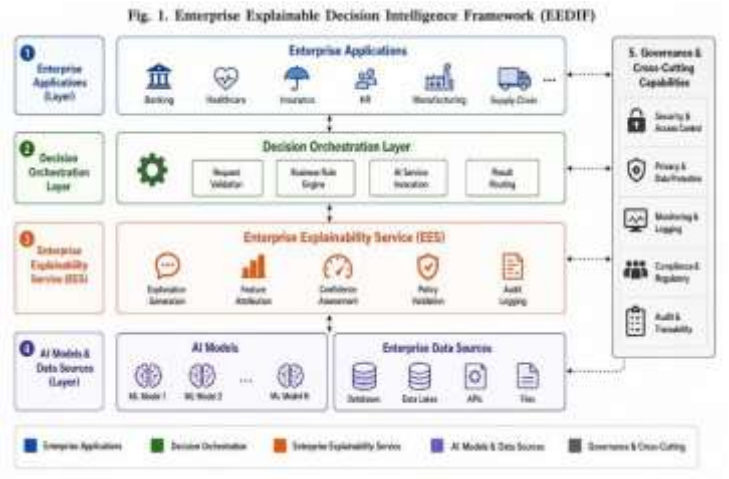


Fig. 1. Enterprise Explainable Decision Intelligence Framework (EEDIF) illustrates the layered service-oriented architecture and the interaction between enterprise applications, decision orchestration, the Enterprise Explainability Service (EES), AI models, enterprise data sources, and cross-cutting governance capabilities.

The **Enterprise Applications** layer represents the business systems that consume AI-assisted decision support, including banking, healthcare, insurance, human resources, manufacturing, and supply chain applications. These systems initiate decision requests through standardized interfaces without requiring knowledge of the underlying AI implementation or explainability mechanisms.

The **Decision Orchestration Layer** coordinates the end-to-end execution of enterprise decisions. It validates incoming requests, applies business rules, invokes the appropriate AI prediction service, and routes prediction results to the Enterprise Explainability Service. This orchestration layer enables consistent workflow management while isolating business processes from AI-specific implementation details.

The **Enterprise Explainability Service (EES)** forms the central component of the proposed architecture. Rather than embedding explainability within individual AI models, EES provides a reusable enterprise service responsible for generating human-readable explanations, identifying influential features, assessing prediction confidence, validating organizational policies, and maintaining complete audit records. This separation promotes architectural consistency, simplifies maintenance, and enables enterprise-wide reuse across heterogeneous AI applications.

The **AI Models and Enterprise Data Sources** layer contains predictive models together with enterprise databases, historical records, and domain knowledge repositories. EEDIF remains model-agnostic, allowing different machine learning technologies to be integrated without modifying the explainability service.

The proposed layered architecture enables organizations to introduce explainability as an enterprise capability while preserving existing AI investments. By separating prediction from explanation, EEDIF improves interoperability, governance, maintainability, and scalability across enterprise decision support systems.

IV. Enterprise Explainability Service (EES)

The **Enterprise Explainability Service (EES)** is the principal architectural component of EEDIF, providing a standardized mechanism for generating explainable and trustworthy AI decisions. Implemented as an independent enterprise service, EES enables multiple business applications to share a common explainability capability while remaining independent of individual AI models.

As illustrated in **Fig. 2**, EES consists of five cooperating components that transform AI predictions into enterprise-ready explainable decision packages.



Fig. 2. Internal architecture of the Enterprise Explainability Service (EES) showing explanation generation, feature attribution, confidence assessment, policy validation, audit logging, and standardized enterprise interfaces.

Table I. Enterprise Explainability Service Components

Component	Primary Responsibility
Explanation Generator	Produces human-readable explanations for AI predictions.
Feature Attribution	Identifies the most influential factors contributing to a prediction.
Confidence Analyzer	Evaluates prediction confidence and reliability indicators.
Policy Validator	Verifies compliance with organizational rules and regulatory requirements.
Audit Logger	Records explainability artifacts and decision metadata for governance and traceability.

Together, these components transform raw AI predictions into transparent, auditable, and business-friendly decision outputs. By centralizing explainability within a reusable enterprise service, EEDIF promotes consistent governance and simplifies the integration of trustworthy AI across multiple enterprise applications.

V. Enterprise Explainable Decision Workflow

The Enterprise Explainable Decision Intelligence Framework (EEDIF) integrates explainability into the enterprise decision lifecycle through the Enterprise Explainability Service (EES). Rather than returning AI predictions directly to business applications, every prediction is processed by EES, where explanations, confidence indicators, policy validation, and audit information are generated before the final decision is presented to business users.

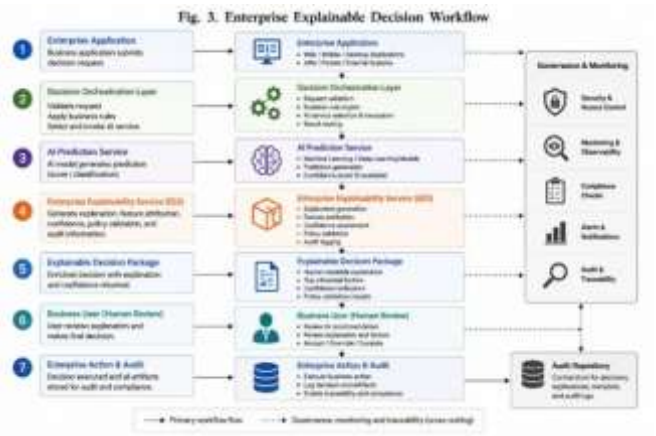


Fig. 3. Enterprise Explainable Decision Workflow illustrates the end-to-end interaction between enterprise applications, the Decision Orchestration Layer, AI prediction services, the Enterprise Explainability Service (EES), governance functions, human decision makers, and enterprise audit repositories.

As illustrated in **Fig. 3**, the workflow begins when an enterprise application submits a decision request to the Decision Orchestration Layer. After validating business rules, the orchestration layer invokes the appropriate AI prediction service. The resulting prediction is forwarded to the Enterprise Explainability Service, where explainability and governance activities are performed.

EES generates a comprehensive explainable decision package consisting of a human-readable explanation, feature importance, confidence assessment, policy validation

results, and audit information. The enriched decision is then presented to the business user, enabling transparent human review before the final business action is executed.

By separating explainability from AI prediction services, the proposed workflow provides a standardized and reusable decision process that supports transparency, governance, regulatory compliance, and trustworthy AI adoption across heterogeneous enterprise applications.

VI. AI-Assisted Loan Decision Support Case Study

To demonstrate the practical applicability of the proposed architecture, EEDIF is applied to an AI-assisted loan approval process. The case study illustrates how explainability can be incorporated into an enterprise decision support workflow without modifying the underlying AI prediction model.

A customer submits a loan application containing information such as income, employment status, repayment history, credit score, existing liabilities, and requested loan amount. After initial validation, the Decision Orchestration Layer invokes the AI prediction service to estimate the applicant's creditworthiness.

Rather than returning the prediction directly to the loan management system, the result is forwarded to the Enterprise Explainability Service (EES). As illustrated in **Fig. 4**, EES generates a human-readable explanation, identifies the most influential decision factors, evaluates prediction confidence, validates organizational lending policies, and records an audit trail before returning an explainable decision package.

The loan officer reviews both the AI recommendation and its supporting explanation before making the final approval decision. This human-in-the-loop approach improves transparency and enables informed decision-making while preserving organizational accountability.

The case study demonstrates how EEDIF integrates explainability, governance, and auditability into enterprise AI workflows through a reusable service-oriented architecture that can be extended to other domains such as healthcare, insurance, recruitment, and supply chain management.

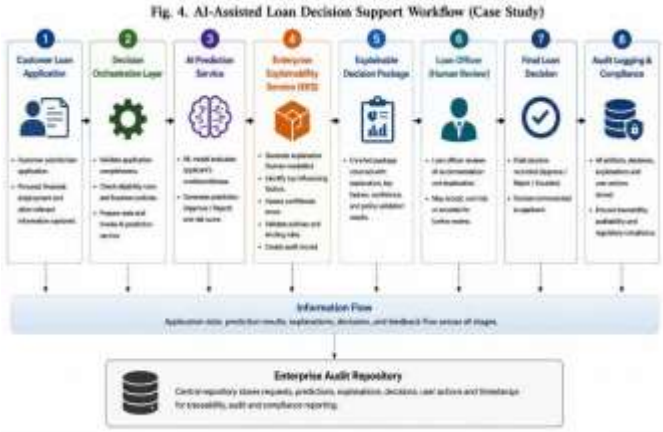


Fig. 4. AI-assisted loan decision support workflow demonstrating the application of EEDIF from customer request through AI prediction, Enterprise Explainability Service (EES), human review, final decision, and enterprise audit logging.

VII. Framework Evaluation

The proposed Enterprise Explainable Decision Intelligence Framework (EEDIF) is qualitatively evaluated against key enterprise requirements commonly associated with trustworthy AI systems. The evaluation focuses on the architectural capabilities of the framework rather than predictive model performance, as the primary contribution of this work is the integration of explainability into enterprise decision support through reusable service-oriented architecture.

As illustrated in **Fig. 5**, EEDIF addresses seven fundamental enterprise requirements: transparency, explainability, governance, auditability, reusability, scalability, and trustworthiness. By separating explainability from AI prediction models, the framework provides a standardized mechanism for generating consistent explanations while supporting organizational governance and regulatory compliance across heterogeneous enterprise applications.



Fig. 5. Qualitative evaluation of EEDIF showing how the proposed architecture supports transparency, explainability, governance, auditability, reusability, scalability, and trustworthy AI adoption within enterprise decision support systems.

Table II. Qualitative Evaluation of EEDIF

Evaluation Criterion	Assessment
Transparency	High
Explainability	High
Governance	High
Auditability	High
Reusability	High
Scalability	High
Trustworthiness	High

The evaluation indicates that EEDIF provides a consistent architectural foundation for integrating explainability into enterprise AI systems. Rather than replacing existing AI models, the framework complements them by introducing standardized explainability services that support transparent, auditable, and human-centered decision support.

VIII. Conclusion

This paper presented the **Enterprise Explainable Decision Intelligence Framework (EEDIF)**, a service-oriented architecture that introduces explainability as a reusable enterprise capability for trustworthy AI decision support.

The proposed architecture separates explainability from AI prediction models through the **Enterprise Explainability Service (EES)**, enabling standardized explanation generation, feature attribution, confidence assessment, policy validation, and audit logging across enterprise applications.

A loan approval decision support case study demonstrated how the proposed architecture can be integrated into enterprise workflows without modifying underlying AI models. The qualitative evaluation further showed that EEDIF supports transparency, governance, auditability, scalability, and reusability, providing a practical foundation for enterprise adoption of trustworthy AI.

Future work includes implementing the proposed architecture in operational enterprise environments, conducting quantitative evaluations across multiple application domains, and extending the framework to support emerging explainability techniques and evolving AI governance standards.

IX. References

[1] A. Adadi and M. Berrada, "Peeking Inside the Black-Box: A Survey on Explainable Artificial Intelligence (XAI)," *IEEE Access*, vol. 6, pp. 52138–52160, 2018.

[2] S. M. Lundberg and S.-I. Lee, "A Unified Approach to Interpreting Model Predictions," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2017, pp. 4765–4774.

[3] M. T. Ribeiro, S. Singh, and C. Guestrin, "Why Should I Trust You? Explaining the Predictions of Any Classifier," in *Proc. ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining*, 2016, pp. 1135–1144.

[4] R. Guidotti, A. Monreale, F. Turini, D. Pedreschi, and F. Giannotti, "A Survey of Methods for Explaining Black Box Models," *ACM Computing Surveys*, vol. 51, no. 5, pp. 1–42, 2018.

[5] N. Barredo Arrieta *et al.*, "Explainable Artificial Intelligence (XAI): Concepts, Taxonomies, Opportunities and Challenges Toward Responsible AI," *Information Fusion*, vol. 58, pp. 82–115, 2020.

[6] D. Gunning and D. Aha, "DARPA's Explainable Artificial Intelligence (XAI) Program," *AI Magazine*, vol. 40, no. 2, pp. 44–58, 2019.

[7] T. Erl, *Service-Oriented Architecture: Concepts, Technology, and Design*. Upper Saddle River, NJ, USA: Prentice Hall, 2005.

[8] M. Fowler, *Patterns of Enterprise Application Architecture*. Boston, MA, USA: Addison-Wesley, 2002.

[9] ISO/IEC 23053:2022, *Framework for Artificial Intelligence Systems Using Machine Learning*. Geneva, Switzerland: International Organization for Standardization, 2022.

[10] National Institute of Standards and Technology, *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. Gaithersburg, MD, USA: NIST, 2023.

[11] IEEE Standards Association, *IEEE Std 7001-2021: Transparency of Autonomous Systems*. Piscataway, NJ, USA, 2021.

[12] ISO/IEC 22989:2022, *Artificial Intelligence—Concepts and Terminology*. Geneva, Switzerland: International Organization for Standardization, 2022.

[13] European Commission High-Level Expert Group on Artificial Intelligence, *Ethics Guidelines for Trustworthy AI*.Brussels, Belgium, 2019.

[14] Organisation for Economic Co-operation and Development (OECD), *OECD Principles on Artificial Intelligence*. Paris, France, 2019.

[15] T. H. Davenport and R. Ronanki, "Artificial Intelligence for the Real World," *Harvard Business Review*,vol. 96, no. 1, pp. 108–116, Jan.–Feb. 2018.

[16] Microsoft, *Microsoft Responsible AI Standard*, Version 2, Redmond, WA, USA, 2022.

[17] IBM Research, *AI Explainability 360: An Extensible Toolkit for Understanding Machine Learning Models*, 2019.

[18] J. Pearl, *Causality: Models, Reasoning, and Inference*, 2nd ed. Cambridge, U.K.: Cambridge University Press, 2009.