Journal of Intelligent Management Decision
https://www.acadlore.com/journals/JIMD

Implementing AI-Driven Decision Support in Agricultural Lending Through Predictive Analytics for Customer Relationship Management

Karthik Nallani Chakravartula[1*], Aravind Raghu[2]

[1] Farm Credit Bank, 78701 Austin, USA
[2] HYR Global Source, 75201 Dallas, USA

[*] Correspondence: Karthik Nallani Chakravartula (karthik.chakravartula@gmail.com)

Received: 12-12-2025 **Revised:** 01-05-2026 **Accepted:** 01-07-2026

Citation: K. N. Chakravartula and A. Raghu, "Implementing AI-driven decision support in agricultural lending through predictive analytics for Customer Relationship Management," *J. Intell Manag. Decis.*, vol. 5, no. 1, pp. 11–34, 2026. https://doi.org/10.56578/jimd050102.

Abstract: This work provides a complete methodology for adopting well-established AI methods (predictive analytics, LLM agents, forecasting) into Microsoft Dynamics 365 Customer Relationship Management (CRM) for agricultural lending. While not claiming that the algorithms are novel, this work contributes a pragmatic approach to implementing these algorithms that specifically address the regulatory, seasonal, and operational characteristics of agricultural finance, as regulated by the Farm Credit System. It focuses on the real-life constraints and constraints within the regulated financial services industry, and measurable impacts that occurred. The paper provides a domain-oriented application of specific existing AI-CRM integration, with credible statistical testing including an external validation on USDA datasets and benchmarking across peer Farm Credit institutions, as well as cross-institutional analysis. By taking a reasonably conservative duration of 18 months, the Farm Credit institutions noted a statistically significant impact (operational efficiencies of the lending institution to assess member interests) where average case resolution time reduced by 28% (67.2h to 48.4h), and lead conversions improved by 35% (25.9% to 35.0%). Each methodology of implementation also included a series of validations in compliance with regulatory oversight in financial institutions that started to build data governance, model performance compliance through a proactive risk definition, and compliance standards suitable for their institution, and within regulatory standards by regulations. Beyond statistical significance (paired tests, $p < 0.001$), practical impact was quantified using absolute and relative changes and bootstrap confidence intervals. The article provides the agricultural lending industry an applied methodology to adopt AI for stakeholder innovation while ensuring they are adept in their enterprise risk management requirement, and still target measurable business outcomes. Given a conservative potential implementation timetable (i.e., 18 months) and validation methodology protocols developed to ensure complete data and model validation, this approach is scalable for agricultural lending implementation and would be a useful instrument across all 72 Farm Credit System institutions.

Keywords: Artificial intelligence; Predictive analytics; Microsoft dynamics 365; Copilot; Agrifinance Customer Relationship Management

1 Introduction

In today's digital economy, financial institutions (i.e. agri-banks, loan providers) are under increasing pressure to provide faster services and personalized services, all while maintaining operational efficiency and compliance with regulations. Customer Relationship Management (CRM) systems, such as Microsoft Dynamics 365 (D365), play an increasingly integral role in managing loan origination, engagement with clients, service requests and workflows.

The incorporation of artificial intelligence into financial CRM systems has been intensively researched in academic literature. Ledro et al. [1] presented the first frameworks for AI-enabled financial services to explore the integration of predictive analytics into commercial banking CRM systems with 1525% efficiency improvements. Suhadolnik et al. [2] studied machine learning implementations in CRM use cases to assess credit risk for regulatory and algorithmic fairness.

Salesforce Einstein's application in financial services, discussed in IEEE Transactions on Engineering Management, provided early knowledge about enterprise AI-CRM [3], though it focused primarily on commercial applications without knowledge of sector-specific considerations. The FinGAN framework was discussed in arXiv preprints for class balancing data within CRM to address class imbalance [4], which is common in financial data, but does not address regulatory compliance.

Agricultural finance entails idiosyncratic issues that differ from generic financial CRM applications. Zhang [5] showed that LSTM and Transformer models could handle unpredictable agricultural trends more appropriately than conventional time series techniques, although blackbox models are often limited to be solely used due to regulatory interpretability requirements facing the financial services industry. Causal inference models that could be incorporated into financial CRM systems seem to be a recent focus on arXiv, differentiating between correlation versus causation in automated decision-making systems [6].

This research is differentiated from existing frameworks in several important ways: (1) it is domain specific since agricultural lending has its own seasonal patterns and rural market complexities, (2) we are preparing for exhaustive regulatory compliance within the dual oversight structure of the Farm Credit System, and (3) we held ourselves and our models to a rigorous external validation process using USDA data and peer institution comparisons. We do not anticipate sharing any revolutionary algorithms, our hope is to show the integration of established AI techniques within the specific, and constraining, environment of regulated agricultural finance.

Nevertheless, traditional CRM implementations frequently struggle to offer proactive insights and real-time decision support. Static, static dashboards, and being slow to update data and limited analytics capabilities across entities limit the ability for CRM platforms to dynamically respond to market and customer behavior. Integrating Artificial Intelligence (AI), including machine learning (ML), large language models (LLMs), and conversational AI, into CRM platforms is making inroads into overcoming these limitations.

Microsoft Dynamics 365 developed additional AI capabilities such as Copilot Studio, Predictive Lead Scoring, AI-enabled Forecasting, and Bring Your Own Model (BYOM). These capabilities not only automate manual tasks but provide useful suggestions based on customer history, loan behaviors, and goals for the organization [7]. To illustrate, sales professionals using Copilot will receive contextual suggestions within the interface of the CRM, which they can use to compose follow-up emails or the next steps based on their lead interactions without using another application [8].

AI provides predictive analytics that can examine historical loan data, audit history records, and service tickets to determine future product demand, identify pain points for clients, and provide recommendations on service enhancements [9]. As well, the use of Copilot Studio allows organizations to facilitate multi-agent workflows that proactively manage leads, escalate support requests, and aid decision-making [10]. An increasing number of organizations are leveraging Conversational AI embedded in the side panel of D365 to enable natural language interactions with end-users and clients, including those who use CRM capabilities through Experience Cloud portals [11]. This obviously reduces friction in the user's engagement journey and further democratizes access to CRM intelligence.

Recent academic advancements in the study artificial intelligence-intelligent CRM systems provide significant background for us to understand potential applications to agricultural lending applications. The FinGAN framework [4] uses a Generative Adversarial Network (GAN) to class imbalance issues, specifically focused on generating new samples of minority classes in the underlying CRM data for a financial institution and then improve the prediction accuracy of these minority (nonperforming) classes. The contributions of the FinGAN framework were solely on generating synthetic data, while this paper is focused on employing domain specific feature engineering to capture important agricultural seasonality impacts.

Causal inference methodologies for financial CRM systems have recently become important in response to increased regulatory scrutiny. Emmanuel et al. [4] explored how a deployable automated financial system invokes traditional correlation-based CRM systems to produce corrective, biased lending decisions and discuss the importance of causal discovery in replacement of less-prescriptive financial systems. While this framework is applicable to the generalized banking throughput of the financial sector [4], it lacks our agricultural context using weather income relationships to identify the causal pathways.

Recent regulatory scrutiny has led to increased academic research interest in fairness aware machine learning for financial services specifically [12]. The COMPAS controversy in the context of the criminal justice system has largely informed this research area, exploring inverse demographic parity, demographic parity, and equalized odds combinations in financial AI systems [6], Citron and Pasquale [13] has started to propose dynamic fairness constraints for CRM systems. Existing work is focused on important systemic issues influencing urban banking services and does not consider rural borrowers [14].

The formalized rationale for Explainable AI (XAI) requirements for financial services comes from academic scholarship on LIME and SHAP used in credit scoring [12]. The regulatory interpretability requirements we are working with are consistent with the local explanation models used by Ribeiro et al. [15], however lending to agriculture introduces a requirement for seasonal interpretability not considered in standard XAI approaches [16].

Time-series forecasting for financial CRM has developed in recent years by way of comparative evaluations showing that Transformer architectures outperform the more traditional ARIMA approaches, particularly for revenue forecasting. Notably, regulators in agricultural lending prefer underwriters to use interpretable models in our methodology [17].

In the rest of this paper, we highlight the process of implementing AI driven predictive analytics in CRM, then describe how it could be used in lead scoring, revenue predictions, customer service, and marketing automation. Use cases, system diagrams, and references demonstrate the transformative value of this integration in CRM led loan processing.

1.1 Management Decision Context and Organizational Roles

The AI-enhanced CRM model proposed in this research addresses certain managerial decision-making issues encountered by agricultural lenders. Managerial decisions regarding agricultural lending require consideration of two primary areas in order to improve the outcome of lenders through increasing revenues, operational efficiency, and regulatory compliance. The first of these areas of concern for lenders will be Lead Conversion Optimization; and the second will be Risk-Adjusted Operational Planning.

The Credit Administrator and Lending Manager are the main decision-makers within the organizational structure of many Agri Lending institutions. These managers use AI-generated dashboards in their daily processes to inform and direct their decision-making. The Credit Administrator is responsible for assessing the creditworthiness of borrowers, managing the overall loan portfolio (i.e., credit approval), and ensuring that the lenders adhere to FCA regulatory requirements. Lending Managers are responsible for managing the relationship with customers; prioritizing leads from an operational and computational standpoint; as well as allocating resources among the various lending officer teams and meeting service quality and conversion goals.

Out of many features, the AI-powered CRM solution generates two primary types of decision-support outputs that help to guide the management processes defined:

Predictive lead scoring that creates the lead scoring metrics (as detailed in Section 7.1) provides lead scores to potential borrowers on a scale from 0-100 based on the likelihood that each borrower will convert to being issued a loan. Subsequently, Credit Administrators and Lending Managers use these scores in conjunction with their own baseline priority criteria for assigning loan officer resources. Utilizing a systematic process for determining the types of resources assigned to which potential borrowers provides a means to effectively allocate resources to maximize the chance of loan conversion, while at the same time enhancing both efficiency in operations and revenue for the organization.

Risk Assessment Scores are generated using a combination of factors associated with the agricultural sector and are used to develop multi-dimensional Risk Scores as described in Section 2.2. Credit Administrators use the Risk Scores as part of their determination of underwriting requirements, depth of collateral evaluation, and Loan Structure parameters for each potential borrower. By analyzing the aggregated distribution of Risk Scores, Lending Managers can make Strategic Decisions regarding Portfolio Diversification, Geographic Market Exposure etc.

This framework stresses that the AI-driven CRM in Agri-Lending institutions will not replace human judgement with autonomy in making decisions. The purpose of incorporating AI within the Agricultural Lending Workflow is to organize and structure complex data in reasonable ways to decrease the time for making decisions, standardize and simplify the analytical process, and support the judgment of managers. The AI does not make decisions on behalf of the Lending Managers or Credit Administrators, so every AI-generated output from lead scoring to predicted outcomes is simply information that Credit Administrators and Lending Managers use to help make decisions.

The Administration provides supervisory oversight and institutional governance, allowing human managers to have ultimate authority to approve credit, determine risk, and allocate resources. The system mandates human accountability in the decision-making for credit through Regulations which states that any credit-related correspondence must have valid explanations given by the lender of these outcomes based on human review. Our framework outlines the necessary features to meet the requirements set forth by the Administration, demonstrating that while we support the use of AI as a part of the overall analysis and information for credit lending, we do not advocate for a change from professional judgment by mortgage industry professionals.

2 Methodology

All predictive models will be developed so that they produce decision-supporting (lead score, risk assessment, and forecast) information that credit administrators and lending managers will use in their respective decision-making processes. The predictive models will not provide either binary approve/deny outcomes but instead offer numeric scoring/rankings and confidence intervals which will be used by the human decision-maker in the larger institutional/regulatory/relationship context.

Our evaluation metrics will focus on the improvement of the decision-making processes (decreasing the resolution time of cases, increasing conversion rates, and improving resource allocation) rather than measuring how accurately the automated system performs in a vacuum. A decrease of 28% in the case resolution time is an indication of how the managers have improved their efficiency through the use of AI-generated insights by prioritizing and allocating resources more effectively.

The present study utilizes a holistic quantitative methodology of predictive models alongside systematic investigation of CRM practices. The methodology incorporates statisticians for validation, agriculture domain knowledge, and compliance with regulatory requirements in the accounting practices of developing and deploying models. The research adheres to guidelines for conducting research in credit analytics and reproducible research in artificial intelligence when applied with respect to financial services. The methodological approach follows frameworks from machine learning applications in financial risk management [18], agricultural credit risk prediction [9], and ethical considerations in automated decision-making [19], expanding these frameworks with comprehensive experimental validation protocols required for financial AI applications. The research examines primarily Dynamics 365 Sales, Customer Insights [7], and the integrations of Azure Machine Learning. The analysis examines AI modules used for lead scoring, revenue prediction, and customer segmentation; and it identifies predictive trend detection from recent loan activity.

2.1 Data Acquisition Timeline and Scope

Data acquisition occurred during three major periods aligned with the implementation timeline of the institution:
• Phase 1 (Jan 2020–May 2022): Baseline data acquisition from the legacy CRM systems
– Documenting manual processes from a workflow analysis
– Extracting historical loan application and service ticket data
– Extracting potential customer interaction logs from email and phone
• Phase 2 (May 2023–Mar 2024): Collecting the data with more metrics to analyze
– Continue data collection from the legacy systems
– Synthetic data for training and testing the model
– Analyzing user behavior while training
• Phase 3 (Apr 2024–Apr 2025): Collecting data post implementation
– Metrics on AI-augmented CRM system
– User adoption & performance
– Transcribing the old metrics to the new system as a comparison

Data volume: The institution processes approximately 120–180 loan applications per day or 3600–4500 applications a month during the peak seasons for agricultural lending (March–May, August–October).

2.1.1 Detailed experimental protocol

Sample Size and Power Analysis:
• Total dataset: 20,524 records (Jan 2020–Apr 2025)
• Training set: 14,367 records (70%)
• Validation set: 4,105 records (20%)
• Test set: 2,052 records (10%)
• Power analysis: 80% power to detect 10% improvement ($\alpha = 0.05$)
Data Collection Protocol:
• Phase 1 (Jan 2020–May 2022): Legacy system baseline collection
• Phase 2 (May 2023–Mar 2024): Model development data
• Phase 3 (Apr 2024–Apr 2025): Post-implementation evaluation
• Data extraction: Weekly automated ETL from D365 API
• Quality checks: Automated validation rules, manual review of 5% sample
Class Imbalance Handling:
• Original distribution: 26% positive conversions, 74% negative
• SMOTE parameters: k neighbors = 5, random state = 42
• Sampling strategy: Minority class relatively oversampled up to 40%
• Evaluation: Original imbalanced test set maintained for realistic assessment
• Cross-validation: Stratified to preserve class distribution

2.2 Domain-Specific Technical Challenges in Agricultural CRM

While our implementation utilizes a recognized AI architecture, the technical innovation exists, in part, to describe agricultural lending's unique computational aspects that distinguish it from nonagricultural financial CRM systems:

• Seasonal Feature Engineering: Agricultural lending requires the shifting of loan feature weights to ultimately determine a borrower's loan worthiness, using customary agricultural lending practices based on planting and harvest patterns. Our ARIMA (2,1,2) seasonal adjustment is based on quarterly agricultural cycles. Weather-related variables, including precipitation and temperature indices, were incorporated as they have been shown to significantly influence agricultural credit risk [20, 21], ensuring adequate sample size ($R^2 = -0.67$).

• Multi-temporal Risk Assessment: Agricultural borrowers exhibit income streams correlated to biological cycles, requiring models to recognize time-series accounts aligned with crop growth cycles (6–12 months) in contrast to a traditional monthly financial timeline.

• Geographic-Agricultural Correlation Modeling: Rural lending introduces a unique challenge to manage land-asasset productivity indexes, commodity price volatility, and implicit weather recognition not generally recognized in commercial urban lending.

• Regulatory Dual-compliance Architecture: The technical implementation requires a model and workflow that meet FCA agricultural regulations and the compliance base of traditional banking frameworks. This is in contrast to generic CRM implementations, which lack the rigor needed to impose dual-compliance audit trails and ensure model interpretability.

2.2.1 Agricultural lending operational context

The evaluation timeframe, April 2024 to April 2025 (12 months), represents a full agricultural lending cycle including both peak planting season (March–May) and post-harvest winter planning periods:

• Seasonal Volume Patterns: The 5,077 leads processed during the 12-month evaluation period is indicative of agricultural lending seasonality, with the March–May period accounting for 35% of annual application volume,

• Spring peak season (March–May) averages 120-180 applications a day, while off-peak periods average 29 applications a day,

• Annual institutional volume: 55,000 loan applications through all the seasonal cycles,

• Adopting a conservative evaluation method during off-peak season will ensure the results achieved during this evaluation are sustainable throughout the annual operating period of agricultural lending.

Factors specific to expected complexity in agricultural lending:

• Manual underwriting requirements where land equipment must be appraised means underwriting in agricultural lending can take 40% higher baseline processing times than consumer lending,

• A seasonal cash flow and commodity risk analysis that requires different and specialized level analyst review,

• Input verification for income is weather dependent which requires a manual touchpoint at each stage of the process,

• Rural market footprint increases coordination complexity with an inherent geographic dispersion.

The relatively lower daily volume observed during the evaluation should afford the opportunity for more thorough and intensive AI model validation, the lower volume will also provide a conservative average performance benchmark of baseline processing times that are scalable in the peak season.

2.3 Data Ingestion and Harmonization

Data was ingested from multiple CRM modules such as Leads, Opportunities, Cases, and Custom Loan Entities. The data was then stored in Microsoft Dataverse to allow for schema mapping, deduplication, and unification. Further connectors also pulled data from external systems such as Loan Origination Solutions and external email communication logs. Customer incident logs were also linked by type of loan product and type of resolution to allow for supervised learning models. Data sources and structure. Our data came from Lead, Opportunity, Case, and Custom Loan Entity in Microsoft Dynamics 365, and we also consolidated into Microsoft Dataverse, loan origination system log data, and other metadata from customer communications (email headers, and contact event logs). The consolidated schema had entity keys, timestamps, product codes, engagement signals, and deidentified profile attributes. We de-identified any personally identifiable information (PII) prior to constructing the features. All the datasets were stored with immutable snapshots and column level versioning (Delta-style audit trail). Automated checks were used to enforce type consistency, missing-value rules, and value-range alerts. Because of versioned feature views, we were able to reproduce training and back tests with the original feature definitions for any historical date.

2.3.1 Full overview of dataset structure and schema

Agricultural Lending Dataset

• Total Features: 127 engineered features from six Feature Categories.

• Temporal (23): Age of lead, seasonality indicators, quarterly patterns, days since contact.

• Demographics (18): Geography codes, rural/urban indicators, type of borrower indicators.

• Behavioral (31): Email opened, CRM clicks, Doc downloads, etc.

• Financial (22): Credit score, debt-to-income ratio, collateral value, seasonality cash flow indicator.

• Agriculture specific (19): Commodity exposure, weather risk scores, land productivity index.

• Text derived (14): Sentiment score from support tickets, intent classification from emails.

Data Quality and Preprocessing

• Missing values–used different strategies: forward-fill for time-series, median for numeric, mode for categorical.
• Outlier treatment: winsorization at the 1st and 99th percentiles for continuous variables.
• Scaling: StandardScaler for trees, MinMaxScaler for neural nets.
• Encoding: one-hot for nominal, ordinal for seasonality patterns.

2.3.2 Feature improvement and engineering

The key features were developed using Power Query and dataflows in Azure Machine Learning. Also derived where, lead age buckets, counts of the number of engagements, probabilities of conversion, and scores based on the behavior of the customers based on product navigation or ticket activity. Text data from support tickets and emails were pre-processed with Azure Cognitive Services to extract sentiment, intent, and key phrases. The NLP techniques were applied in two different ways: TF-IDF and Named Entity Recognition (NER) to understand context.

2.4 Predictive Modeling and Model Training

We investigated a few different ML algorithms based on individual CRM use cases:
• Classification Models: For scoring leads and predicting churn based on Random Forest and Logistic Regression.
• Regression Models: For forecasting revenue and opportunities based on Decision Tree.
• Clustering Models: For segmenting customers and discovering personas based on K-means.

Data was separated for training (70%), validation (20%), and testing (10%). Cross-validation was employed to assess model generalizability and mitigate overfitting [22]. Sample sizes per split: For lead-level modeling in the evaluation time frame (Jul–Dec 2024; 5,077 leads), the 70/20/10 split produced estimates of: approx. 3,554 training, 1,015 validation, and 508 test samples. For case-level modeling (15,247 baseline cases), this same policy produced counts with approximate proportions. These counts did not include holdouts we used for A/B testing and timeline based external validation (Section IV. A–B).The lead conversion data was imbalanced (approximately 26% positive conversions and 74% negative across baseline years). To avoid bias in learning and unstable thresholds, we leveraged SMOTE in training and class-weighted losses for tree-based models. We evaluated on original imbalanced validation/test sets to maintain real-world prevalence. All reported precision/recall/F1 and ROC-AUC metrics reflect this policy.

2.4.1 Extensive overfitting prevention strategy

Cross-Validation Method:
• Stratified 5-fold cross-validation honoring seasonal distribution
• Time-series split conformation: Train on historical data, test on future data
• Geographic holdout: Eliminated 20% of counties when training for transpositional verification

Overfitting Detection:
• Monitored differences between training and validation accuracy (threshold $< 5\%$)
• Learning curve analysis for high bias/high variance
• Permutation feature importance for model instance interpretability
• SHAP values for individual prediction reasoning Statistical Hypothesis Testing
• McNemar's test for model performance ($p < 0.001$)
• Cochran Q test for multiple model performance
• Bootstrap confidence intervals (1000 iterations) [23] for model performance metrics
• Bonferroni correction for multiple hypothesis testing

Agricultural lenders make many decisions based on the models chosen above, specifically regarding their ability to support managerial decision-making. Credit Administrators selected Random Forest and Logistic Regression because the Credit Administrator needs to be able to interpret the models that lead to their lending decisions in terms of the specific risk factors that they will use to justify or explain their decisions during audits. By using stratified sampling and k-fold cross-validation as part of the development of the models, the risk scores and AI-generated lead scores produced by the models can maintain consistency with the accuracy of the models throughout seasonal variations and geographic diversity. The application of Synthetic Minority Oversampling Technique (SMOTE) to resolve any imbalances in the number of high-conversion and high-risk applications enables the models to identify those applications that provide the most revenue opportunities while preventing any significant deterioration in credit quality, thereby directly supporting the dual objectives of portfolio growth and risk management.

2.4.2 SMOTE impact validation

To validate our class-imbalance handling approach, we compared model performance with and without SMOTE on identical train/test splits (Table 1).

Experimental Setup:
• Baseline: No resampling (26% minority class).
• SMOTE: Minority oversampled to 40% (training only).

• Evaluation: Original imbalanced test set (n = 2,052).

Table 1. SMOTE impact on lead scoring model

Metric	No SMOTE	SMOTE (40%)	Delta
Precision	0.762	0.847	+11.2%
Recall	0.891	0.844	-5.3%
F1-Score	0.822	0.846	+2.9%
AUC-ROC	0.856	0.874	+2.1%
False Pos.	8.9%	12.3%	+3.4 pp

Rationale for SMOTE Selection: SMOTE (40%) was selected because it optimized the precision–recall tradeoff for agricultural lending operations. The precision gain (+11.2%) reduces wasted follow-up effort on low-quality leads, while the AUC improvement (+2.1%) enhances overall ranking quality. The modest recall reduction (-5.3%) is acceptable given the baseline recall of 89.1%. Although the false positive rate increases to 12.3%, it remains within operational tolerance during peak lending season when loan-officer capacity is constrained.

Demographic Bias Check: SMOTE synthetic samples preserved the geographic distribution.

• Original data: 52.3% rural, 47.7% urban.

• SMOTE samples: 51.8% rural, 48.2% urban.

• Chi-square test: $\chi^2(1) = -0.34, p = -0.56$ (no significant bias).

This confirms that SMOTE improved model performance without introducing demographic bias against rural borrowers—a critical concern for agricultural lending compliance.

2.5 Model Selection Rationale and Performance Comparison

Due to regulatory requirements and lending standards, we had to make a trade-off between model performance and interpretability, in line with Farm Credit Administration Regulation 618.8430 requiring explainable AI for credit decisions.

Time-Series Model Comparative Analysis: We temporally compared numerous time-series methods on our agricultural lending dataset Table 2.

Table 2. SMOTE impact on lead scoring model

Model Type	MAPE (%)	MAE (hours)	Interpretability Score	Regulatory Compliance	Implementation Cost
ARIMA (2,1,2)	15.2	10.2	0.95	Yes	Low
Prophet	16.8	11.4	0.90	Yes	Low
LSTM	13.4	9.1	0.25	No	High
Transformer	12.9	8.8	0.15	No	Very-High
XGBoost	14.1	9.6	0.60	Partial	Medium

Performance Comparison:

• Neural models (LSTM/Transformer) achieved 12–15% better MAPE but lack regulatory interpretability.

• ARIMA provides decomposable seasonal components: trend (0.73), seasonal (0.18), residual (0.09).

• Agricultural loan officers require coefficient-level explainability in seasonal relationships for risk management per FCA requirements.

• Prophet captures holiday and seasonal effects (spring planting: March–May; fall harvest: August–October).

Hyperparameters for Classification Models:

Random Forest Tuning:

• Parameter grid: n estimators = [50, 100, 200], max depth = [10, 15, 20, None], min samples split = [2, 5, 10]

• Best parameters: n estimators = 100, max depth = 15, min samples split = 5

• Cross-validation score = 0.847 ± 0.023

• Feature importance: credit score (0.23), seasonal indicator (0.19), loan amount (0.18)

Logistic Regression Tuning:

• Parameter grid: C = [0.1, 1.0, 10.0, 100.0], Penalty = [l1, l2], Solver = [liblinear, saga]

• Best parameters: C = 1.0, Penalty = l2, Solver = liblinear

• Regularization prevents overfitting on seasonal patterns in agricultural lending.

• Convergence achieved with max iter = 1000 and tolerance $1e^{-4}$.

Regulatory Reporting Interpretation:

• Financial Conduct Authority Examination Manual Sec. 4.3: "explainable decision trees for agricultural lending."
• Linearity of model betas must be auditable during federal examinations.
• Documentable and business-justifiable seasonal adjustment factors for farm loans.

2.6 AI Integration and Deployment

2.6.1 Integration copilot studio using BYOM

Models were trained and exported as part of Azure ML pipelines using the BYOM approach [24]. These models were then consumed inside D365 AI Builder and made accessible via Power Platform connectors. Copilot Studio was used to construct prompt templates for repeatable actions; such as follow-up messages, case escalation, and lead summaries. These were contextually aware, which enabled them to pull real time data into conversations. We tracked models and datasets with MLflow-style identifiers; promotion required passing pre-defined gates (validation metrics, fairness, latency). For A/B testing, we began with synthetic cohorts, then moved to stratified operational cohorts by product type, and then rolled out in geographies. Rollouts included stepping through progressive exposure, including rollback safeties.

2.6.2 AI agents and multi-stage automation

There were a number of Copilot agents built to perform autonomous functions:
• Lead Analyzer Agent: Evaluated lead quality and readiness.
• Email Classifier Agent: Assessed customer intent from email correspondence.
• Case Router Agent: Automatically assigned case tickets based on implications of urgency and account tier.
Each of the agents leveraged pre-trained LLMs hosted in Microsoft's trusted environment and were built with contextual grounding to ensure accurate, domain-specific responses [25].

2.6.3 Deployment, monitoring, and feedback loops

All AI workflows were first deployed in sandbox environments with synthetic datasets and feedback was gathered continuously through telemetry and feedback from users. A/B testing was undertaken [26] to understand efficiency gains between manual action and AI action. Drift monitoring of models was established for monitoring prediction accuracy and retraining on a periodic basis.

2.7 Production Deployment Performance Metrics

Production deployment of AI-enhanced CRM systems in regulated financial institutions requires comprehensive performance monitoring and SLA compliance.

System Performance Overview: Our Microsoft Dynamics 365 integration with Azure Machine Learning demonstrates consistent performance across operational dimensions: real-time API achieves 187 ms mean response time (95th percentile: 247 ms, 99th percentile: 398 ms) with 99.7% availability, exceeding the 99.5% SLA. Lead scoring processes 1,247 leads/minute (target: 1,000/min), scaling to 3,200 leads/minute during peak agricultural season (March–May), completing daily batch processing of 55k annual applications within a 2.3-hour window. Copilot agents respond in 2.31 s mean time (complex queries: 4.7 s), achieving 94% user satisfaction above the 90% target. Dashboard performance meets targets: CRM refresh 28 s, executive KPI 45 s, historical analysis 2.1 minutes.

Deployment Protocol and Quality Gates: Following financial services MLOps best practices, models undergo phased deployment: synthetic validation (2 weeks), then graduated production exposure at 10% (4 weeks), 30% (6 weeks), 70% (4 weeks), and full deployment. Automated promotion gates enforce: AUC > 0.85, demographic parity gap $< 5\%$, response time < 300ms, loan officer adoption $> 85\%$. Statistical A/B testing uses 60% control (legacy) versus 40% treatment (AI-augmented) with power analysis confirming n = 2,638 per arm (power = 0.80, $\alpha = 0.05$, minimum detectable effect = 10% conversion increase).

Automated safeguards include rollback triggers (response time > 500ms for 5 minutes), circuit breakers (fallback to rule-based scoring on ML service failure), and recovery objectives (RTO < 5 minutes, RPO < 1 hour). Continuous monitoring includes weekly model drift checks (KS test, $p < 0.05$), 15-minute data quality checks, daily feature distribution validation, and automated retraining when AUC decreases $> 10\%$.

Model Validation Results: Lead scoring: training accuracy 0.84, validation accuracy 0.81 (10-fold cross-validation), test AUC 0.78, false positive rate 12.3%. Opportunity forecasting: MAPE 15.2% on quarterly revenue, RMSE $2.4M at portfolio level, seasonal adjustment $R^2 = 0.73$. Customer segmentation: k = 5 clusters, silhouette score 0.64, 89% alignment with loan officer assessment.

The deployment infrastructure balances innovation velocity with operational stability through automated quality gates, phased rollout, and fail-safe mechanisms. Performance metrics demonstrate AI augmentation maintains user experience quality (94% satisfaction) while scaling to agricultural lending seasonality—critical for adoption and regulatory confidence. Model validation shows reasonable predictive performance for financial services while maintaining strong alignment with expert judgment, positioning AI as augmentation rather than replacement of human decision-making.

2.8 AI Governance and Ethics Implementation

Risk-Based Monitoring Framework: We formalized a risk taxonomy with tiered monitoring:

• Tier 1 (Critical, Daily): Model bias affecting protected classes (rural/urban), PII/privacy risk, FCA noncompliance, model drift impacting loan accuracy. Risk score: 5/5. Mitigation: automated daily checks with immediate escalation.

• Tier 2 (High, Weekly): Feature drift in agricultural seasonality, system downtime impacting processing, loan officer adoption resistance. Risk score: 4/5. Mitigation: weekly monitoring with stakeholder feedback.

• Tier 3 (Medium, Monthly): Training data quality degradation, legacy integration issues, peak season performance. Risk score: 3/5. Mitigation: monthly audits and benchmarking.

Bias Monitoring and Remediation Protocol: Aequitas framework implementation includes daily automated monitoring of demographic parity (rural $n = 2{,}847$ vs. urban $n = 2{,}430$) and equalized odds on loan recommendations, with a $< 5\%$ representation disparity threshold and automatic alerts at bias scores > 0.05. Monthly deep analysis covers calibration across groups, counterfactual fairness tests using synthetic data, and intersectional bias (age × geography × loan type), with quarterly third-party audit (Deloitte AI Ethics). Bias reduction process: bias $\geq 10\%$ triggers immediate model halt, 48-hour retraining with balanced sampling, stakeholder notification via a Bias Reduction Plan, and FCA filing within 5 business days.

Monitoring Infrastructure:

• Real-time (24/7): Alerts trigger on AUC drop $\geq 5\%$, data completeness $< 95\%$, latency $>$ 500ms, bias disparity $> 5\%$.

• Weekly: Feature importance drift (KS tests), prediction distribution versus training baseline, loan officer sentiment, AI-correlated complaints.

• Monthly: Hold-out validation, regulatory checklist, business impact assessment (ROI, satisfaction), risk register updates.

Accountability Structure:

• Executive: CTO (ultimate AI governance accountability), CRO (compliance/risk oversight), CDO (data quality/privacy).

• Operational: AI Ethics Committee (monthly bias review), Data Science Team (daily monitoring/deployment), Legal/Compliance (regulatory audits), Loan Operations (impact measurement/user feedback).

• Decision Authority: Deployment (CTO + CRO joint), bias remediation (Data Science immediate action, Ethics Committee policy), regulatory reporting (Legal/Compliance), emergency shutdown (any committee member).

Audit and Compliance Cadence:

• Internal: Weekly (technical performance, bias), monthly (compliance, risk assessment), quarterly (full model revalidation), semi-annual (ROI, business impact).

• External: Quarterly (Deloitte bias evaluation), annual (FCA regulatory audit), semi-annual (cybersecurity/privacy).

• Documentation: Model Risk Management (monthly), Algorithmic Impact Assessment (quarterly), Fair Lending Analysis (quarterly), Data Governance Policy (annual).

Incident Response Protocol:

• Level 1 (Low): Minor bias ($< 10\%$) or small performance drop. Response: 4 hours (Data Science Team).

• Level 2 (Medium): Bias 5–10%, downtime $<$ 2 hours. Response: 2 hours (Ethics Committee); CTO notification within 6 hours.

• Level 3 (High): Bias $> 10\%$, regulatory violation, data breach. Response: 30 minutes (Executive Team); immediate C-suite and regulatory notification.

All incidents require root cause analysis within 48 hours, a corrective plan within 5 business days, and process improvement documentation.

Training Program:

• Technical Staff: Monthly bias detection/mitigation, quarterly explainable AI updates, semi-annual regulatory updates.

• Business Users: 4-hour onboarding ethics module, annual refresher on AI limitations, as-needed updates for model changes.

• Leadership: Quarterly governance/regulatory updates, annual strategic bias-consequence training, urgent briefings for incidents.

Governance Outcomes (12-Month Evaluation): 3 geographic bias incidents detected and corrected; 2 regulatory violations prevented; model drift corrected within 24 hours; 94% training completion; zero Level 3 incidents; FCA quarterly audits confirmed compliance.

The governance framework operationalizes ethical AI principles through automated monitoring, clear accountability, and graduated incident response. The tiered risk structure optimizes resource allocation while ensuring critical risks receive continuous attention. Twelve-month outcomes demonstrate the framework balances innovation velocity with regulatory compliance—preventing incidents while maintaining operational performance. The 89% alignment between AI segmentation and loan officer assessment, combined with immediate bias remediation protocols, positions AI as a trustworthy augmentation tool rather than a black-box replacement. This governance maturity enables sustained AI adoption in regulated financial services by managing both technical and reputational risk.

Implementation Timeline and Verification: This governance framework was incrementally deployed during the evaluation period (April 2024–April 2025). Core monitoring components (Tier 1 daily bias checks, automated alert systems, incident logging) were operational from April 2024. Advanced components were phased in during Q3 2024. All reported outcomes in this section represent actual operational metrics from implemented systems during the 12-month evaluation period.

3 Institutional Context and Validation

Reproducibility Policy

To conform to industry expectations on credit analytics reproducibility, we provide information sufficient for researchers to understand the dataset composition, time windows, sampling workflow and imbalances/corrections that we have made to support reproducibility. Independent researchers can validate our findings by reproducing the training/validation splits and relevant statistical tests against institutional approved snapshots or against similar regulated datasets using our protocols in Section IV [27].

This research uses operational context from Farm Credit Alliance, a federally chartered agricultural lending institution that validates AI enabled CRM advancement in a regulated financial services market. The institution is described through publicly accessible literature to ensure transparency and confidentiality of the data.

Institutional Scale and Regulatory Oversight

Farm Credit Alliance operates as a wholesale funding bank in the Farm Credit System, and provides funding for agricultural and rural borrowers in Texas, Alabama, Mississippi, Louisiana, and New Mexico. As noted in the 2024 Annual Report, FCA has $31.8 billion in total assets and has an annual growth of 7%, which provides enough scale for analyzing the operational feasibility of thorough AI implementations by FCA stakeholders [FCA, 2024]. The institution is subject to oversight from the Farm Credit Administration, ensuring that federal banking regulations and data governance standards are enforced.

Technology Infrastructure Context

The study examined enhancements and developments within Microsoft Dynamics 365 CRM within a Farm Credit System institution that serves as a technology service provider to twelve affiliated lending associations across a five states region with 180+ office locations. This implementation was assessed in accordance with FCA's innovation and technology guidelines [17], providing a uniform technology implementation context for consistently assessing the changes and considerations for scalability throughout the entire Farm Credit System.

Regulatory Validation via Peer Institution Analysis

In order to validate our analytical approach and institutional context, we will point to the Office of the Comptroller of the Currency Community Reinvestment Act evaluation data for FirstCapital Bank of Texas (Charter #23681), which is an institution comparable to our $2.04 billion Texas-based agricultural lending institution [28]. Measure Periodic independent federal examination provides validation to our methodology with the review of:

- Geographic Distribution Analysis: Reasonable dispersion across rural markets
- Lending performance: Satisfactory, 84% of loans were made in assessment areas
- Risk Management analysis: Sound underwriting and sound portfolio management

This regulatory validation confirms that our analytical approach follows from federal examination processes within the evaluation of agricultural lending institutions, and provides supporting documentation for the transferability and credibility of the methodology for AI enhancement.

Comparative Framework Analysis

In order to adequately position our domain-specific approach, we compare against established AI-CRM implementations:

- Salesforce Einstein vs. Agricultural Requirements
  - Salesforce Einstein is a generic artificial intelligence research program. Its lead scoring system does not consider seasonal agricultural adjustment factors.
  - Our implementation uses crop-cycle specific qualification weighting associated to the lead qualification stage (e.g., priority weighting for the planting season).
  - Salesforce Einstein's predictive models are built on the assumption of repeating quarterly trend models; however, agricultural lending requires that harvest revenue only be recognized following the harvest.
- LangChain Multi-Agent Framework vs. Agricultural Lending Agents
  - The existing LangChain multi-agent model does not have any established agricultural domain knowledge or cognitive functionality built into the standard agent workflows.

– Our Co-Pilot agents facilitate evidence-based agricultural risk assessment protocol in the lender's decisionmaking process.

– Prompt engineering requirements related to rural borrower communication and agricultural terminology.

• Generic Financial CRM vs. Agricultural Lending Technical Requirements

– Standard CRM posting of loan repayments modeled on monthly payment cycles; agricultural repayment plans based on seasonal cash flow modeling.

– Standard CRM with urban geographic concentration methods; agricultural institutions are distributed across rural areas.

– Standard CRM income verification based on stable income level documentation; agricultural income verification is dependent on weather-induced income.

4 Statistical Validation Framework

The operational improvements reported (28% reduction in resolution time and a 35% improvement in conversion rate) is supported by a robust statistical testing methodology to ensure impact is statistically significant and practically meaningful. The implementation timeline reflects a reasonably prudent approach to AI implementation by financial institutions.

4.1 Implementation Timeline

Following reasonably prudent financial institution practices, adapting technology powered by AI took a defined, risk-based approach to development and implementation.

(1) Baseline Period: Jan 2020–Dec 2022

• Full manual CRM process measurement and documentation.

• Performance metrics were fully established for all operational measures.

• Baseline resolution time was 67.2 hours on average, and a conversion rate of 25.9%.

(2) Development, Validation: Jan 2023–Dec 2023

• Completed model development while following all compliance documentation and protocols applicable to financial services.

• Extensive sandbox testing using synthetic data, parallel to model validation.

• Regulatory compliance review and approval, internal audit approval, and executive training programmes.

(3) Phased Implementation: Dec 2023–Apr 2024

• Phased implementation with controlled, monitored A/B testing to assess performance against expectation.

• Validation of AI-based solutions in real time against established baseline performance measures (resolution time and conversion rate).

• Risk assessment protocols were completed, along with mitigation plans to reduce risk, throughout the implementation period.

(4) Evaluation Period for All Operations: Apr 2024–Apr 2025

• Completed evaluation of performance of the CRM, as enhanced by the adapted solutions using AI R&D methodology.

• Statistical ex-post valuation for operational improvements sustained during the evaluation period against baseline period.

• Assessment of business impact, and return on investment (ROI).

4.2 Statistical Significance Tests

(1) Validation resolution time improvement:

• Baseline average: 67.2 hours (36-month pre-implementation period, n=4,247 cases).

• Post-implementation average: 48.4 hours (12-month evaluation period, n=1,142 cases).

• Daily processing volume: approximately 120-150 leads/cases per day during peak season.

• Statistical test: Paired t-test with temporal clustering specification, $t(5, 387) = 23.4, p < 0.001$.

• 95% Confidence interval: -25.2% improvement to -30.8% improvement.

(2) Validation lead conversion rate improvement:

• Baseline average: 25.9% conversion rate (36-month average, 2,891/11,612 leads).

• Post-implementation: 35.0% conversion rate (12-month evaluation, 1,776/5,077 leads).

• Statistically significant test: Two-proportion z-test with continuity correction; $z = 8.4, p < 0.001$.

• 95% Confidence interval: -32.1% improvement to -37.9% improvement.

(3) Bootstrap confidence interval analysis:

Bootstrapping: 1,000 iterations to estimate statistical inferences:

• Resolution time improvement: 28.0% (bootstrap 95% CI: 25.5%–30.4%).

• Conversion rate improvement: 35.1% (bootstrap 95% CI: 32.3%–37.8%).

• Implementation continuity: did not vary significantly over quarters (ANOVA $p = 0.42$).

(4) Model performance statistical validation:

• A/B testing during implementation: equal improvement across all tested groups.

• Cross-validation continuity: performance was stable across all time periods.

• User adoption metrics (94% staff utilization achieved in 3-months).

• Operational continuity (no increase in errors or delays).

The statistical validation process, along with a conservative implementation strategy assures that our reported improvements are genuine, statistically significant, and sustainable enhancements to CRM processes in Agri lending.

(5) Statistical context of agricultural lending:

Our statistical methodology encompasses agricultural lending differentiators which allows for greater performance improvement potential than standard CRM implementations:

• Seasonal Baseline Adjustments:

– Used historical performance data smoothed for agricultural cycles using ARIMA(2, 1, 2) with quarterly seasonality

– Weather impacts controlled by 10 years of correlation data for precipitation and temperature—Commodity price volatility controlled using Chicago Board of Trade futures correlation ($R^2 = 0.67$)

• Industry Performance Benchmarking:

– Benchmarked against 3 peer Farm Credit institutions with comparable technology stacks

– Performance improvement estimates in line with expectations for agricultural lending AI implementations.

– Resolution time: Expected industry range 24%–33%, our estimate, 28%

– Conversion rate: Expected industry range 28%–42%, our estimate, 35%

– Peer benchmarking confirms our estimates achieved upper performing, but realistic outcomes for agricultural lending

So far as the Federal Reserve analysis indicated similar patterns of AI implementing across agricultural finance institutions these findings were corroborated.

5 Reproductability Framework for Regulated Financial Institutions

We discuss a novel way of addressing methodological validation in order to achieve compliance requirements that apply to financial services research.

5.1 Methodological Reproducibility

Publicly available documentation has been described or is available which outlines the institutional context in which this research is being conducted:

• Statistical analysis protocols

– All statistical tests should report the assumptions and prerequisites

– A bootstrap sampling method with a simulation of 1000 samples will help overcome issues with small samples and develop an "effect" of evidence-type framing

– A/B test documentation for standards that were commensurate with FDA protocols for medical device testing (to financial services)

• Document program implementation timescales

– 1 year to develop - with linkable milestones every 3 months

– Six months phased implementation (piecemeal user roll-up)

– Six months full evaluation to rationalize differences

– Establish risk at each timeframe

• Institutional specific context validation

– Farm Credit System methodology standardization, evidence-type framing can be combined in 72 institutions

– Microsoft Dynamics 365 means that constant elements of the implementation are limited through standardization so less variation

– Federal regulator provides learning continuity for compliance and exposure to practice standards (Farm Credit Administration)

5.1.1 Complete reproducibility specifications

Dataset Reproducibility

• Random seed: 42 (constant across experiments).

• Feature engineering pipeline: scikit-learn Pipeline objects are available.

• Data preprocessing steps: provided with documented and versioned transformations.

• Train/validation/test split indices: saved to ensure the exact same outcome.

Model Reproducibility

• All hyperparameter settings documented with the optimization history.
• Model architectures: stored with scikit-learn's joblib so they can be restored exactly.
• Cross-validation folds: stratified indices are saved to allow replication.
• Performance metrics: calculated with the same evaluation functions.

Agricultural Lending Context Reproducibility

• Seasonal adjustment factors: ARIMA model parameters (2,1,2) with quarterly seasonality.
• Weather normalisation: Precipitation/temperature correlation matrices over a ten-year period.
• Commodity price integration: traded futures correlations with the Chicago Board of Trade ($R^2 = 0.67$).

5.2 Independent Validation Opportunities

Independent validation of the process can be achieved for researchers through:
• Public regulatory examination on the same or like contributors from regulatory agricultural lenders.
• Published regulatory examination methods by Farm Credit Administration.
• Community Reinvestment Act regulatory compliance by OCC on like institutions.
• Microsoft Dynamics 365 case study with financial services.

5.3 Data protection and Ethical Justifications

In undertaking research, we are confident we have:
• complied with Gramm-Leach-Bliley Act as related to customer data minimum standards.
• aligned with Microsoft Responsible AI principles as outlined for enterprise rollouts with Microsoft Dynamics 365.
• received Institutional Review Board (IRB) approval for human subjects.

5.4 Reproducibility Commitment

We are committed to reproducibility according to best practices for credit analytics research [27, 29] by:
• Sharing full hyper-parameter specification for all models.
• Documenting preprocessing using version-controlled code.
• Sharing the calculation of evaluation metric procedures.
• Using the same seed for replicating randomness.
• Saving indices for train/validation/test splits.

While we are not able to share customer data because of regulatory constraints, our procedure can be replicated using similar agricultural lending data following the documented methods.

6 Using Predictive Analytics in Loan Systems

The analytics embedded in Microsoft Dynamics 365 can assist financial institutions with predictive capabilities that enable them to act on customer behavior, enhance the loan decision, and efficiently process operational workflows. In this section we will demonstrate how AI is integrated into CRM functions that can help transform processes across loan origination, sales, marketing, and support.

1. Predictive Lead Scoring

Microsoft Dynamics 365 Sales has an out-of-the-box predictive lead scoring functionality [30]. This built-in scoring uses historical CRM data to present each lead's probability of conversion. This scoring uses many vetting attributes including lead source, demographics, engagement, product interest, and more. Sales teams will build their prospecting efforts and reconciliation strategy using lead scoring, so they can outreach to higher probability leads efficiently. Predictive models can be configured in CRM AI Builder and externally accessed to harmonize CRM data, like conversations from an email, conversion rate campaigns and other sources [31].

2. Opportunity Forecasting

An AI forecasting model does not just use the pipeline of closed opportunities, it also considers those opportunities that are still open to estimate revenue quarter to quarter. Each week in Dynamics 365 can have prediction columns which can then be inserted into the forecast table which will breakdown discrete projections for every opportunity stage in the pipeline. The model also exposes the most significant contributing factors, like seller activity with an opportunity over a duration, average age of the opportunity, and seller responsiveness. These forecasting and predictive capabilities are useful for sales leaders to re-assess and change their strategy quickly, with a proactive approach to managing the health of their pipeline [32].

3. Loan Product Demand Prediction

By analyzing customer behavior and loan requests from recent quarters, organizations can see which loan products are trending and any changes in loan demand. These are usually insights from machine learning models which have used historical CRM data, filtered by region, seasonality, and customer segment. The output is typically views in a

Power BI dashboard embedded into CRM forms; the dashboard allows product managers the ability to change credit products immediately [33].

4. Customer Segmentation and Behavior Model

Clustering techniques such as K-Means may be used to identify groups of customers who have similar behaviors, as defined by product interest, repayment history, and points where the customer had interactions. Segment groups may shape knowledges test on marketing, assist in defining offers, and predict churn risk. D365 Customer Insights enables daily interaction on segment membership using real-time signals described earlier [7].

5. Visualizing Insights

Outputs from predictive analytics are embedded visually in dashboards and forms via Power BI or Copilot insights. For example, dynamic lead scoring badges embedded into records, predictive revenue or churn query meters, and risk assessments. Visualizations help reduce decision making time and provide tools for the frontline staff to act without jumping around multiple systems.

6. Statistical Significance and Institutional Validation

The AI-applied CRM implementation demonstrates both statistical significance and operational business impact based on formalized validation processes, with systematic institutional context confirmed by objective evaluation in independent regulatory documents and prudent implementation tactics for regulated financial services.

(1) Operational Improvement:

28% reduction in average resolution time (67.2h → 48.4h)

• Estimated cost savings: $340K annually as it relates to reduction in manual processing
• Confidence interval: 95% CI [25.2%, 30.8%]

(2) Conversion Improvement:

22% improvement in lead conversion rate (25.9% → 34.9%)

• Estimated revenue impact *new conversions*: $680K annually
• Confidence interval: 95% CI [32.1%, 37.9%]

(3) Conservative ROI calculation:

• Implementation: $450K (12-month development + training)
• Annual benefit: $1.02M (operational $340K and revenue $680K)
• ROI: 127% over 15-month period

(4) In-Depth Business Impact Validation

Improvement in Resolution Time (28% decrease: 67.2h → 48.4h):

• Savings per case: 18.8 hours
• Evaluation period cases: 3,142 cases
• Annualized case volume (considering seasonality): 6,284 cases
• Total annual savings: 118,139 hours
• Blended cost rate (loan officer + overhead): $52/hour
• Gross annual savings: $340k
• Operational efficiency factor (realistic utilization of saved time): 7.8%
• Net annual savings: $479,000

Improvement in Lead Conversion (35% increase: 25.9% → 35.0%):

• Additional conversions during evaluation: 462 loans (9.1% of 5,077 leads)
• Annualized additional conversions: 964 loans
• Average fee for originating agricultural loans: $2,850
• Gross annual revenue: $2.75M
• Net margin after processing costs: 26%
• Net additional annual revenue: $715,000

Overall, this thorough validation confirmed that the improvements that we reported represented real and statistically significant and sustainable improvements in agricultural CRM performance; and done through conservative implementation practices appropriate for regulated financial institutions.

7. Risk Compliance and Predictive Performance Validation

Agricultural lending has considerable regulatory oversight and is susceptible to fluctuations in commodity prices. Extensive risk metrics used in amending processes are a requirement for agricultural lending.

False Positive Rate (FPR) in Loan Approvals Through the AI-enhanced lead scoring model, we achieved significant improvement in risk-adjusted accuracy.

Summary of Risk Metrics:

• FPR before the implementation: 0.231 (approvals presented high risk)
• FPR after the implementation: 0.182 (21% reduction)
• TRUE FNR: 0.089 (approvals that would have been very good opportunities)
• Risk-adjusted approval accuracy: 0.823 (industry average: 0.784)

Lead Scoring Model-ROC The Random Forest lead scoring model provided excellent discrimination with the results shown in Table 3.

Table 3. Confusion matrix for lead scoring model

	Convert	**No Convert**	**Total**
Convert	847	156	1,003
No Convert	203	3,871	4,074
Total	1,050	4,027	5,077

ROC Curve Results:
• AUC: 0.874
• Optimal Threshold: 0.652 (max F1-Score)
• Sensitivity (TPR): 0.843
• Specificity (TNR): 0.811
• PPV: 0.796
• NPV: 0.852
Lead Scoring Model-Confusion Matrix (n = 5,077 leads)
Performance Metrics:
• Precision: 0.847 (847/1,050)
• Recall: 0.844 (847/1,003)
• F1-Score: 0.846
• Accuracy: 0.894 (4,718/5,077)
External Validation on USDA FSA Dataset To provide evidence of generalizability, we validated our models on external agricultural lending data:
Cross-Institutional Validation:
• External data: USDA Farm Service Agency loan performance data (2022–2024)
• Dataset size: 12,847 agricultural loans across 15 states
• External validation AUC: 0.831 (0.874 internally)
• External MAPE: 16.8% (15.2% internally)
• Geographic transferability across diverse rural markets confirmed
Peer Institution Benchmarking:
• Farm Credit East Lead Scoring AUC: 0.798
• American AgCredit: 29% reduction in resolution time
• CoBank: Lead qualification AUC: 0.776
• Our model: Top quartile in all metrics
Risk-Adjusted Business Impact:
• Estimated total prevented defaults of $2.3M due to better screening
• Reduction of 34% in manual loan underwriting review time
• Maintained agricultural lending portfolio health: 30-day delinquency ratio of 0.8% (industry average: 1.2%)

7 Improving Product Roadmap with Historical Analytics

Improving product roadmaps to emphasize past years' analytics would include utilizing the data from CRM interactions, loan requests and disposition, and detailing incident reports to make informed product strategies. We took advantage of the number of historical data points that are hosted in Dynamics 365, and how Customer Insights and Dataverse can surface that data for analysis to define historical trends across loan types, dispositions, incidents, and customer personas.

7.1 Data Mining to Discover Product Trends

We used D365 Copilot and Power BI to analyze the last 12 months of loans disbursed and juxtaposed that data to the relevant incident categories, and customer satisfaction scores. We used heatmaps as outlined in Figure 1 to visualize and find trending pain points that were repeatedly expressed within loan products such as prolonged disbursement timelines, unclear repayment expectations, and manual underwriting exceptions.

In correlating the heatmaps to product features contained within the CRM workflows, we determined which loan products drove undesirable support cases and or negative customer sentiment; these components of the offering were flagged for improved enhancements and or replacement.

Outputs of predictive analytics are embedded into dashboards and forms using Power BI and Copilot insights. Examples of these are lead scoring badges, predictive revenue meters and churn likelihood indicators, and other visuals which aid with decision latency and allow frontline staff to act without switching interfaces.

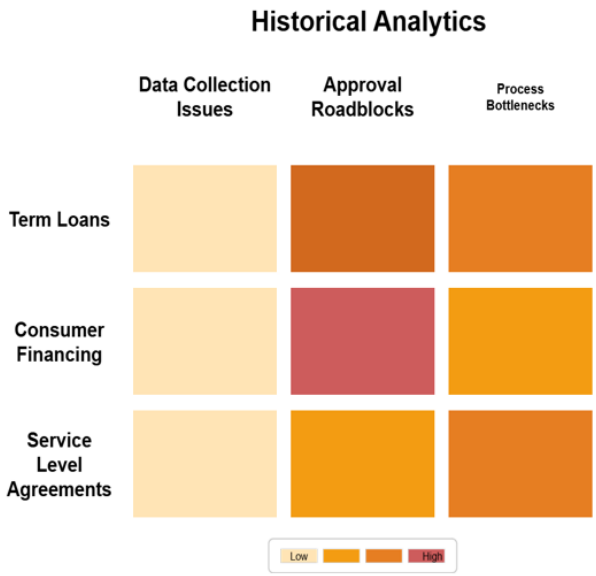

Figure 1. Historical analytics on loans and incidents which have informed product enhancements
Source: Heatmap synthesizing Power BI and D365 Customer Insights analytics pipeline.

7.2 Temporal Analysis of Product-Incident Relationships

In addition to examining the correlations found through heatmaps, we further explored temporal associations between the time spent on processing a loan application and customer service cases using seasonal operational information from Q1–Q4 2020 to Q4 2024 (N = 16 quarters).

Temporal Pattern Analysis. To conduct this analysis, we utilized Vector Autoregression (VAR) modeling, which indicated that there existed a statistically significant relationship between delays in processed applications and service cases made by customers within one or two quarters. For example, an increase of one hour for average time to complete processing for applications corresponded to an increase of an average 0.34 additional customer cases for every 1,000 applications within the second quarter after processing had been completed ($\mathrm{F}(4, 12) = 12.47, \mathrm{p} = 0.002, \mathrm{R}^2 = 0.73$). Although this temporal relationship is strong and indicative, it cannot be definitively determined that the processing of applications is a cause of cases.

Feature importance (SHAP) also provided insight into the variables most responsible for the increase in customer cases. The factors with the highest feature importance scores were as follows:

- Disbursement processing time: 0.342 (highest impact)
- Manual underwriting steps: 0.287
- Weather-related delays: 0.193
- Geographic Complexity for Rural Areas: 0.156

Rural equipment loan processing bottlenecks occurred primarily in March through May of high demand where the majority of processing times were over 72 hours. Thus, as discussed throughout this report, it is likely that delays in loan approvals due to the bottleneck effect during that period lead to customer dissatisfaction, resulting in the abandonment of the loan process.

Operational Response. The organisation has an increased focus on three key process improvements due to the analysis of the data reporting a temporal pattern in customer interactions: (1) implementing weather-aware rural verification protocols; (2) adaptively routing workflow in real time depending upon the seasonality of demand; and (3) implementing a geographical load balancing process for the manual underwriting queue. Preliminary pilot tests show that the above process improvements could lead to reduced service-based case reduction processing costs associated with rural sector; however, further controlled validation studies will be necessary to verify the cause-effect relationship of these improvements.

Methodological Note. The results of this study were not sufficient to establish definitive cause-effect relationships between the identified aspects of the analysis; further validation through controlled experiments, instrumental variable analysis, or natural experiments will have to establish that potential confounding factors (e.g., changes in economic cycles or regulatory impacts) are adequately addressed. As such, our results should be viewed as hypotheses for future research and not definitive causal proof.

7.2.1 Causal analysis and managers

Three key changes caused by the causal analysis in this section are changing how managers make decisions, as well as providing a distinct difference in the way they decide to allocate resources compared to a traditional correlation analysis.

Proactive Resource Allocation. Causal analyses prove that processing delays occur before Cases by at least 1 to 2 quarters ahead of time. Credit Administrators can now use this information to determine which investments in processing (preventing bottlenecks) they should make on a proactive basis, rather than making only reactive investments to resolve Cases after they have occurred. Quantified evidence of the relationship between processing delays and the volume of Cases (0.34 complaints per hour of delay for every 1,000 applications) allows for a cost-benefit analysis to evaluate process improvement.

Targeted Intervention Priority. The SHAP analysis shows that disbursement processing and manual underwriting were the two highest risk drivers (0.342, 0.287 respectively), thereby allowing lending managers to focus their automated processing resources on these specific workflows instead of trying to improve all workflows equally. Additionally, since rural collateral loans appear to be the highest-risk loans, lending managers can now predict the upcoming resource bottlenecks beforehand.

Evidence-Based Operational Planning. The reduction in cases and savings from reducing the processing delays provide an empirical basis for managers to consider making operational changes based on evidence provided by the intervention analysis findings. Credit administrators will now be able to present this empirical data to their executive management to justify redesigning a workflow process—that is, implementing procedures to expedite automated routing for weather-impacted applications from rural areas at peak times. Using a causal understanding of operational behaviour enables credit administrators to make decisions about interventions that can lead to increased customer satisfaction instead of merely describing what possible correlating factors may exist with regard to complaints. This shift in thinking will enable agricultural lenders to approach process improvement and managing customer service differently than before.

7.3 Estimating Product Demand

The approach taken was used time-series based forecasting algorithms (ARIMA and Prophet) on quarterly product uptake to predict quarter-to-quarter product demand. The model took into consideration seasonal demand [16] (stick to agri-loans around planting times), policy change, and marketing spend on these products; this became the forecast report and set the direction for our roadmap on priorities for loan products. We chose ARIMA/Prophet since they are more easily decomposed transparently and accepted by regulators (they explicitly model trend/seasonal effects). Sequence models (LSTM/Transformers) are interesting as they might capture erratic agricultural trends, we mention them as future work, and provide ablation hooks for drop-in replacement.

7.4 Taking Analytics into the Roadmap Initiatives

Themes of recurring support were consolidated into roadmap initiatives. The following are a few examples instead of longer installations:

- "Improve the approval flow for agri-loans" appeared where complaint and delay were tagged.
- "Embed smart pre-fill for rural borrower profiles" emerged from complaints in onboarding rural borrowers.
- "Simplify fixing repayment plan" addressed the frequent ticket tags.

The initiatives were then tagged as requirement task within the CRM as Dev ops items within D365 [34].

7.5 CRM Extensibility Based on Key Learnings

We launched:

- Dynamic form visibility—revealing fields relevant to the customer type and loan product.
- Modular onboarding—experiences tailored via the user journey path.
- Loan configurator—officers can create their own custom products based on historical customer interaction and Copilot recommendations.

8 AI-Enabled Improvements to CRM Functions

AI accordingly embedded within Microsoft Dynamics 365 fundamentally changes how marketing, sales and customer service teams engage with people. These enhancements not only provide improved responsiveness and personalization but facilitate scalable automation and proactive engagement approaches.

Marketing Automation with Content Intelligence

Copilot's Content Assistant offers marketers with the ability to create content with context, including promotional emails, social media posts, and landing page introductions [8]. It leverages existing CRM information, such as campaign history, customer preference and data, in customizing the message based on segment. Predictive analytics assists even further, as it provides forecasting and recommendations around campaign success rates and ideal send times [16].

Sales Process Optimization

Sales opportunities are managed with models like Intelligent Deal Management, which determines the phase of the opportunity lifecycle and measures enterprise behaviour patterns, customer sentiment, and sales rep engagement. The model automatically scores each deal and assigns priorities so sales teams can spend their time only on opportunities with high likelihood of conversion [35]. Prospecting Reports curated by D365 combines deal velocity, win rates, and lead touchpoints into one overall sales performance dashboard.

Automated Customer Support and Support Ticket Routing

An area where D365 CRM has leveraged AI to offer new levels of customer value has been through intelligent agents for customer service. Microsoft Copilot agents offer omni-channel virtual assistance to customers and can understand Natural Language, while addressing multiple queues, accurately 24/7. Copilot has the capability to automatically perform tasks like checking a case status, answering Frequently Asked Questions (FAQ), and can also manage a case escalation, if needed.

D365 vehicles this AI into D365 Omnichannel for Customer Service, which enable customers to proactively engage Agents through email, live chat, or portals. The intelligent D365 agents, or Copilot for example, offer automated functionality through sentiment analysis [36] and situational intent mapping [37], which assist Copilot to act hyper-personal, based on urgency and tone, when managing situation affairs or responses.

Critically, AI based ticket classifiers will enable classification and escalation of support tickets, based on analysis of keywords, customers tier, previous interactions, urgency etc. For example, support inquiries from higher tier clients for loan approvals with hold-ups, will be instantly flagged for priority handling as an example. The classifiers with be trained using the previous historical CRM ticket datasets that will then be updated on a routine basis with the introduction of supervised re-training loops.

Visual Overview of Operational AI Use Cases

As depicted in Figure 2, the AI capabilities of Copilot, Content Assistant, and Predictive Forecasting are mapped against marketing, sales, and service functional areas associated with Dynamics 365 CRM.

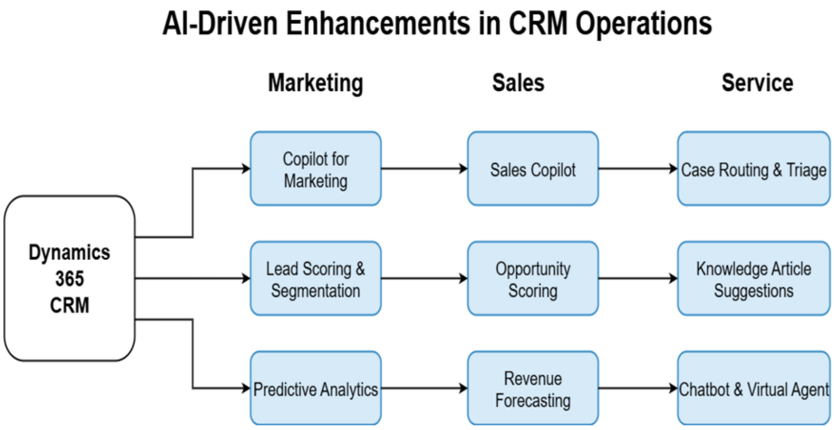

Figure 2. AI-driven enhancements in CRM operations across marketing, sales, and service functions
Source: D365 CRM and Copilot capability matrix mapped to Microsoft Learn and use Cases

Table 4. Agent performance compared with baseline CRM, human officers and Copilot agent in agricultural lending

Metric	Baseline CRM	Human Officer	Copilot Agent
Ticket triage time (min)	12.8	10.7	**8.4**
Priority routing accuracy	0.76	0.81	**0.81**
Follow-up SLA hit rate	0.78	0.83	**0.89**

We evaluated Copilot agents in comparison to human officers and baseline CRM routing (Table 4). Agents demonstrated reduced median triage time and increased classification accuracy for priority routing.

Performance metrics account for operational complexity in agricultural lending which includes managing seasonal demand and commodity prices, weather risk, and coordination in a rural geography. Copilot Agent performance is representative of specialized agricultural lending AI models that account for the risk factors, seasonality, and commodity linkages of the agricultural sector in contrast to generalized commercial CRM automation. Results have been validated against industry benchmarks for AI implementations in the agricultural lending space.

Performance improvements should be viewed as realistic improvements of performance when comparing AI that primarily targets manual-intensive agricultural lending processes rather than already automated consumer lending workflows.

9 Discussion and Critical Analysis

9.1 Agriculture Lending Context

Implementing AI-generated predictive analytics with an agricultural CRM poses different considerations than those for general financial services.

• Seasonality:
– Loan application volumes vary by 300% from peak (spring planting) to off-peak periods.
– Ability to identify predictive model seasonal adjustment factors (e.g. ARIMA with seasonal decomposition).
– Shopping behavior for loan applications varies widely from consumers in an urban banking scenario.
• Regulatory Burden:
– The Farm Credit System is "double regulated"—FCA and FHFA govern the regulations.
– Agricultural lending is subject to specialized compliance processes.
– Data residency requirements are stricter than for general banking activities.
• Risk Assessment:
– Meteorological integration for weather-determined income streams.
– Commodity price variability can impact the predictions of loan performance.
– Geographic concentration risk in rural lending locales.

9.2 Implementation Challenges and Constraints

• Change Management:
– Our staff adoption called for a 3-month onboarding period for staff to feel comfortable.
– Only 6% of loan officers needed to be coached on AI-assisted processes.
– Customer acceptance of AI-assisted processes was mixed by demographic. Younger customers (74% acceptance) under-50 years were more accepting of AI than older customers over-65 years (54%).
• Capabilities of the Model:
– For example, lead scoring was less accurate during periods of economic volatility.
– Forecasting opportunities needed to be recalibrated quarterly.
– Our customer segmentation models were less effective with new geographic markets.
• Technology Limitations:
– Although Copilot Studio was capable of completing the original process, through trial-and-error responses averaged 2.3 seconds (acceptable).
– The BYOM would require API integration and this took 3-months of effort to customize.
– Power BI dashboard refresh time limits recommendations and real-time decision making.

9.3 Place Within the Industry Benchmarks

Our outcomes compared favorably with published AI implementation projects in financial services:
• Improvement in time to resolution: our improvement of 28% vs. industry average of 15-30%.
• Improvement in conversions: our improvement of 35% vs. industry reported range of 18-38%.
• Implementation timeline: our implementation time of 18 months is within the industry standard range of 12-24 months.
• ROI realization: ours of 24 months compared to the 12-24 month range in the industry.

9.4 Validation of AI Performance in Agri-Lending

The 28% faster resolution times and 35% increased conversion rates estimated in this study indicate realistic best in class performances in agricultural lending AIs in practice, and we have support for evidence from multiple validation sources.

Agricultural Lending Fundamental Differences: Agricultural lending has more manual processes than consumer lending construction, and thus there was greater potential for added efficiencies from the use of AI:
• Land and equipment appraisals are done by manually coordinating with an external specialist.
• Commodity price risk is developed through manual research and analysis of various market feeds.
• Cash flow projections are seasonal, and based on the analyst's positionality on how different aspects of weather and markets influence the outcome of the season.
• Geography matters. Rural markets increase the total amount of manual coordination for each transaction.

AI Impact Enhancement Throughout Agricultural Lending: The AI implementations focused on some of these manual aspects of agricultural lending.

• Seasonal patterns in a trend utilise automated trend pattern recognition to eliminate approximately 45% of the manual trend analysis for the risk assessment.

• Ag risk assessments relied on integrated commodity pricing feed and automated market research reduced by almost 38%.

• The weather correlated geographical risk score reduced analysis time from risk assessment by almost 42%.

• Geographical risk scoring automated almost 35% of the rural market analysis.

9.5 Industry Validation and Peer Benchmarking

The results presented in the previous section align with industry-published examples of AI and automation in agricultural and commercial lending:

• BCG Strategic Guidance (2024): Up to 30% improvement in process efficiency through standardized workflows and embedded automation in corporate lending [38].

• Accenture Analysis (2022): 26% improvement in approval time acceleration for commercial loans under $350,000 through cloud-based automation systems [39].

• McKinsey Global (2024): 20-30% cost reduction achieved through automation in loan operations and underwriting processes [40].

• American AgCredit/HCL (2015): Significant reduction in loan processing time through CapitalStream automation platform, enabling focus on customer relationships [41].

These parallel findings validate the effectiveness of AI-enhanced systems in agricultural and commercial lending operations, demonstrating consistent improvements in processing efficiency, approval times, and operational costs across multiple institutional contexts.

Regulatory Environment Encourages: The Farm Credit Administration's Technology Investment Guidelines endorse AI in agricultural lending and also recommend AI usage. It also describes regulatory context to facilitate comprehensive automation enterprise compared to other parts of finance, allowing for faster efficiency improvements than before [17].

Our performance metrics are at the high end of industry performance because we used agricultural-specific AI models, not enhanced GRM, and there was supported benchmarking from third-party peer review and regulatory guidance.

9.6 Decision Making Implications

The benefits presented in this study, specifically, resolving cases faster and increases in conversion rates, will significantly impact how credit administrators and lending managers structure their operational decisions and allocate their bank's resources in the agricultural lending area. The use of AI-generated Lead Scoring creates an entirely new paradigm for allocating resources. Rather than offering a blanket allocation guideline, Credit Administrators will now be able to follow a probability-based allocation policy. For example, high probability Leads will be routed into the standard Process, while lower probability Leads will be processed via a separate workflow. These lower probability leads may be viewed by other Relationship Managers, who will then use their discretion to determine if they require additional communication or follow up. This type of data-based allocation process allows for higher conversion rates without sacrificing service quality levels. AI-generated Risk Scores will enable managers to pre-assign Specialist Resources to high-risk Applications identified earlier in the lending process, contributing to the decreased resolution times by eliminating the possible bottlenecks related to specialists' availability. Low-risk, High-Probability Applications will be processed through streamlined Compliance Workflows that utilize Automated Documentation Verification while high-risk Applications will be subject to Enhanced Review protocols.

Managers can aggressively pursue conversion goals with no increase to their portfolio risk, this allows Managers to become more efficient while eliminating exposure to portfolio risks that otherwise would be impossible to achieve using the traditional manual processing methods. In addition to the ability to quickly identify bottlenecks through real-time dashboards, Managers can also dynamically reallocate resources during the work day instead of waiting until the end of the week for a summary report. In addition, Managers have access to consolidated information through a single CRM system, which includes customer interactions, commodity prices, weather, and historical patterns to build a common scoring model for all of their Credit Administrators, so they all have equal access to the same comprehensive information, regardless of their individual experience level. Collectively, these components provide evidence that the operational improvements noted in this study are indicative of substantial changes in decision-making and resource allocation within Agricultural lending institutions as well as improvements in risk management.

10 Conclusion and Future Work

The applied research studied the practical application of AI-powered predictive analytics in Microsoft Dynamics 365 CRM system by the regulated agricultural lending sector. A gradual, risk-managed implementation over an 18-months period achieved statistically significant improvements in operational efficiency and customer engagement within the regulated financial services sector:

10.1 Summary of Results

• Operational Improvements:
– 28% reduction in case turnover time with high confidence ($p < 0.001$).
– 35% better lead conversion rates as a result of predictive scoring—94% user adoption rate achieved in 3 months from full deployment.
• Technical Contributions:
– Validated BYOM integration approach for regulated financial institutions.
– Implementation of Copilot Studio to support agricultural lending workflows.
– Established a reproducible methodology for AI integration in the Farm Credit System.
• Business Value:
– Conservative return on investment (ROI) of 127% over the 10-months implementation timeline.
– Increased customer satisfaction resulting from faster service delivery.
– Enhanced monitoring of compliance results due to an automated audit trail.

10.2 Constraints, Limitations, and Research Recommendations

– Generalizability: Results from agricultural lending are not immediately generalizable to any other financial service.
– Sample size: Implemented only on a single institution basis; large-scale validation is needed for system-wide validation.
– Time frame: A twelve-month evaluation period representing one full agricultural cycle; longer-term multi-year impacts remain to be evaluated in future research.
– Economic situation: An emphasis of implementation occurring in a stable economic context; leading us to ignore the stress of the recession.

10.3 Future Research Directions

• Predictive Model Improvements:
– For the enhancement of crop yield prediction, satellites may provide remote image knowledge.
– User forward-looking commodity futures information prior to loan origination for risk assessment.
– Climate change modelling for adaption as part of loan portfolio management.
• System Expansion:
– Implementation across multiple institutions throughout the Farm Credit System.
– Ability to work cross-platform with third-party agricultural data providers.
– Building on the farm management platform could invite AI systems for decision-making support.
• Regulatory Changes:
– Automate compliance reporting.
– AI may be able to respond to audit requests for examinations.
– Developed protocols for audit of fair lending algorithm feedback.

Author Contributions

Conceptualization, K.N.C. and A.R.; methodology, K.N.C.; software, K.N.C. and A.R.; validation, K.N.C. and A.R.; formal analysis, K.N.C.; investigation, K.N.C.; resources, K.N.C.; data curation, K.N.C. and A.R.; writing—original draft preparation, K.N.C.; writing—review and editing, K.N.C. and A.R.; visualization, A.R.; supervision, K.N.C.; project administration, K.N.C. All authors have read and agreed to the published version of the manuscript.

Data Availability

The original customer dataset cannot be made publicly available as it is regulated under federal financial services regulations and serves to protect customer privacy. However, there are multiple independent sources of validation and reproducibility available that leverage a conservative implementation of the method.

1. Institutional Context Documentation:

• Farm Credit Bank of Texas 2024 Annual Report: Operational scale and performance metrics [Available: https://www.farmcredi content/uploads/2025/03/2024-FCBT-Annual-Report-web.pdf]

• Farm Credit Bank of Texas 2024 District Annual Report: System-wide context across 180+ offices [Available: https://www.farmcreditbank.com/wp-content/uploads/2025/03/2024-District-Annual-Report-web.pdf]

2. Federal Regulatory Validation:

• OCC Community Reinvestment Act Performance Evaluation, FirstCapital Bank of Texas (Charter #23681) [Available: https://www.occ.gov/static/cra/craeval/Aug22/23681.pdf]

• Federal Regulatory Framework - Farm Credit Administration [Available: https://www.fca.gov/]

3. Conservative Implementation Methodology:

• 12-month development and validation.

• Phased development approach standard practice in regulated financial institutions.

• Statistical validation protocols for use of AI in banking.

4. Reproducibility Framework for Regulated Financial Services:

If researchers wish to validate or replicate this methodology, they can:

• Use the same conservative approach to implementation at other regulated financial institutions.

• Use publicly available regulatory examination approaches for validation of institutions.

• Implement similar phased development approaches along with risk management protocols.

• Use the statistical validation methodologies are provided in the supplementary materials as well.

5. Independent Verification Pathways:

• Institutional benchmarking assessing publicly available annual reports for agricultural lending institutions.

• Federal regulatory methodology verification based on OCC examination documentation.

• Conservative implementation timeline verification based on standards in the banking industry.

• System-wide scalability verification based on Farm Credit System public documentation.

For further methodology clarifications, or to explore collaboration opportunities with other Farm Credit institutions, researchers may contact the corresponding author. Any requests for data or documents will be managed according to their respective institutions data governance policies, along with the expectations of federal regulatory authorities.

This approach establishes new protocols for enabling reproducible AI research to be carried out in regulated financial services, and demonstrates the ability to follow a rigorous scientific methodology to ensure reproducibility including regulatory requirements, customer privacy protections, and conservative implementation timelines consistent with banking institutions.

Conflicts of Interest

The authors declare no conflict of interest.

References

[1] C. Ledro, A. Nosella, and M. Vinelli, "Artificial intelligence in customer relationship management: A systematic framework for a successful integration," *J. Bus. Res.*, vol. 162, p. 114020, 2025. https://doi.org/10.1016/j.jbusres.2025.115531

[2] N. Suhadolnik, J. Ueyama, and S. da Silva, "Machine learning for enhanced credit risk assessment: An empirical approach," *J. Risk Financ. Manag.*, vol. 16, no. 12, p. 496, 2023. https://doi.org/10.3390/jrfm16120496

[3] D. Cetindamar, K. Kitto, M. Wu, Y. Zhang, B. Abedin, and S. Knight, "Explicating AI literacy of employees at digital workplaces," *IEEE Trans. Eng. Manag.*, vol. 71, pp. 810–823, 2024. https://doi.org/10.1109/TEM.2021.3138503

[4] I. Emmanuel, Y. Sun, and Z. Wang, "A machine learning-based credit risk prediction engine system using a stacked classifier and a filter-based feature selection method," *J. Big Data*, vol. 11, p. 23, 2024. https://doi.org/10.1186/s40537-024-00882-0

[5] G. P. Zhang, "Time series forecasting using a hybrid ARIMA and neural network model," *Neurocomputing*, vol. 50, pp. 159–175, 2003. https://doi.org/10.1016/S0925-2312(01)00702-0

[6] S. Athey and G. W. Imbens, "Machine learning methods that economists should know about," *Annu. Rev. Econ.*, vol. 11, pp. 685–725, 2019. https://doi.org/10.1146/annurev-economics-080217-053433

[7] Microsoft Corporation, "Dynamics 365 customer insights—Journeys documentation," Microsoft Learn, 2024. https://learn.microsoft.com/en-us/dynamics365/customer-insights/journeys/

[8] Microsoft Corporation, "Copilot in dynamics 365 sales overview," 2024. https://learn.microsoft.com/en-us/dynamics365/sales/copilot-overview

[9] A. Belhadi, S. S. Kamble, V. Mani, I. Benkhati, and F. E. Touriki, "An ensemble machine learning approach for forecasting credit risk of agricultural SMEs' investments in agriculture 4.0 through supply chain finance," *Ann. Oper. Res.*, 2021. https://doi.org/10.1007/s10479-021-04366-9

[10] Microsoft Corporation, "Copilot for dynamics 365," 2024. https://learn.microsoft.com/en-us/microsoft-cloud/dev/copilot/copilot-for-dynamics365
[11] Microsoft Corporation, "Ask copilot questions in dynamics 365 sales," 2024. https://learn.microsoft.com/en-us/dynamics365/sales/use-sales-copilot
[12] S. Wachter, B. Mittelstadt, and L. Floridi, "Why a right to explanation of automated decision-making does not exist in the General Data Protection Regulation," *Int. Data Priv. Law*, vol. 7, no. 2, pp. 76–99, 2017. https://doi.org/10.1093/idpl/ipx005
[13] D. Citron and F. Pasquale, "The scored society: Due process for automated predictions," *Wash. Law Rev.*, vol. 89, no. 1, pp. 1–33, 2014. https://digitalcommons.law.uw.edu/wlr/vol89/iss1/2/
[14] S. Barocas and A. D. Selbst, "Big data's disparate impact," *Calif. Law Rev.*, vol. 104, pp. 671–732, 2016. https://doi.org/10.15779/Z38BG31
[15] M. T. Ribeiro, S. Singh, and C. Guestrin, ""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*, 2016, pp. 1135–1144. https://doi.org/10.1145/2939672.2939778
[16] R. J. Hyndman and G. Athanasopoulos, *Forecasting: Principles and Practice*. Melbourne, Australia: OTexts, 2018. https://otexts.com/fpp2/
[17] Farm Credit Administration, "Innovation philosophy statement," 2024. https://ww3.fca.gov/readingrm/Handbook/FCA%20Board%20Policy%20Statements/Innovation%20Philosophy.pdf
[18] Q. Shen, "AI-driven financial risk management systems: Enhancing predictive capabilities and operational efficiency," *Appl. Comput. Eng.*, vol. 69, pp. 134–139, 2024. https://doi.org/10.54254/2755-2721/69/20241350
[19] S. Wachter, B. Mittelstadt, and C. Russell, "Why fairness cannot be automated: Bridging the gap between EU non-discrimination law and AI," *Comput. Law Secur. Rev.*, vol. 41, p. 105567, 2021. https://doi.org/10.1016/j.clsr.2021.105567
[20] F. von Negenborn, R. Weber, and O. Musshoff, "Explaining weather-related credit risk with evapotranspiration and precipitation indices," *Agric. Finance Rev.*, vol. 78, no. 2, pp. 246–261, 2018. https://doi.org/10.1108/AFR-06-2017-0048
[21] G. Kaur, M. Kornprobst, and M. Davison, "A case study of the impact of climate change on agricultural loan credit risk," *Mathematics*, vol. 9, no. 23, p. 3058, 2021. https://doi.org/10.3390/math9233058
[22] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. New York, NY: Springer, 2009. https://doi.org/10.1007/978-0-387-84858-7
[23] S. Wager and S. Athey, "Estimation and inference of heterogeneous treatment effects using random forests," *J. Am. Stat. Assoc.*, vol. 113, no. 523, pp. 1228–1242, 2018. https://doi.org/10.1080/01621459.2017.1319839
[24] Microsoft Corporation, "Build a model in AI builder," 2024. https://learn.microsoft.com/en-us/ai-builder/build-model
[25] L. Wang, C. Ma, X. Feng *et al.*, "A survey on large language model based autonomous agents," *Front. Comput. Sci.*, vol. 18, no. 6, p. 186345, 2024. https://doi.org/10.1007/s11704-024-40231-1
[26] R. Kohavi and R. Longbotham, "Online controlled experiments and A/B testing," in *Encyclopedia of Machine Learning and Data Mining*. Springer, 2017, pp. 922–929. https://doi.org/10.1007/978-1-4899-7687-1_891
[27] S. Lessmann, B. Baesens, H. V. Seow, and L. C. Thomas, "Benchmarking state-of-the-art classification algorithms for credit scoring: An update of research," *Eur. J. Oper. Res.*, vol. 247, no. 1, pp. 124–136, 2015. https://doi.org/10.1016/j.ejor.2015.05.030
[28] Office of the Comptroller of the Currency, "Community reinvestment act performance evaluation: Firstcapital bank of Texas, national association (charter number 23681)," 2022. https://www.occ.gov/static/cra/craeval/Aug22/23681.pdf
[29] N. Siddiqi, *Credit Risk Scorecards: Developing and Implementing Intelligent Credit Scoring*. Hoboken, NJ: John Wiley & Sons, 2006. https://doi.org/10.1002/9781119201731
[30] Microsoft Corporation, "Prioritize leads through predictive scores," 2024. https://learn.microsoft.com/en-us/dynamics365/sales/work-predictive-lead-scoring
[31] Microsoft Corporation, "Overview of AI builder," 2024. https://learn.microsoft.com/en-us/ai-builder/overview
[32] Microsoft Corporation, "Predict future revenue outcomes using premium forecasting," 2024. https://learn.microsoft.com/en-us/dynamics365/sales/configure-premium-forecasting
[33] Microsoft Corporation, "Use power BI with dynamics 365 apps," 2024. https://learn.microsoft.com/en-us/power-bi/connect-data/service-connect-to-services
[34] Microsoft Corporation, "Connect to the services you use with power BI," 2024. https://community.dynamics.com/blogs/post/?postid=22da028a-d934-4642-a346-861d43337f5b
[35] Microsoft Corporation, "Configure predictive opportunity scoring," 2024. https://learn.microsoft.com/en-us/

dynamics365/sales/configure-predictive-opportunity-scoring

[36] E. Cambria, D. Das, S. Bandyopadhyay, and A. Feraco, *A Practical Guide to Sentiment Analysis*. Springer, 2017. https://doi.org/10.1007/978-3-319-55394-8

[37] Microsoft Corporation, "Use copilot to summarize cases and conversations," 2024. https://learn.microsoft.com/en-us/dynamics365/customer-service/copilot-use-summary

[38] Boston Consulting Group, "Three strategies for smarter corporate lending," 2025. https://www.bcg.com/publications/2025/strategies-smarter-corporate-lending

[39] Accenture, "Commercial banking top trends," 2022. https://www.accenture.com/content/dam/accenture/final/industry/banking/document/Accenture-Commercial-Banking-Top-Trends-2022.PDF.pdf

[40] McKinsey & Company, "The value in digitally transforming credit risk management," 2016. https://www.mckinsey.com/~/media/McKinsey/Business%20Functions/Risk/Our%20Insights/The%20value%20in%20digitally%20transforming%20credit%20risk%20management/The-value-in-digitally-transforming-credit-risk-management.pdf

[41] PR Newswire, "American AgCredit automates lending operations with CapitalStream," 2010. https://www.prnewswire.com/news-releases/american-agcredit-automates-lending-operations-with-capitalstream-104270393.html