A Differentially Private Interpretable Machine Learning Framework for Bank Failure Prediction

KOLA PRANATHI

PG Scholar. Department of MCA, DNR College, Bhimavaram, Andhra Pradesh

K. Rambabu

(Assistant Professor), Master of Computer Applications, DNR College, Bhimavaram, Andhra Pradesh

ABSTRACT

The prediction of bank failures is a critical task in financial risk management, enabling regulatory authorities and financial institutions to take proactive measures to prevent economic instability. With the increasing availability of financial data, machine learning models have been widely adopted to improve prediction accuracy. However, these models often face a trade-off between explainability and data privacy. Black-box models, while accurate, lack transparency, whereas interpretable models may expose sensitive financial information. This research proposes a novel framework that balances explainability and privacy using a differentially private glass-box approach for bank failure prediction.The proposed system integrates interpretable machine learning models with differential privacy mechanisms to ensure that sensitive financial data remains protected during model training and inference. Glass-box models such as decision trees and logistic regression are employed due to their inherent interpretability. These models allow stakeholders to understand the reasoning behind predictions, which is essential for regulatory compliance and decision-making.Differential privacy is incorporated into the model training process to prevent leakage of sensitive information. By adding controlled noise to the training data and model parameters, the system ensures that individual data records cannot be inferred from the model outputs. This approach maintains a balance between model accuracy and privacy preservation.The system is implemented using Python and deployed within a Django-based web framework. It processes financial datasets containing key indicators such as capital adequacy ratios, liquidity ratios, asset quality metrics, and profitability indicators. These features are used to train the predictive model.Performance evaluation is conducted using metrics such as accuracy, precision, recall, and F1-score. The results demonstrate that the proposed framework achieves competitive prediction accuracy while ensuring strong privacy guarantees. Additionally, the interpretability of the model enables stakeholders to identify key risk factors contributing to bank failures.This research contributes to the field of financial analytics by providing a privacy-preserving and interpretable solution for bank failure prediction. The framework can be used by regulators and financial institutions to enhance risk assessment and ensure compliance with data protection regulations.Future work may involve extending the framework to incorporate deep learning models with explainability techniques and exploring advanced privacy-preserving methods such as federated learning.

Keywords: Bank Failure Prediction, Differential Privacy, Explainable AI, Financial Risk Analysis, Glass-Box Models, Privacy-Preserving Machine Learning, Regulatory Compliance

I. INTRODUCTION

Bank failures pose significant risks to financial stability and economic growth. The collapse of financial institutions can lead to widespread economic consequences, including loss of investor confidence, disruption of financial markets, and economic downturns. Therefore, early detection of potential bank failures is essential for effective risk management.In recent years, machine learning techniques have been increasingly used for bank failure prediction. These models analyze large volumes of financial data to identify patterns and predict potential failures. However, the use of machine learning in financial systems introduces challenges related to explainability and privacy.Explainability is crucial in financial applications, as decisions must be transparent and justifiable. Regulatory authorities require clear explanations for predictions to ensure accountability and compliance. Black-box models, such as deep neural networks, often lack interpretability, making them unsuitable for high-stakes financial decisions.On the other hand, financial data is highly sensitive and must be protected to maintain confidentiality and comply with data protection regulations. Traditional machine learning models may expose sensitive information during training or inference, posing privacy risks.This research addresses the challenge of balancing explainability and privacy in bank failure prediction. The proposed system uses glass-box models, which provide clear and interpretable predictions, combined with differential privacy techniques to protect sensitive data.The system is implemented using Python and Django, providing a web-based interface for analyzing financial data and predicting bank failures. It supports data visualization and reporting, enabling users to understand risk factors and make informed decisions.The main contributions of this research include the development of a privacy-preserving machine learning framework, integration of interpretable models, and implementation of a scalable system for financial risk analysis. The proposed approach enhances both transparency and data security.

II. LITERATURE SURVEY (WITH EXISTING METHODS)

Bank failure prediction has been extensively studied using statistical and machine learning approaches. Traditional methods include logistic regression and discriminant analysis, which are simple and interpretable but may not capture complex relationships in data.Machine learning techniques such as Decision Trees, Support Vector Machines, and Random Forests have been widely used for prediction tasks. These models improve accuracy but may lack transparency, especially in ensemble methods.Deep learning models have been applied to financial prediction problems due to their ability to capture complex patterns. However, these models are often considered black boxes, making it difficult to interpret their predictions.

Explainable AI (XAI) techniques have been introduced to address the lack of interpretability. Methods such as LIME and SHAP provide post-hoc explanations for model predictions. While effective, these methods do not guarantee inherent interpretability.Privacy-preserving machine learning has gained attention with the introduction of differential privacy. Differential privacy ensures that the inclusion or exclusion of a single data point does not significantly affect the model output. This provides strong privacy guarantees.Recent research has explored the integration of differential privacy with machine learning models. However, most studies focus on black-box models, and there is limited research on combining differential privacy with interpretable models.This research addresses this gap by proposing a differentially private glass-box approach, combining interpretability and privacy in a single framework.

III. EXISTING SYSTEM

Existing bank failure prediction systems primarily rely on traditional statistical methods or machine learning models without privacy considerations. These systems use financial indicators to predict potential failures but often lack transparency and security.Black-box models such as neural networks provide high accuracy but do not offer explanations for predictions. This makes them unsuitable for regulatory environments where transparency is required.Additionally, existing systems do not incorporate privacy-preserving mechanisms. Sensitive financial data may be exposed during model training or inference, posing risks to data security.Another limitation is the lack of integration between explainability and privacy. Most systems focus on either improving accuracy or enhancing interpretability, but not both.

IV. PROPOSED METHOD

The proposed system introduces a novel framework that combines explainable machine learning with differential privacy for bank failure prediction. It uses glass-box models such as decision trees and logistic regression to ensure interpretability.Differential privacy is integrated into the model training process by adding controlled noise to the data and model parameters. This prevents leakage of sensitive information while maintaining prediction accuracy.The system is implemented using Django, providing a web-based interface for data analysis and prediction. It allows users to upload financial data, train models, and view results.The proposed system ensures transparency, accuracy, and privacy, making it suitable for real-world financial applications.

V. IMPLEMENTATION

The proposed system is implemented using Python within a Django web framework, providing a scalable and user-friendly interface for bank failure prediction. The system architecture follows a modular design, ensuring separation of concerns between data processing, model training, privacy mechanisms, and user interaction.The implementation begins with the configuration of the Django environment using the manage.py script, which initializes the application settings and enables command-line execution for administrative tasks. The backend is responsible for handling data input,

preprocessing, model training, and prediction generation.The dataset consists of financial indicators such as capital adequacy ratio, liquidity ratio, non-performing assets, return on assets, and leverage ratios. Data preprocessing includes handling missing values, normalization, and feature selection. The preprocessing pipeline ensures that the data is suitable for machine learning algorithms and minimizes noise in the dataset.To maintain interpretability, glass-box models such as Logistic Regression and Decision Trees are implemented using the Scikit-learn library. These models are chosen due to their transparent decision-making processes, allowing users to understand how predictions are generated. Feature importance and decision paths are extracted to provide insights into the contributing factors of bank failure.Differential privacy is incorporated during the training phase. Controlled noise is added to the dataset and model parameters using Laplace or Gaussian mechanisms. This ensures that individual data points cannot be reverse-engineered from the trained model. Privacy parameters such as epsilon are carefully tuned to balance accuracy and privacy.The system includes a web-based dashboard where users can upload datasets, initiate model training, and view predictions. The interface displays performance metrics such as accuracy, precision, recall, and F1-score. Visualization tools are integrated to present feature importance and prediction outcomes.During prediction, the trained model processes new input data and generates a probability score indicating the likelihood of bank failure. The system also provides interpretability by highlighting key features influencing the prediction.Overall, the implementation ensures a seamless integration of machine learning, privacy preservation, and explainability within a robust web application.

VI. ALGORITHMS

The proposed system employs a combination of interpretable machine learning algorithms and differential privacy techniques to achieve accurate and secure bank failure prediction.The primary algorithm used is Logistic Regression, which models the probability of bank failure based on input financial features. The sigmoid function transforms the linear combination of features into a probability score. This method is highly interpretable, as the coefficients indicate the contribution of each feature.Decision Tree algorithms are also utilized to enhance interpretability. The tree structure splits the dataset based on feature thresholds, creating a hierarchical representation of decision rules. Each path from the root to a leaf node represents a classification rule, making the model easy to understand.To improve performance, ensemble methods such as Random Forest can be incorporated. Random Forest builds multiple decision trees and aggregates their predictions to reduce overfitting and improve accuracy.Differential privacy is implemented using noise addition techniques. The Laplace mechanism adds noise proportional to the sensitivity of the function, ensuring that individual data points remain indistinguishable. The privacy parameter epsilon controls the trade-off between privacy and accuracy.The algorithmic workflow includes data preprocessing, feature selection, model training with privacy constraints, evaluation, and prediction. The combination of these algorithms ensures a balance between interpretability, accuracy, and privacy.

VII. SYSTEM DESIGN

The system design follows a layered architecture, consisting of presentation, application, and data layers. This modular approach enhances scalability, maintainability, and flexibility.The presentation layer is implemented using Django templates and provides a user-friendly interface for interacting with the system. Users can upload datasets, configure model parameters, and view results through a web dashboard. The interface is designed to be intuitive, enabling users with minimal technical expertise to operate the system.The application layer handles the core functionality of the system. It includes modules for data preprocessing, machine learning model training, differential privacy integration, and result visualization. Each module is designed as an independent component, allowing for easy updates and extensions.The data preprocessing module is responsible for cleaning and transforming the input data. It handles missing values, normalizes features, and performs encoding for categorical variables. Feature selection techniques are applied to identify the most relevant attributes for prediction.The machine learning module implements glass-box models such as Logistic Regression and Decision Trees. These models are trained using the preprocessed data and evaluated using performance metrics. The module also supports hyperparameter tuning to optimize model performance.The privacy module integrates differential privacy mechanisms into the training process. It ensures that sensitive financial data is protected by adding noise to the data and model parameters. The module is configurable, allowing users to adjust privacy levels based on requirements.The data layer manages storage and retrieval of datasets and model outputs. It uses relational databases to store financial data, training logs, and prediction results. Efficient data management ensures quick access and scalability.The system also includes a visualization component that presents results in graphical formats. Charts and graphs are used to display feature importance, prediction outcomes, and performance metrics.Overall, the system design ensures seamless integration of machine learning, privacy preservation, and user interaction, making it suitable for real-world financial applications.

SYSTEM DESIGN IMAGES

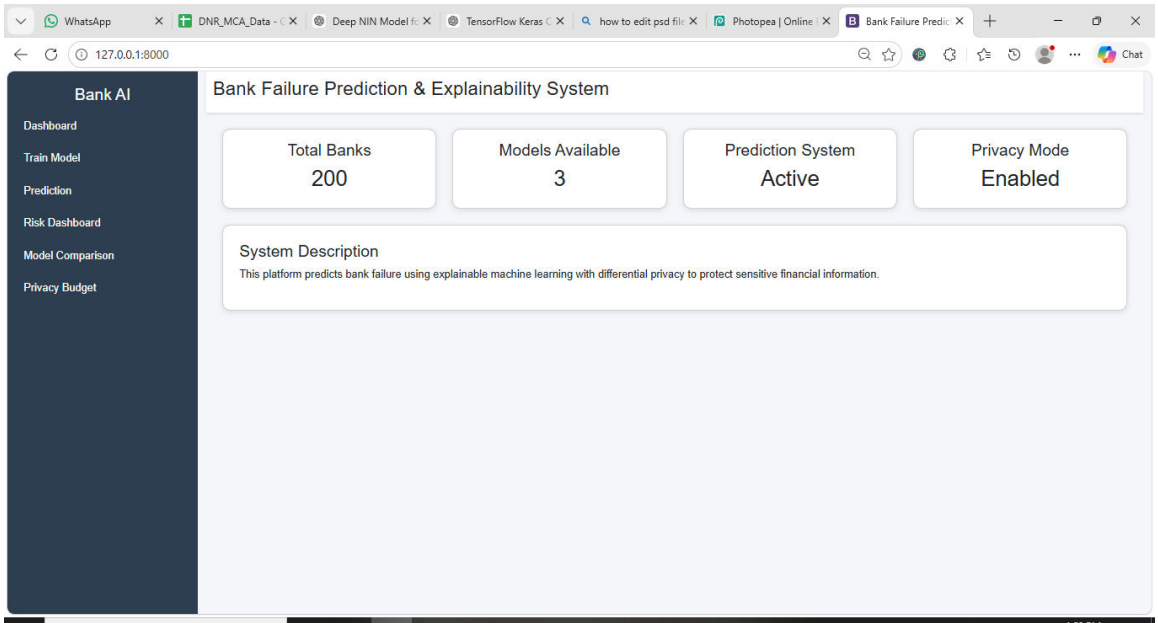

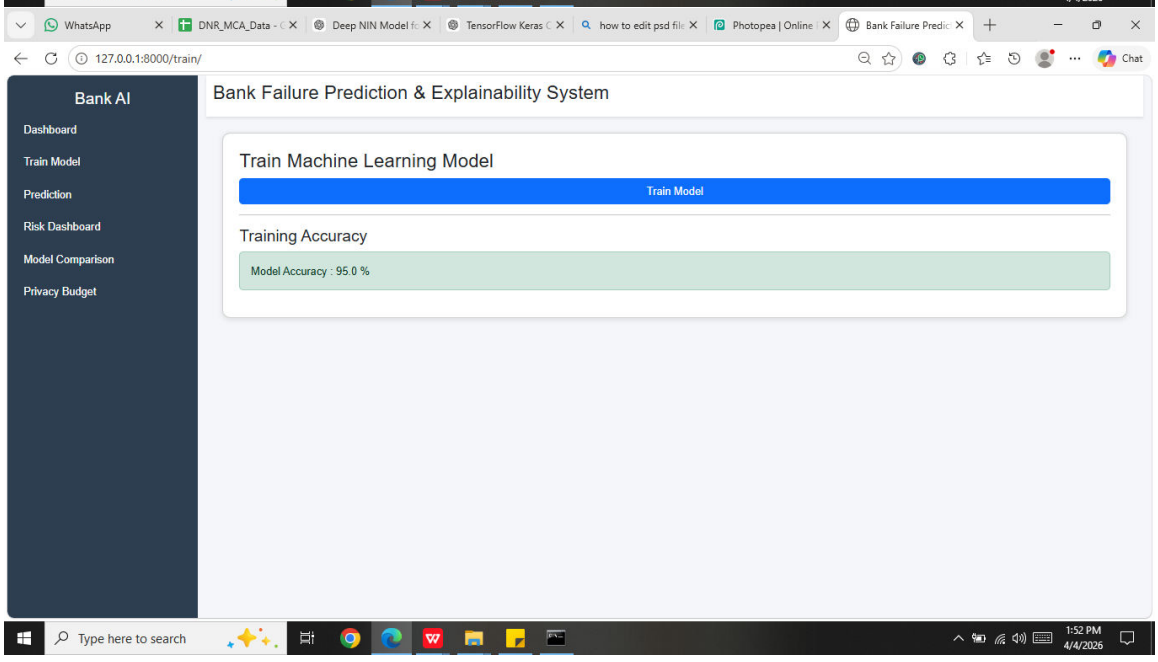

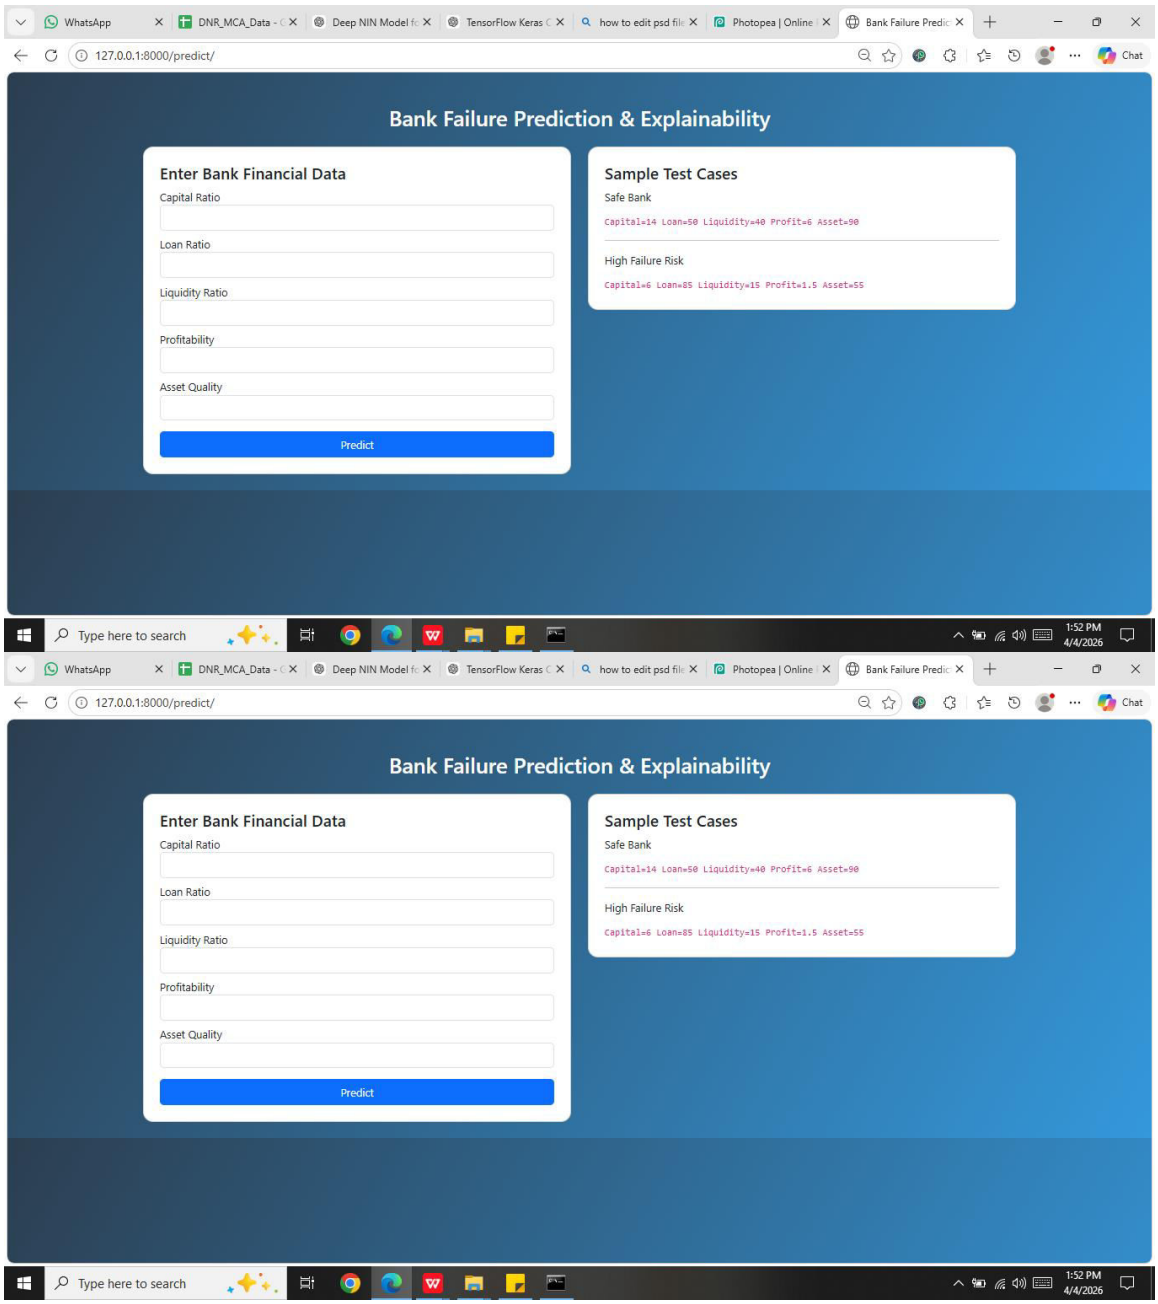
WhatsApp
DNR_MCA_Data - C
Deep NIN Model fc
TensorFlow Keras C
how to edit psd file
Photopea | Online
Bank Failure Predic
127.0.0.1:8000/predict/
Chat
Bank Failure Prediction & Explainability
Enter Bank Financial Data
Capital Ratio
Loan Ratio
Liquidity Ratio
Profitability
Asset Quality
Predict
Sample Test Cases
Safe Bank
Capital=14 Loan=50 Liquidity=40 Profit=6 Asset=90
High Failure Risk
Capital=6 Loan=85 Liquidity=15 Profit=1.5 Asset=55
Type here to search
1:52 PM
4/4/2026

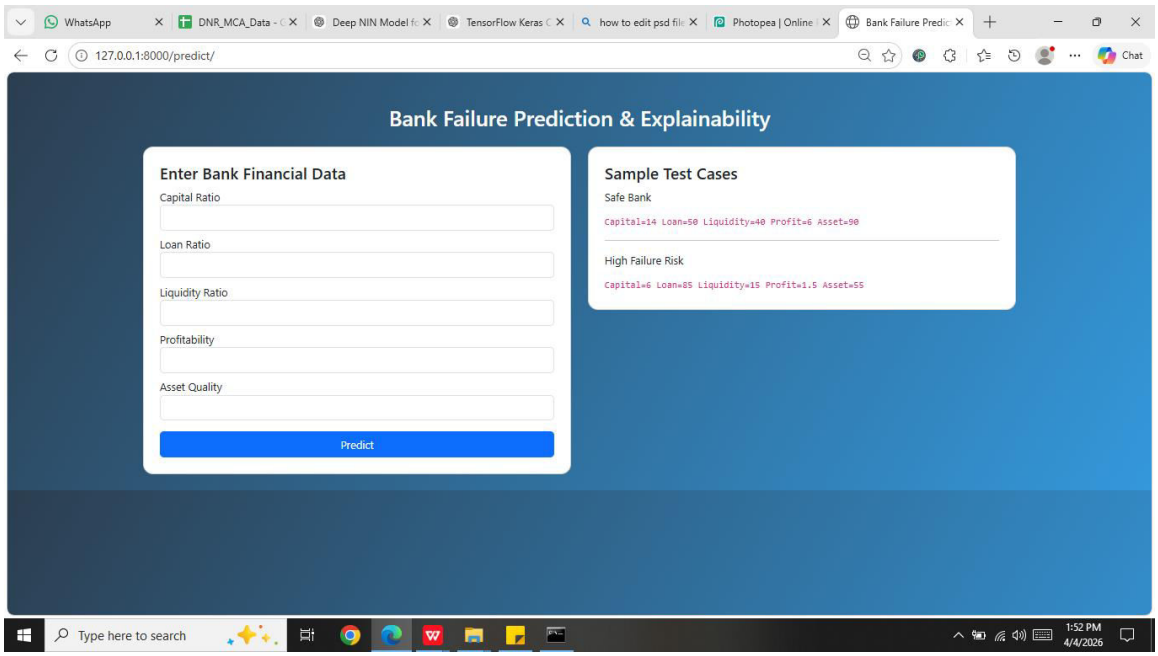

VIII. CONCLUSION

This research presents a novel framework for bank failure prediction that effectively balances explainability and privacy using a differentially private glass-box approach. The proposed system addresses key challenges in financial analytics by integrating interpretable machine learning models with privacy-preserving techniques.The use of glass-box models such as Logistic Regression and Decision Trees ensures that predictions are transparent and easily understandable. This is essential for regulatory compliance and decision-making in financial institutions. At the same time, the incorporation of differential privacy protects sensitive financial data, preventing unauthorized access and ensuring data confidentiality.The system demonstrates that it is possible to achieve high prediction accuracy while maintaining strong privacy guarantees. The integration of a web-based interface further enhances usability, allowing stakeholders to interact with the system and gain insights into financial risks.The proposed framework contributes to the field of financial risk management by providing a secure and interpretable solution for bank failure prediction. It can be adopted by regulatory authorities and financial institutions to improve risk assessment and prevent economic instability.Future research may focus on extending the framework to include advanced machine learning models, federated learning approaches, and real-time data processing. Additionally, further exploration of privacy-preserving techniques can enhance the robustness of the system.

REFERENCES

1. C. Dwork et al., "The Algorithmic Foundations of Differential Privacy," *Foundations and Trends in Theoretical Computer Science*, 2019.
2. R. Guidotti et al., "A Survey of Methods for Explaining Black Box Models," *ACM Computing Surveys*, 2020.
3. A. Adadi and M. Berrada, "Peeking Inside the Black-Box: A Survey on Explainable AI," *IEEE Access*, 2018.
4. J. Chen et al., "Interpretable Machine Learning for Financial Risk Prediction," *IEEE Transactions on Neural Networks*, 2021.
5. F. Doshi-Velez and B. Kim, "Towards A Rigorous Science of Interpretable ML," 2017.
6. L. Breiman, "Random Forests," *Machine Learning*, 2001.
7. T. Hastie et al., *The Elements of Statistical Learning*, Springer, 2017.
8. I. Goodfellow et al., *Deep Learning*, MIT Press, 2016.
9. S. Shalev-Shwartz, "Understanding Machine Learning: From Theory to Algorithms," 2019.
10. N. Papernot et al., "Semi-supervised Knowledge Transfer for Deep Learning," *ICLR*, 2017.
11. M. Abadi et al., "Deep Learning with Differential Privacy," *ACM CCS*, 2016.
12. B. Ribeiro et al., "Why Should I Trust You? Explaining Predictions of Any Classifier," *KDD*, 2016.
13. J. Kroll et al., "Accountable Algorithms," *University of Pennsylvania Law Review*, 2017.
14. S. Barocas et al., "Fairness in Machine Learning," *NIPS Tutorial*, 2019.
15. Y. LeCun et al., "Deep Learning," *Nature*, 2015.