

# Credit Wise Loan System

## An Intelligent Loan Approval Prediction System

Live App: [creditoptima.streamlit.app](https://creditoptima.streamlit.app)

Source Code: [github.com/Butkii025/CreditOptima](https://github.com/Butkii025/CreditOptima)

---

## 1. Problem Statement

XYZ is a mid-sized financial company that offers personal and home loans to customers across urban and rural regions of India. Every day, hundreds of customers apply for loans through online and branch applications.

Until now, XYZ has relied on a **manual verification process**, where loan officers evaluate applications by checking income proofs, employment details, credit history, and other supporting documents. This process is time-consuming, inconsistent, and prone to human bias.

As a result, the bank faces two major challenges:

1. **Good customers sometimes get rejected**, leading to loss of business.
2. **High-risk customers sometimes get approved**, leading to financial losses.

To solve this problem, the bank wants to introduce an **intelligent loan approval system powered by Machine Learning** that can automatically analyze applicant details and predict whether a loan should be **Approved** or **Rejected**, before final human verification.

As the Machine Learning Engineer on this project, the goal is to design and develop this intelligent system using historical loan application data — one that learns hidden patterns from previous customer records and delivers accurate, fast, and unbiased loan approval decisions.

---

## 2. Objectives

- Analyze historical loan application data to identify the key factors that influence loan approval.
  - Build a machine learning classification model that predicts loan approval status (**Approved** / **Rejected**) with high accuracy.
  - Reduce human bias and inconsistency in the loan evaluation process.
  - Speed up the loan approval workflow, cutting down decision time from days to seconds.
  - Provide an interpretable, data-driven decision support tool for loan officers — not a black box that replaces human judgment entirely.
  - Deploy the model behind an interactive interface so non-technical staff can use it without writing code.
-

### 3. Dataset Description

Each row in the dataset represents a single loan applicant and contains personal, financial, and credit-related attributes collected at the time of application.

| Column             | Description                                                           | Type        |
|--------------------|-----------------------------------------------------------------------|-------------|
| Applicant_ID       | Unique identifier for each applicant                                  | Identifier  |
| Applicant_Income   | Primary applicant's monthly/annual income                             | Numeric     |
| Coapplicant_Income | Co-applicant's income, if any                                         | Numeric     |
| Employment_Status  | Employment type (e.g., Salaried, Self-employed, Contract, Unemployed) | Categorical |
| Age                | Applicant's age in years                                              | Numeric     |
| Marital_Status     | Married or Single                                                     | Categorical |
| Dependents         | Number of dependents financially reliant on the applicant             | Numeric     |
| Credit_Score       | Applicant's credit score, reflecting creditworthiness                 | Numeric     |
| Existing_Loans     | Number of loans the applicant currently holds                         | Numeric     |
| DTI_Ratio          | Debt-to-Income ratio — proportion of income already committed to debt | Numeric     |
| Savings            | Applicant's savings balance                                           | Numeric     |
| Collateral_Value   | Value of collateral offered against the loan, if any                  | Numeric     |
| Loan_Amount        | Amount of loan requested                                              | Numeric     |
| Loan_Term          | Requested loan repayment term, in months                              | Numeric     |
| Loan_Purpose       | Reason for the loan (e.g., Home, Personal, Car, Business, Education)  | Categorical |
| Property_Area      | Applicant's residential area type (Urban, Semiurban, Rural)           | Categorical |
| Education_Level    | Graduate or Not Graduate                                              | Categorical |
| Gender             | Male or Female                                                        | Categorical |

|                   |                                                             |             |
|-------------------|-------------------------------------------------------------|-------------|
| Employer_Category | Type of employer (e.g., Private, Government, MNC, Business) | Categorical |
|-------------------|-------------------------------------------------------------|-------------|

The dataset contains **1,000 applicant records**, with a small proportion of missing values (~5%) spread across **20 columns**, requiring imputation before modeling.

## 4. Project Workflow

The system follows a standard end-to-end machine learning pipeline:

Raw Data → Data Cleaning → Feature Engineering → Model Training  
→ Model Evaluation → Model Selection → Deployment (Streamlit App)

### 4.1 Data Cleaning & Preprocessing

- Missing numeric values (e.g., income, credit score) are imputed using the **mean** of each column.
- Missing categorical values (e.g., employment status, property area) are imputed using the **most frequent category**.
- Rows with a missing target label (**Loan\_Approved**) are excluded from training, since they can't be used to evaluate model performance.

### 4.2 Feature Engineering

- **Credit\_Score** and **DTI\_Ratio** are transformed into squared terms (**Credit\_Score\_sq**, **DTI\_Ratio\_sq**) to help the model capture non-linear relationships — for instance, the effect of credit score on approval likelihood isn't perfectly linear near the high and low ends of the scale.
- Categorical variables are **one-hot encoded** so they can be used by models that require numeric input.
- All numeric features are **standardized** (zero mean, unit variance) to ensure fair contribution across features with different scales (e.g., income in thousands vs. dependents in single digits).

### 4.3 Model Training

Three classification algorithms are trained and compared on an 80/20 train-test split:

| Model                      | Why it's used                                                                                                        |
|----------------------------|----------------------------------------------------------------------------------------------------------------------|
| <b>Logistic Regression</b> | A strong, interpretable baseline for binary classification; coefficients show how each feature drives approval odds. |

|                                  |                                                                                                   |
|----------------------------------|---------------------------------------------------------------------------------------------------|
| <b>K-Nearest Neighbors (KNN)</b> | Captures local patterns and similarity between applicants without assuming a linear relationship. |
| <b>Naive Bayes</b>               | Fast, probabilistic model that performs well even with limited data and correlated features.      |

## 4.4 Model Evaluation

Models are evaluated using:

- **Accuracy** — overall correctness of predictions
- **Precision** — of predicted approvals, how many were correctly approved
- **Recall** — of actual approvals, how many the model correctly identified
- **F1 Score** — harmonic balance between precision and recall
- **Confusion Matrix** — a breakdown of true/false positives and negatives

Precision and recall matter here more than raw accuracy: a false approval (high-risk applicant approved) carries financial risk, while a false rejection (good applicant rejected) carries business/reputational cost.

## 5. Results

| Model                      | Accuracy | Precision | Recall | F1 Score |
|----------------------------|----------|-----------|--------|----------|
| <b>Logistic Regression</b> | ~0.87    | ~0.82     | ~0.75  | ~0.78    |
| <b>Naive Bayes</b>         | ~0.86    | ~0.82     | ~0.73  | ~0.77    |
| <b>K-Nearest Neighbors</b> | ~0.76    | ~0.63     | ~0.60  | ~0.62    |

**Logistic Regression** was selected as the primary model, offering the best balance of accuracy and interpretability — an important factor for a financial institution that needs to explain its decisions to regulators and customers.

## 6. System Design

The final system is deployed as an interactive web application ("**CreditOptima**") — live at [creditoptima.streamlit.app](https://creditoptima.streamlit.app) — with three core views:

1. **Predict** — loan officers enter an applicant's details and receive an instant Approved/Rejected prediction with a confidence score.
2. **Model Comparison** — transparency into how each model performs, so stakeholders can see the reasoning behind model selection.
3. **Dataset Overview** — a snapshot of the training data and class balance, for auditing and trust-building.

## 7. Tools & Technologies

| Category         | Tools                                                |
|------------------|------------------------------------------------------|
| Language         | Python                                               |
| Data Handling    | Pandas, NumPy                                        |
| Machine Learning | Scikit-Learn (Logistic Regression, KNN, Naive Bayes) |
| Web Interface    | Streamlit                                            |
| Deployment       | Streamlit Community Cloud                            |
| Version Control  | Git & GitHub                                         |

## 8. Business Impact

- **Faster decisions** — reduces evaluation time from days to seconds.
- **Consistency** — removes officer-to-officer variability in judgment.
- **Reduced bias** — decisions are based on data patterns, not subjective impressions.
- **Scalability** — the same system can evaluate hundreds of applications simultaneously without added staffing.
- **Human-in-the-loop** — the model supports, rather than replaces, final human verification, keeping accountability with loan officers.

## 9. Limitations & Future Scope

- The model is trained on a limited, synthetic-style dataset (1,000 records) and would need retraining on real, larger-scale historical data before production use.
- Fairness auditing (e.g., checking for indirect bias against gender, region, or employment type) should be conducted before deployment in a real lending environment.
- Future improvements could include:
  - Adding ensemble models (Random Forest, Gradient Boosting) for potentially higher accuracy.
  - Incorporating explainability tools (e.g., SHAP values) so loan officers can see *why* a specific application was flagged.
  - Building a feedback loop where actual loan outcomes are used to periodically retrain and improve the model.

## 10. Conclusion

The Credit Wise Loan System demonstrates how machine learning can transform a manual, inconsistent loan evaluation process into a fast, data-driven, and more objective decision-support tool. By combining solid preprocessing, thoughtful feature engineering, and a transparent choice of model, the system gives XYZ a practical foundation for reducing both missed business opportunities and financial risk — while keeping loan officers in control of the final decision.

---