Loan Pirates

AI-supported Loan Default Prediction for Faster and Fairer Credit Decisions

Frank Roth, Linda Carmichael, Christina Ntenekou
MAS ETH in AI and Digital Technology

Problem Statement

Business context. SwissCredit Bank processes more than 50,000 loan applications per month across Switzerland and the European Union. The current manual review process requires approximately 3–5 business days, creating customer churn and operational cost pressure.

Core challenge. Build an AI-supported credit risk system that reduces decision time to under 30 minutes for most applications while keeping portfolio default rates below the current 20% threshold and meeting regulatory expectations.

Objectives.
- Accelerate low-risk loan approvals without weakening risk discipline.
- Support fair and explainable decisions for customers and loan officers.
- Provide an auditable model lifecycle suitable for regulated finance.
- Balance predictive quality, interpretability, operational feasibility, and compliance.

Stakeholder	Primary concern	Poster metric
CEO / Board	Competitive position, compliance	Approval time, audit readiness
Chief Risk Officer	Model quality, portfolio risk	AUC, recall, default-rate control
Retail Banking	Customer experience	Decision latency, churn reduction
Compliance	EU AI Act, GDPR, FINMA, revFADP	Explainability, human review, documentation
Loan Officers	Usability and trust	Local explanations, override workflow
Customers	Fair treatment and clarity	Bias metrics, adverse-action reasons

Data & Preprocessing

Dataset. Historical consumer lending data with financial indicators, repayment outcomes, temporal fields, categorical variables, and a binary target indicating default vs. non-default.

Preprocessing journey.
- Cleaned missing values and standardized mixed feature types.
- Encoded categorical variables and prepared numerical financial features.
- Addressed class imbalance through random oversampling of the minority class.
- Excluded direct protected race attributes from model training.
- Retained protected attributes separately for fairness and bias analysis.
- Retained FICO score as a realistic credit-risk signal, while monitoring its influence through interpretability analysis.

Main data risks. Class imbalance, opaque composite indicators, proxy discrimination, missing values, heavy-tailed financial features, and differences between U.S.-style lending data and the intended Swiss/EU deployment context.

Old vs. AI-enhanced Process

Old process

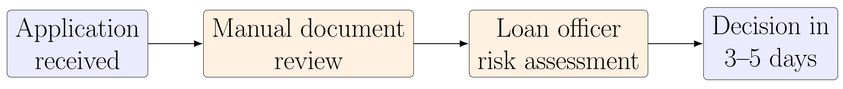

AI-enhanced process

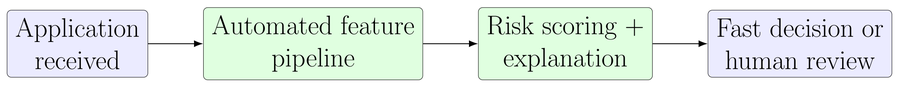

The proposed workflow shifts repetitive assessment tasks into a monitored ML pipeline, while preserving human accountability for ambiguous, risky, or legally sensitive decisions.

Model Architecture & Computational Requirements

Compared model families.
- Logistic Regression as interpretable baseline.
- Random Forest for robust non-linear tabular learning.
- XGBoost as the strongest tree-based candidate.
- MLP implementations in sklearn and PyTorch to test neural-network alternatives.

Computational observations.
- XGBoost provides a strong balance between model quality, training efficiency, and operational feasibility.
- sklearn MLP requires substantially longer iterative training and tuning.
- PyTorch improves engineering flexibility and can leverage hardware acceleration, but faster training did not automatically translate into better predictive quality.

Preferred model. XGBoost is the recommended candidate for structured tabular credit-risk data, subject to explainability, fairness, and governance controls.

Results & Evaluation

Evaluation dimensions.
- Statistical metrics: AUC, F1, recall, precision, confusion matrix.
- Business metrics: decision speed, expected portfolio risk, default-rate control.
- Regulatory metrics: explainability, auditability, fairness monitoring.

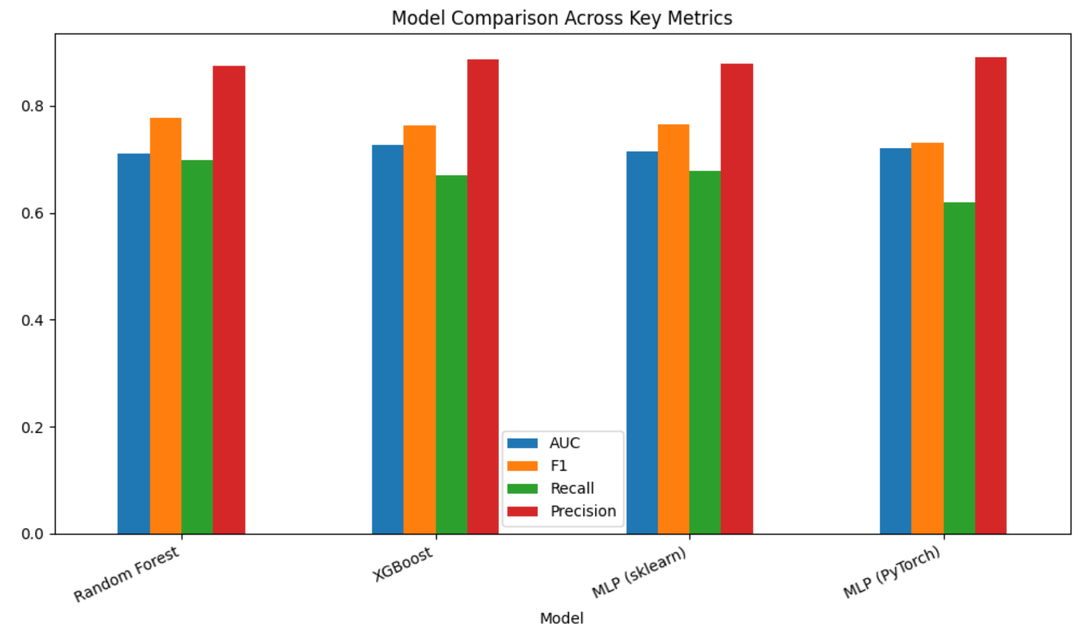

Interpretation. The model comparison highlights a common trade-off in credit-risk modeling: precision is generally high, but recall and false-negative control remain critical because missed defaults directly affect portfolio losses.

Interpretability, Fairness & Safety

Explainability approach.
- Global feature importance to understand portfolio-level model behavior.
- Local explanations for individual loan decisions and adverse-action reasoning.
- Monitoring of FICO score and other strong predictors to identify dependency on opaque or proxy variables.

Fairness and governance.
- Protected race attributes excluded from training but retained for bias testing.
- Bias assessment through demographic parity, equalized odds, and disparate impact.
- Human-in-the-loop review for uncertain or high-impact cases.
- Documentation and lifecycle controls aligned with EU AI Act, GDPR Article 22, FINMA, and Swiss revFADP expectations.

Key Takeaways

- Structured credit-risk data favors strong tree-based methods over generic neural networks.
- XGBoost offers the best practical balance of accuracy, speed, interpretability, and deployability.
- Accuracy alone is insufficient in regulated finance: fairness, explainability, auditability, and human review are design requirements.
- The proposed system should be introduced as governed decision support, not as unchecked fully automated credit approval.

Proposed Technical Implementation

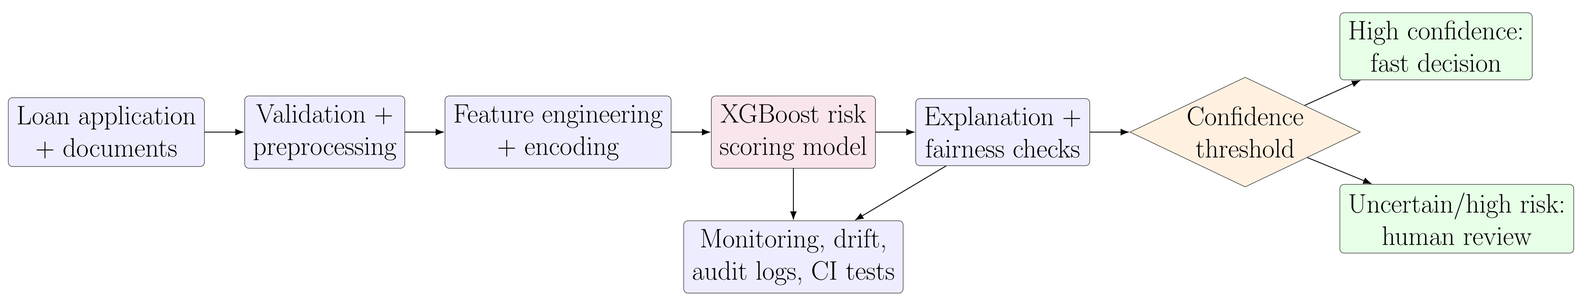

Implementation should expose a controlled prediction service with validated input schema, preprocessing parity between training and inference, logged model decisions, automated tests for preprocessing/model/fairness modules, and monitoring for data drift and performance degradation.