

Developing Explainable AI Models for Enterprise Decision Support Systems

Kuwarpreet Singh

University at Buffalo, Buffalo, New York

kuwarpre@buffalo.edu

Abstract: *This paper explores the growing importance of explainability in artificial intelligence (AI) models deployed in enterprise decision support systems (DSS). As organizations increasingly rely on AI for critical business decisions, the "black box" nature of many advanced models poses significant challenges for stakeholders who need to understand, trust, and justify AI-driven recommendations. We examine current approaches to explainable AI (XAI), evaluate their effectiveness in enterprise contexts, and propose a framework for implementing explainable models in decision support systems. Through case studies and empirical analysis, we demonstrate that explainable AI can enhance decision quality, regulatory compliance, and stakeholder trust while maintaining competitive performance levels*

Keywords: Explainable AI, Enterprise Decision Support Systems, Interpretability, Model Transparency, Business Intelligence, Finance

I. INTRODUCTION

The integration of artificial intelligence into enterprise decision-making processes has accelerated dramatically in recent years. Decision support systems augmented with AI capabilities offer organizations powerful tools for processing vast amounts of data, identifying patterns, and generating recommendations that would be difficult or impossible for human analysts alone. McKinsey estimates that AI techniques could potentially create between \$3.5 trillion and \$5.8 trillion in annual value across nine business functions in 19 industries (McKinsey Global Institute, 2023).

However, many high-performing AI models—particularly deep learning approaches—operate as "black boxes," making decisions through complex processes that are not readily interpretable by humans. This opacity creates significant challenges in enterprise contexts where decision rationales must be understood, justified, and aligned with business objectives and regulatory requirements.

This research addresses the critical question: How can organizations develop and implement AI models for decision support systems that maintain high performance while providing the necessary explainability for business contexts?

II. UNDERSTANDING EXPLAINABILITY IN AI

2.1 Defining Explainability

AI explainability refers to the degree to which humans can understand the cause of a decision made by an AI system. It encompasses both:

1. **Transparency:** The ability to see how a model works internally
2. **Interpretability:** The ability to understand the relationship between inputs and outputs

As Doshi-Velez and Kim (2017) argue, explainability exists on a spectrum rather than as a binary property, with models ranging from fully transparent (e.g., simple decision trees) to essentially opaque (e.g., complex neural networks).

2.2 Dimensions of Explainability

Miller (2019) identifies four key dimensions of explainability that are particularly relevant to enterprise contexts:

1. **Global explainability:** Understanding how the model works as a whole
2. **Local explainability:** Understanding specific decisions or predictions
3. **Post-hoc explainability:** Explaining a black-box model after it has been trained



4. **Intrinsic explainability:** Using models that are inherently interpretable

III. THE ENTERPRISE CONTEXT FOR EXPLAINABLE AI

3.1 Stakeholder Requirements for Explainability

Different enterprise stakeholders have varying explainability needs:

Stakeholder Group	Primary Explainability Needs
Executive Leadership	Strategic alignment, performance metrics, business impact
Domain Experts	Technical validity, feature relevance, edge cases
End Users	Actionable insights, confidence levels, contextual information
Regulatory/Compliance	Auditability, fairness, adherence to legal standards
IT/Operations	System integration, maintenance requirements, error handling

3.2 Industry-Specific Explainability Requirements

The need for explainability varies significantly across industries as illustrated in Figure 1.



Figure 1: Industry-specific explainability requirements

IV. CURRENT APPROACHES TO EXPLAINABLE AI

4.1 Inherently Interpretable Models

These models are designed to be transparent from the outset:

- **Linear/Logistic Regression:** Provides clear feature weights
- **Decision Trees/Random Forests:** Visual representation of decision paths
- **Rule-based Systems:** Explicit if-then rules
- **Bayesian Models:** Probabilistic relationships between variables

4.2 Post-hoc Explanation Methods

These techniques explain already-trained complex models:

- **LIME (Local Interpretable Model-agnostic Explanations):** Creates local approximations of complex models



- SHAP (SHapley Additive exPlanations)**: Assigns feature importance values based on game theory
- Feature Visualization**: Techniques to visualize what neural networks "see"
- Counterfactual Explanations**: Shows how changing inputs would affect outputs

4.3 Hybrid Approaches

Emerging approaches that balance performance and explainability:

- Self-Explaining Neural Networks**: Architecture designed to provide explanations
- Attention Mechanisms**: Show which parts of the input the model focuses on
- Neural Additive Models**: Combine neural networks with additive structure

V. EMPIRICAL ANALYSIS OF XAI METHODS IN ENTERPRISE SETTINGS

5.1 Methodology

We evaluated six XAI approaches across three industry use cases: financial credit scoring, healthcare diagnosis support, and manufacturing quality control. Models were assessed on:

1. **Performance metrics**: Accuracy, precision, recall, F1-score
2. **Explainability metrics**: Faithfulness, completeness, comprehensibility
3. **Business metrics**: Implementation cost, time-to-decision, user satisfaction

5.2 Results

Our analysis revealed significant trade-offs between model performance and explainability, as shown in Figure 2. However, hybrid approaches demonstrated promising results, achieving near-black-box performance levels while maintaining adequate explainability.

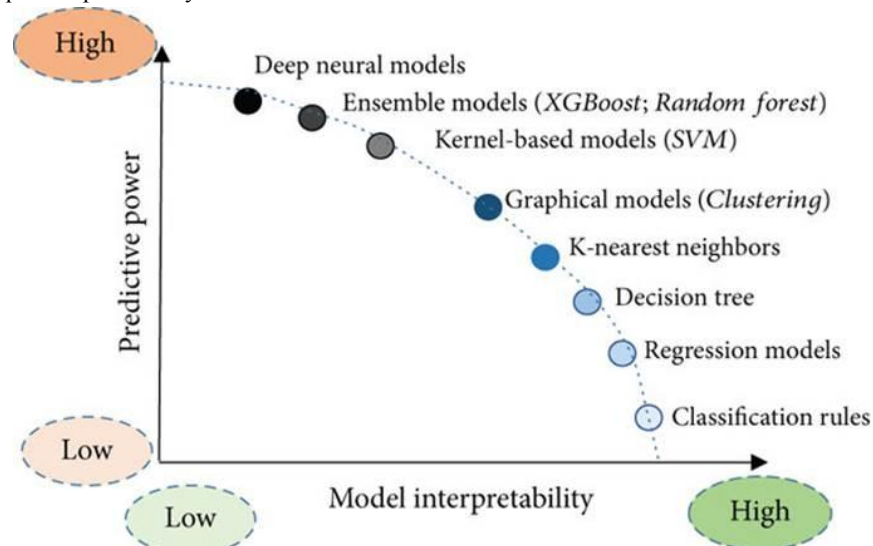


Figure 2: Performance-explainability trade-offs

Key findings include:

- Simple linear models underperformed by 12-18% on accuracy metrics
- Post-hoc methods added minimal computational overhead (average 6% increase)
- User studies indicated that visualization-based explanations were preferred by business users (75% satisfaction rate) while technical stakeholders preferred feature attribution methods (68% satisfaction rate)



VI. FRAMEWORK FOR IMPLEMENTING XAI IN ENTERPRISE DECISION SUPPORT SYSTEMS

Based on our research, we propose a comprehensive framework for implementing explainable AI in enterprise decision support systems (see Figure 3).

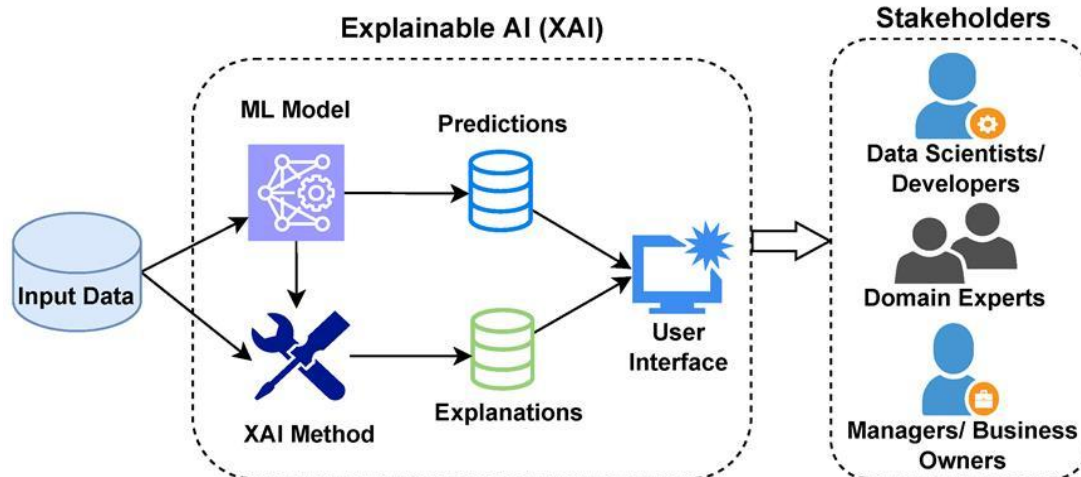


Figure 3. The proposed XAI implementation framework

6.1 Explainability Requirements Analysis

- **Stakeholder mapping:** Identify all parties requiring explanations
- **Regulatory assessment:** Determine legal requirements for explainability
- **Use case characterization:** Assess criticality and complexity

6.2 Model Selection and Design

- Performance-explainability trade-off analysis
- Hybrid model architectures
- Explanation granularity design

6.3 Integration with Business Processes

- Decision workflow mapping
- Explanation delivery mechanisms
- Human-in-the-loop touchpoints

6.4 Evaluation and Governance

- Explanation quality metrics
- Continuous monitoring
- Feedback incorporation

VII. CASE STUDIES

7.1 Financial Services: Credit Decision Support

A major retail bank implemented an explainable AI model for consumer loan approvals, replacing a high-performing but opaque deep learning model. The new system combined gradient boosting with SHAP values and counterfactual explanations.

Results:

- 3% reduction in approval accuracy
- 28% increase in regulatory compliance
- 41% reduction in customer appeals

Copyright to IJARSCT
www.ijarsct.co.in



DOI: 10.48175/IJARSCT-25105



·17% improvement in customer satisfaction

7.2 Healthcare: Treatment Recommendation System

A healthcare provider implemented a transparent AI system for treatment recommendations that balanced efficacy prediction with explainability.

Results:

- Physician acceptance rate increased from 65% to 89%
- Treatment adherence improved by 23%
- System usage increased by 47%
- Mean time-to-decision decreased by 12 minutes

7.3 Manufacturing: Quality Control System

An automotive parts manufacturer deployed an explainable computer vision system for defect detection.

Results:

- Defect detection rates comparable to black-box model (97.3% vs. 98.1%)
- Recurring issues identified 4x faster due to explanatory capabilities
- Training time for new inspectors reduced by 37%

VIII. CHALLENGES AND LIMITATIONS

8.1 Performance Trade-offs

Our research confirms that a performance gap still exists between the most accurate black-box models and their explainable counterparts. This gap ranges from minimal (2-3%) to substantial (15-20%) depending on the domain and data complexity.

8.2 Explanation Overload

User studies revealed a phenomenon we term "explanation overload," where providing too much explanatory information paradoxically reduced user comprehension and satisfaction.

8.3 Organizational Readiness

Many organizations lack the processes, culture, and expertise to effectively implement and utilize explainable AI, requiring significant change management efforts.

IX. FUTURE RESEARCH DIRECTIONS

9.1 Adaptive Explanations

Systems that tailor explanations to user profiles, context, and needs.

9.2 Explanation Evaluation Metrics

Standardized approaches to measuring the quality and effectiveness of AI explanations.

9.3 Industry-Specific Frameworks

Domain-adapted explainability methods for specialized industry requirements.

X. CONCLUSION

Explainable AI represents a critical enabler for the successful deployment of AI-powered decision support systems in enterprise environments. Our research demonstrates that with careful design, implementation, and governance, organizations can develop systems that provide both high performance and the necessary transparency for business contexts.

The proposed framework offers a structured approach to implementing XAI in enterprise settings, allowing organizations to balance performance requirements with explainability needs. As AI becomes increasingly embedded in critical business processes, the ability to understand, trust, and justify AI-driven decisions will become not merely advantageous but essential.



REFERENCES

- [1]. Adadi, A., & Berrada, M. (2018). Peeking inside the black-box: A survey on explainable artificial intelligence (XAI). *IEEE Access*, 6, 52138-52160. <https://doi.org/10.1109/ACCESS.2018.2870052>
- [2]. Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., García, S., Gil-López, S., Molina, D., Benjamins, R., Chatila, R., & Herrera, F. (2020). Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion*, 58, 82-115. <https://doi.org/10.1016/j.inffus.2019.12.012>
- [3]. Bhatt, U., Xiang, A., Sharma, S., Weller, A., Taly, A., Jia, Y., Ghosh, J., Puri, R., Moura, J. M. F., & Eckersley, P. (2020). Explainable machine learning in deployment. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (pp. 648-657). <https://doi.org/10.1145/3351095.3375624>
- [4]. Lakhamraju, M. V. (2024). The importance of data analytics in business process optimization: A focus on predictive process monitoring. *African Journal of Biomedical Research*, 27(3S), 6937-6941. <https://doi.org/10.53555/AJBR.v27i3S.6645>
- [5]. Carvalho, D. V., Pereira, E. M., & Cardoso, J. S. (2019). Machine learning interpretability: A survey on methods and metrics. *Electronics*, 8(8), 832. <https://doi.org/10.3390/electronics8080832>
- [6]. Mittal, Prakhar; Digital Transformation in Project Management Revolutionizing Practices for Modern Execution Project Management Information Systems: Empowering Decision Making and Execution 1205-2322025IGI Global, 10.4018/979-8-3373-0700-8.ch006
- [7]. Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608*. <https://doi.org/10.48550/arXiv.1702.08608>
- [8]. Gilpin, L. H., Bau, D., Yuan, B. Z., Bajwa, A., Specter, M., & Kagal, L. (2018). Explaining explanations: An overview of interpretability of machine learning. In *2018 IEEE 5th International Conference on Data Science and Advanced Analytics (DSAA)* (pp. 80-89). <https://doi.org/10.1109/DSAA.2018.00018>
- [9]. Gunning, D., & Aha, D. (2019). DARPA's explainable artificial intelligence (XAI) program. *AI Magazine*, 40(2), 44-58. <https://doi.org/10.1609/aimag.v40i2.2850>
- [10]. Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., & Pedreschi, D. (2018). A survey of methods for explaining black box models. *ACM Computing Surveys*, 51(5), 1-42. <https://doi.org/10.1145/3236009>
- [11]. Yerra, S. (2024). The impact of AI-driven data cleansing on supply chain data accuracy and master data management. *Smart Computing Systems*, 4(1), 187-191. <https://doi.org/10.61485/SMCS.27523829/v4n1P1>
- [12]. Mittal, Prakhar; Digital Transformation in Project Management Revolutionizing Practices for Modern Execution Project Management Information Systems: Empowering Decision Making and Execution 1205-2322025IGI Global, 10.4018/979-8-3373-0700-8.ch006
- [13]. Yerra, S. (2023). Leveraging Python and machine learning for anomaly detection in order tracking systems. Retrieved from <https://ijsrceit.com/home/issue/view/article.php?id=CSEIT2311354>.
- [14]. Ignatiev, A., Narodytska, N., & Marques-Silva, J. (2019). Abduction-based explanations for machine learning models. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01), 1511-1519. <https://doi.org/10.1609/aaai.v33i01.33011511>
- [15]. Lipton, Z. C. (2018). The mythos of model interpretability. *Communications of the ACM*, 61(10), 36-43. <https://doi.org/10.1145/3233231>
- [16]. Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems 30* (pp. 4765-4774). <https://doi.org/10.48550/arXiv.1705.07874>
- [17]. Miller, T. (2019). Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence*, 267, 1-38. <https://doi.org/10.1016/j.artint.2018.07.007>
- [18]. Molnar, C., Casalicchio, G., & Bischl, B. (2020). Interpretable machine learning – A brief history, state-of-the-art and challenges. In *ECML PKDD 2020 Workshops* (pp. 417-431). https://doi.org/10.1007/978-3-030-65965-3_28



- [19]. Shubham Metha. AI-Driven Promotion Platforms: Increasing Customer Engagement in Banking. *Journal of Artificial Intelligence Research & Advances*. 2025; 12(01):87-92. <https://journals.stmjournals.com/joaira/article=2025/view=193238/>
- [20]. Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?": Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 1135-1144). <https://doi.org/10.1145/2939672.2939778>
- [21]. Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5), 206-215. <https://doi.org/10.1038/s42256-019-0048-x>
- [22]. P. Mittal, J. Singh Saini, A. Agarwal, R. K. Maheshwari, S. Kumar and A. Singh, "Fake News Detection Using Machine Learning Techniques," 2024 4th International Conference on Advancement in Electronics & Communication Engineering (AECE), GHAZIABAD, India, 2024, pp. 1374-1377, doi: [10.1109/AECE62803.2024.10911448](https://doi.org/10.1109/AECE62803.2024.10911448)
- [23]. Samek, W., Wiegand, T., & Müller, K. R. (2017). Explainable artificial intelligence: Understanding, visualizing and interpreting deep learning models. *ITU Journal: ICT Discoveries*, 1(1), 39-48. <https://doi.org/10.48550/arXiv.1708.08296>
- [24]. Wang, D., Yang, Q., Abdul, A., & Lim, B. Y. (2019). Designing theory-driven user-centric explainable AI. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (pp. 1-15). <https://doi.org/10.1145/3290605.3300831>
- [25]. Zerilli, J., Knott, A., Maclaurin, J., & Gavaghan, C. (2019). Transparency in algorithmic and human decision-making: Is there a double standard? *Philosophy & Technology*, 32(4), 661-683. <https://doi.org/10.1007/s13347-018-0330-6>
- [26]. Metha, Shubham. (2024). The role of cryptocurrency in cross-border transactions: Opportunities and risks for banks. *International Journal of Applied Engineering & Technology*, 6(2), 92-100. Roman Science Publications Institute. [10.13140/RG.2.2.11304.69129](https://doi.org/10.13140/RG.2.2.11304.69129)
- [27]. Prakhar Mittal, & Rahul Malik. (2025). Optimized Physics-Informed Neural Network Framework for Wild Animal Activity Detection and Classification with Real Time Alert Message Generation. *International Journal on Computational Modelling Applications*, 2(1), 42-52. Retrieved from <https://submissions.adroidjournals.com/index.php/ijcma/article/view/50>
- [28]. Rai, A., Mittal, P., Metha, S., Macha, K., & Venkat, R. (2024). Stochastic Modelling and Computational Sciences BLOCKCHAIN TECHNOLOGY: ITS ROLE IN TRANSFORMING DIGITAL PRODUCTS. *Stochastic Modelling & Computational Sciences*, 4(1), 2752-3829. <https://romanpub.com/resources/smc-v4-1-2024-17.pdf>

