MACHINE LEARNING CASE STUDY

Loan Acceptance Model

Targeting Cross-Sell with a Decision Tree

Roger Estrada · Machine Learning in Business · AI Strategy · 2026

A pruned decision tree identified which bank customers would accept a personal loan with 98.1% test accuracy and an AUC of 0.990. Income, education, and family size were the strongest signals, so outreach can be ranked and prioritized.

Business problem

A retail bank wants to convert deposit customers into loan customers. In this data only 9.6% of customers accepted a loan offer, so marketing to everyone wastes most of the outreach budget and creates unwanted offers. A propensity model scores each customer, so the cross-sell team contacts the likeliest prospects first.

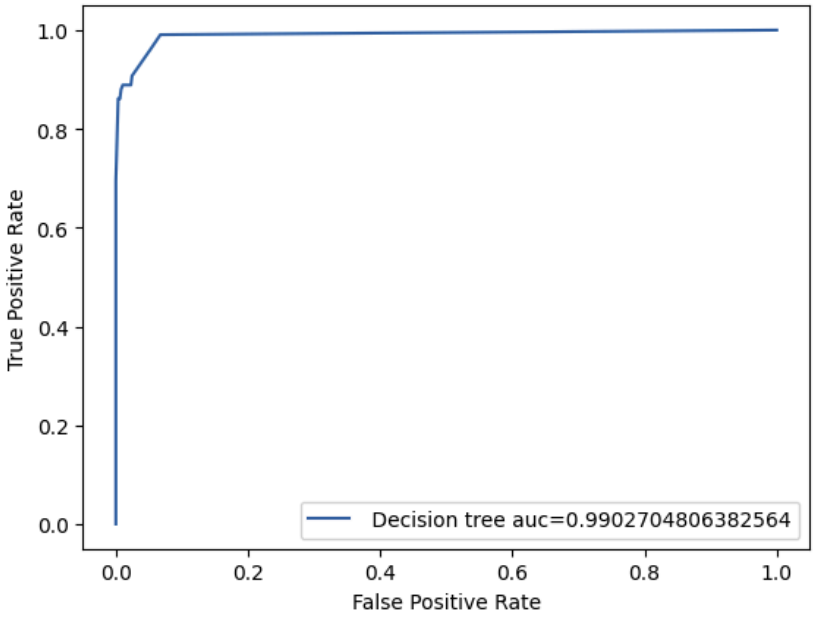
Approach and results

Data. 5,000 bank customers described by 11 attributes: age, experience, income, family size, average credit card spending, education level, mortgage balance, and indicators for securities account, CD account, online banking, and credit card. 480 customers (9.6%) accepted a personal loan. There were no missing values. The data was split 75/25, leaving 1,250 customers for testing.

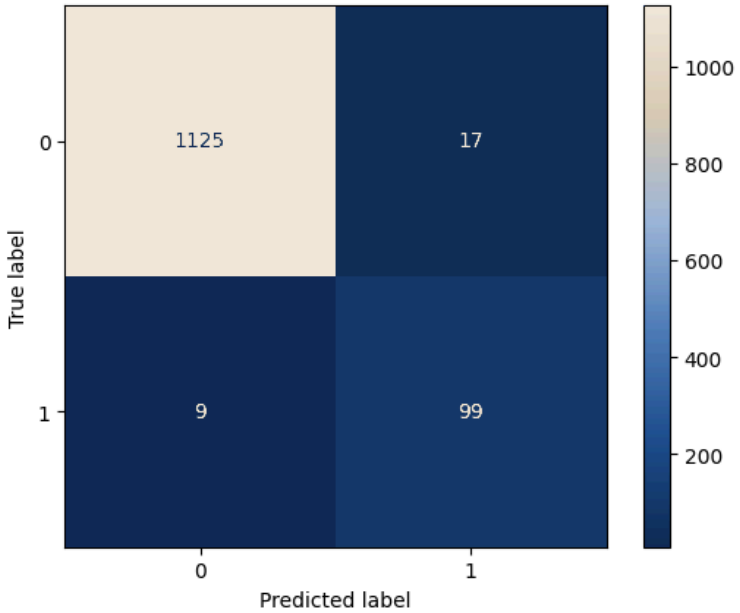
Baseline. In the test set, 1,142 of 1,250 customers declined, so offering to no one is already 91.4% accurate. As with any imbalanced target, the confusion matrix and AUC tell the real story.

Model	Train accuracy	Test accuracy	Test AUC
Unpruned decision tree	1.000	0.979	not reported
Pruned decision tree (max depth 4)	0.987	0.981	0.990

The pruned tree matched the unpruned tree on new customers while using far fewer rules, which makes it easier to explain and less likely to memorize noise.



The unpruned tree's confusion matrix shows the business trade-off. It found 99 of the 108 customers who accepted (91.7%). It made 17 offers to customers who declined, 1.5% of the 1,142 decliners. About 85% of the customers it flagged went on to accept.



What drives acceptance.

Feature	Importance
Income	0.447
Education level	0.284
Family size	0.129
Average credit card spending	0.082

Together these four account for about 94% of the tree's decisions. The other seven attributes each contribute less than 2%.

Enterprise application

The pattern applies to any next-best-product decision: deposit to loan, card to personal loan, checking to home equity, or member to wealth services. A ranked list lets the cross-sell team spend a fixed outreach budget on the strongest prospects. That raises conversion on the same spend, reduces unwanted offers, and supports both lending growth and member experience.

The tree also gives marketing plain rules for why a customer is prioritized, such as an income and education profile, which supports trust in the model and quick adjustment of campaigns.

Governance and limits

- **Teaching dataset.** This is a clean, well-known dataset. An AUC of 0.99 reflects that, and production results on real customer data are typically lower.
- **Fair lending.** Marketing models that influence credit offers need a review of features and outcomes. Age and family size are sensitive attributes and need controls, including checks for proxies.
- **Propensity is not eligibility.** A high acceptance score does not mean a customer qualifies. Approval stays under the credit policy.
- **Monitoring.** Campaign response shifts over time, so scores need performance tracking and periodic retraining.