

AI Use Cases



Smartera 3S
Delivering the future now

Empowering Banks with Scalable, Ready-to-Deploy AI Solutions

We help financial institutions unlock the full potential of AI with pre-built, production-ready machine learning and generative AI use cases designed specifically for the banking sector. By leveraging our extensive library of industry-proven models and the Smartera 3S NextGen data platform, banks can accelerate their AI journey, reduce time-to-market, and cut implementation costs—without compromising on performance or regulatory alignment.

This approach delivers tangible value across four strategic pillars:

Accelerate AI Adoption

Leverage ready-made models for rapid market penetration.

Cost-Efficient

Reduce investment costs through industry-based ready-made models.

Comprehensive NextGen data platform

Optimize Smartera 3S next generation EDW platform with AI pre-built products.

Extensive Use-Case Library

A comprehensive repository of AI-driven solutions tailored for the financial services industry.



Machine Learning banking use cases

Machine Learning (ML) in banking is revolutionizing the financial industry by enabling data-driven decision-making. ML models leverage structured and unstructured data to generate predictive insights, automate processes, and personalize customer experiences. The key technologies involved include:

Accelerate AI Adoption

Leverage ready-made models for rapid market penetration.

Model-Driven Approaches


XGBoost, LightGBM, Random Forest, KMeans, and XGBRegressor.

Techniques Used

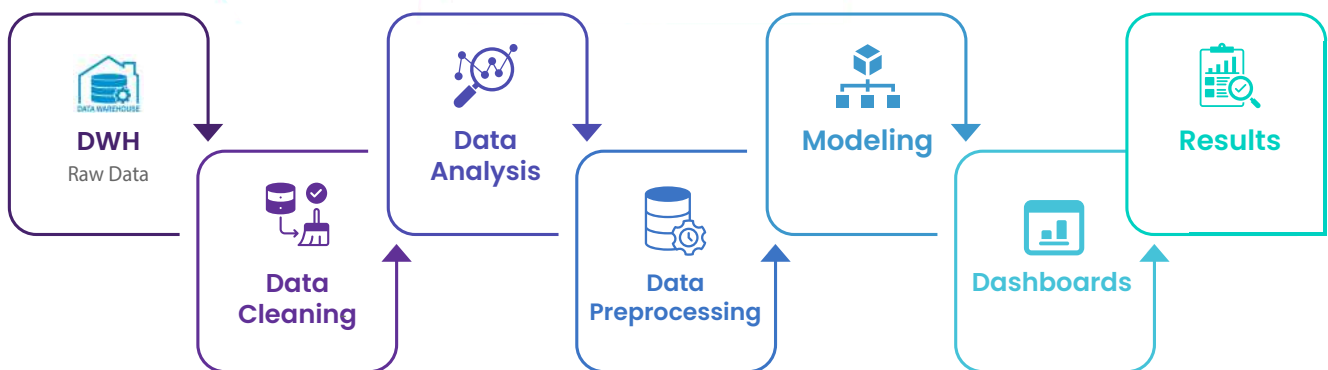
Classification, Regression, Clustering and Association Rules.

Core Components

Feature engineering, model evaluation, hyper-parameter tuning & dimensionality reduction.



The machine learning lifecycle is a structured, step-by-step process that transforms raw data into actionable insights through:



Raw Data Collection

Extract customer, transaction, and financial data from the Data Warehouse (DWH) for analysis.

Data Cleaning

Remove inconsistencies, handle missing values, and eliminate outliers to ensure data quality.

Data Analysis

Explore relationships and trends to identify the most impactful features and business patterns.

Data Preprocessing

Transform & prepare data using encoding, scaling, feature engineering, & class imbalance handling.

Modeling

Train machine learning models like XGBoost, or KMeans for predictions or segment customers.

Results & Dashboards

Train machine learning models like XGBoost, or KMeans for predictions or segment customers.

Some of ready made ML Banking use cases:



Credit Card Churn

Predicting credit card churn using Random Forest and XGBoost helps banks identify at-risk customers by analyzing transaction frequency, spending patterns, and engagement history. Feature engineering includes credit utilization, recency of use, and interaction frequency. Class imbalance handling (e.g., SMOTE) improves model accuracy, evaluated using AUC-ROC and F1-score. Insights drive targeted offers, loyalty rewards, and proactive engagement to reduce churn and enhance customer retention.



Term Deposit Churn

Predicting term deposit churn using XGBoost and Random Forest helps banks identify customers likely to withdraw or not renew their deposits. Models analyze tenure, deposit size, interest rate sensitivity, & withdrawal patterns. Class imbalance is addressed using SMOTE or cost-sensitive learning, & models are evaluated with AUC-ROC & F1-score. Insights drive personalized retention strategies, competitive rates, & tailored financial advisory services, enhancing deposit portfolio stability.



Customer Salary Estimation

Customer salary estimation using XGBRegressor predicts income based on demographics, spending patterns, and credit history, improving financial profiling and risk assessment. Key features include age, occupation, monthly deposits, and loan repayment history. RMSE ensures accuracy, enabling banks to personalize financial products, adjust credit limits, and refine risk-based pricing.



Customer Wealth Segmentation

Customer wealth segmentation using KMeans clustering helps banks classify clients based on net worth, income, & financial behavior for tailored financial services. Key features include account balances, investments, credit utilization, & spending patterns. Feature engineering involves scaling, categorical encoding, & PCA for dimensionality reduction, while the elbow method & silhouette score determine optimal clusters. This segmentation enables banks to offer exclusive banking services, personalized wealth management, & premium financial products, enhancing customer retention & revenue.



Spending Habits Segmentation

Customer spending habits segmentation using KMeans clustering groups clients based on transaction patterns, purchase frequency, and spending categories. Key features include spending volume, payment methods, and seasonal trends, with scaling and PCA improving cluster formation. The elbow method determines the optimal clusters. Insights enable personalized financial products, targeted promotions, and customized credit card rewards, enhancing customer engagement and loyalty.



Customer Loyalty Segmentation

Customer loyalty segmentation using KMeans clustering classifies customers based on engagement, tenure, and product usage to enhance retention and rewards. Key features include account tenure, transaction frequency, product holdings, and support interactions. Feature engineering involves scaling, encoding, and PCA for optimization, with elbow method and silhouette score determining clusters. This enables personalized rewards, loyalty programs, and exclusive benefits, strengthening customer relationships and reducing churn.



Corporate Profitability of Default

Predicting corporate Probability of Default (PD) using XGBoost and LightGBM helps banks assess credit risk by analyzing financial metrics, macroeconomic indicators, & repayment history. Feature engineering includes ratio calculations, while class imbalance is handled with SMOTE techniques. Models are evaluated using AUC-ROC & accuracy, enabling early warning systems, and regulatory compliance for proactive credit risk management.



Customer Demographics Segmentation

Customer demographics segmentation using KMeans clustering helps banks tailor services by grouping customers based on age, income, lifestyle, and financial behavior. Key features include account balances, transaction frequency, and credit utilization, with feature engineering involving scaling, encoding, and PCA. The elbow method or silhouette score determines the optimal clusters. Insights drive personalized marketing, product recommendations, and targeted financial services, enhancing customer satisfaction.



Balance Churn

Predicting balance churn using XGBoost and LightGBM helps banks identify customers likely to withdraw or reduce balances, enabling proactive retention. Key features include balance trends, transaction frequency, credit utilization, and engagement metrics, with feature engineering focusing on lag features, ratio calculations, and categorical encoding. SMOTE or cost-sensitive learning handles class imbalance, while AUC-ROC and precision-recall measure performance. Insights drive personalized incentives, financial planning, and engagement strategies to retain valuable customers.



Customer Income Level Prediction

Customer income level prediction using XGBoost helps banks estimate financial status based on demographics, transactions, and credit behavior, enabling personalized services. Key features include age, occupation, education, spending habits, and savings patterns, with feature engineering focusing on log transformations, ratio features, and categorical encoding. Hyperparameter tuning optimizes performance, evaluated using AUC-ROC, accuracy, and F1-score. These insights support customized financial products, credit limit adjustments, and pricing strategies.



Retail Cross/Up-Selling

Retail cross-selling in banking using association rules analyzes customer transactions to recommend complementary financial products. Patterns like credit card holders purchasing travel insurance are identified using support, confidence, and lift metrics. Feature engineering includes demographics, transaction history, and product usage. These insights drive personalized marketing, bundled offers, and automated recommendations, boosting revenue and customer engagement.



Corporate Cross/Up Selling

Corporate cross-selling using association rules identifies patterns in business banking behavior to recommend complementary financial products. Key features include company size, industry, transaction volume, and credit utilization. Frequent itemset mining uncovers relationships, such as firms with corporate loans adopting treasury services. Support, confidence, and lift metrics ensure relevant recommendations. These insights enable tailored financial bundles, automated recommendations, and targeted advisory services, strengthening corporate relationships and driving revenue growth.



Customer Channel Adoption Segmentation

Customer channel adoption segmentation using KMeans clustering classifies clients based on their preferred banking channels—online, mobile, or in-branch. Key features include transaction frequency, digital engagement, and in-branch visits, with scaling and PCA improving cluster formation. The elbow method determines the optimal clusters. Insights help banks enhance digital services, optimize branch experiences, and personalize customer interactions.



Complete Customer Loss

Complete customer loss prediction using XGBoost identifies clients at risk of fully leaving the bank by analyzing transaction frequency, product usage, balance trends, and digital engagement. Feature engineering and hyperparameter tuning enhance model accuracy, evaluated using AUC-ROC and F1-score. Insights enable personalized retention offers, loyalty incentives, and targeted engagement strategies to prevent customer attrition.



Point of Sale Churn

POS churn prediction using XGBoost identifies merchants at risk of discontinuing Point of Sale (POS) services by analyzing transaction volume, chargeback frequency, and terminal usage. Feature engineering and hyperparameter tuning enhance model accuracy, evaluated with AUC-ROC and F1-score. Insights help banks implement targeted retention offers, service improvements, and merchant engagement strategies to reduce churn.



Investment Plans Segmentation

Investment plans segmentation using KMeans clustering categorizes customers based on portfolio size, risk tolerance, and investment behavior. Feature engineering and clustering techniques improve segmentation accuracy. Insights help banks offer personalized investment products, targeted advisory services, and customized wealth management solutions to enhance customer engagement.



High Risk Customer Prediction

High-risk customer prediction using XGBoost analyzes transaction patterns, credit utilization, and financial stability to identify potential defaulters. Key features include income, account balances, debt-to-income ratio, and missed payments, with feature engineering and hyperparameter tuning optimizing model performance. Insights help banks implement proactive risk management, credit monitoring, and regulatory compliance strategies.



Corporate Wealth Segmentation

Corporate wealth segmentation using KMeans clustering categorizes businesses based on assets, revenue, and financial behavior. Key features include transaction volume, credit utilization, and investment activity, with scaling and PCA improving cluster formation. Insights enable banks to offer customized financial services, premium banking privileges, and specialized advisory solutions for high-value corporate clients.



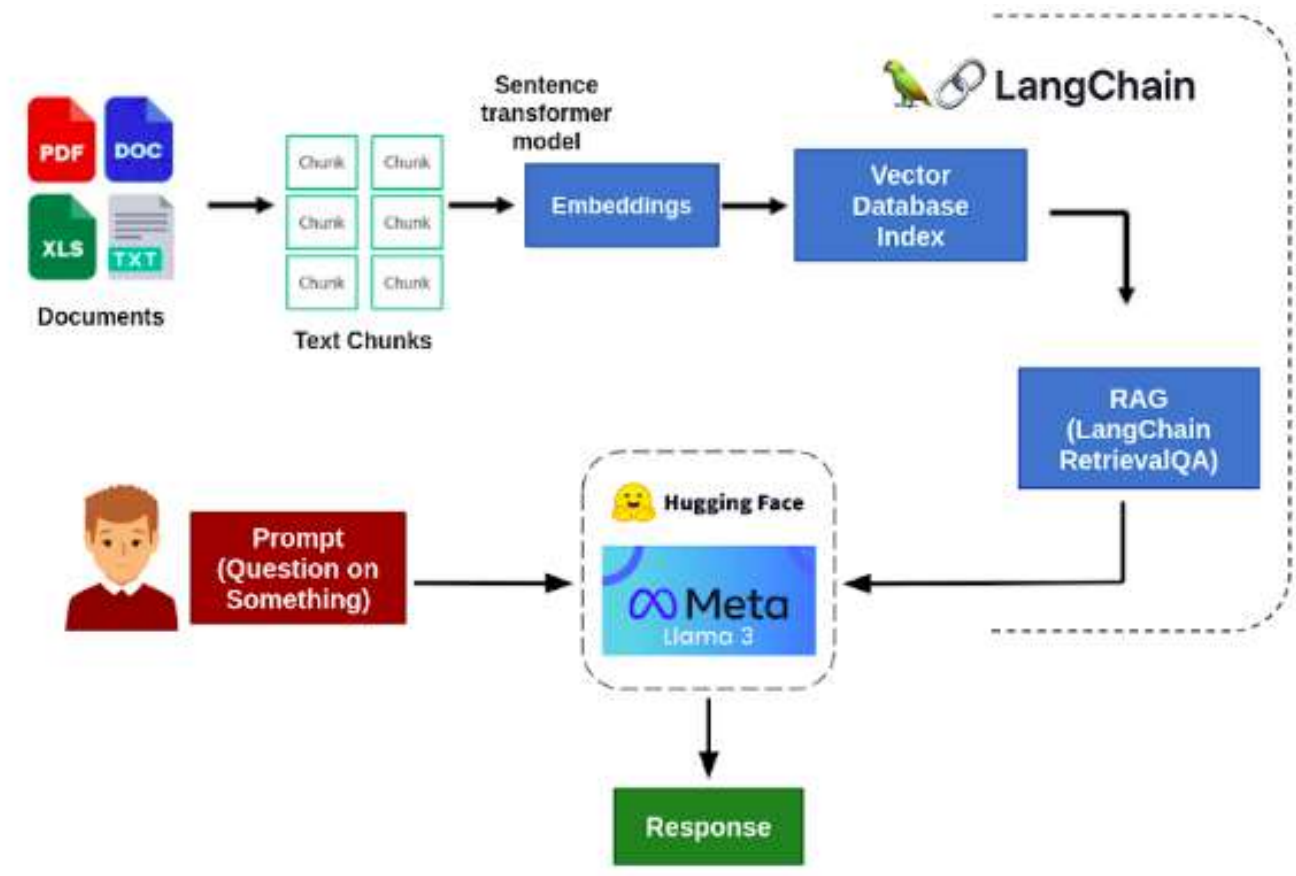
Corporate High Risk Prediction

Corporate high-risk prediction using XGBoost identifies businesses prone to financial instability or default by analyzing revenue, debt ratios, cash flow, and credit history. Feature engineering and hyperparameter tuning enhance model accuracy, evaluated with AUC-ROC and precision-recall scores. Insights help banks implement early risk interventions, stricter credit policies, and tailored financial solutions to mitigate losses.



Generative AI use cases

Modern machine learning advancements, particularly with Large Language Models (LLMs) which are the fruitful output of transformers improvements, have introduced a smart and time-efficient approach to document comprehension. By utilizing models like Llama 3 and Retrieval-Augmented Generation (RAG), we can vastly improve the accessibility and usability of dense, technical documentation. Llama 3, developed by Meta, excels at interpreting user queries without needing exact keywords, making it ideal for navigating dense technical manuals. RAG enhances this by retrieving the most relevant document sections and generating well-structured answers that combine multiple sources. It can respond to queries using data from maintenance logs, troubleshooting guides, and specifications—all in one coherent reply.



Sample of our Ready made Gen AI use cases:

HR Assistant:

Uses AI to generate structured job descriptions and automate candidate screening by evaluating resumes against job criteria. This streamlines recruitment, improves hiring consistency, and reduces HR workload by providing instant candidate assessments. Large language models (LLMs) extract key candidate qualifications for efficient resume matching.

Policy Expert:

An AI assistant that helps employees navigate internal policies, regulatory guidelines, and compliance frameworks. It processes policy documents and offers instant, context-aware responses, ensuring regulatory adherence and reducing reliance on manual policy interpretation. Retrieval-augmented generation (RAG) models improve response accuracy by grounding answers in internal policy documents.

Product Expert:

An AI-powered chatbot that provides real-time, accurate responses to customer inquiries about banking products such as loans, credit cards, and savings accounts. It utilizes natural language understanding (NLU) and integrates with a bank's knowledge base to ensure up-to-date information, improving customer engagement and reducing call center workload. Transformer-based models, such as GPT, enable it to generate contextual responses with high accuracy.



Agentic AI

Agentic AI represents the next evolution in artificial intelligence, where AI agents move beyond simple responses and begin to autonomously plan, decide, and act based on their goals and the context of the data. These agents operate independently across tasks, analyze data from multiple sources, and continuously adapt to achieve desired outcomes.



Financial Assistant

The Financial Assistant is an Agentic AI product that empowers internal bank users with smart, data-driven financial advisory. It connects seamlessly with the bank's Data Warehouse, processes unstructured documents, pulls data from live market feeds and public APIs, and automates financial analysis and communication — all through a dynamic, multi-agent system. At the core of the solution is an Agent Router that intelligently distributes tasks to specialized agents:

Financial Agent

Interprets user queries, calculates metrics and retrieves insights from the banking DWH.

Scraping Agent

Fetches external data from public APIs, and live market feeds to enrich responses with real-world.

Parsing Agent

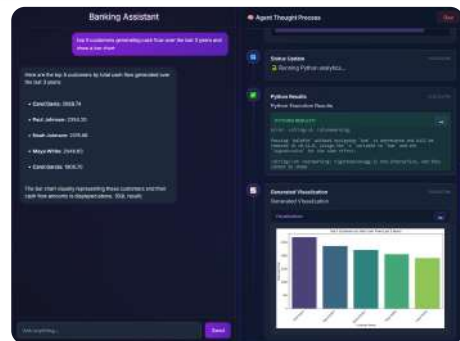
Processes uploaded documents like brochures or reports using OCR and NLP, then uses RAG to find and extract useful information.

Python Agent

Executes advanced calculations, forecasting models, or custom analytics scripts based on financial formulas or ML models using Python.

Communication Agent

Enables sending responses and reports via integrated channels like email and SMS, automating delivery to clients or relationship managers.



Key Advantages



End-to-end pipeline from data ingestion to response generation



Built-in support for multimodal data (structured, unstructured, real-time)



Flexible orchestration for both RAG and agentic tasks



Secure and scalable design ready for enterprise use

