International Journal of Management, IT & Engineering
Vol. 15 Issue 8, August 2025,
ISSN: 2249-0558 Impact Factor: 7.119
Journal Homepage: http://www.ijmra.us, Email: editorijmie@gmail.com
Double-Blind Peer Reviewed Refereed Open Access International Journal - Included in the International Serial Directories Indexed & Listed at: Ulrich's Periodicals Directory ©, U.S.A., Open J-Gate as well as in Cabell's Directories of Publishing Opportunities, U.S.A

LoanPulse AI: Proactive Delinquency Detection via Behavioral Intelligence

Sriranjan Chenoori

Abstract

Faced with heightened delinquency rates and heightened macroeconomic risk, early credit risk detection is now vital to supporting the stability of the financial system. LoanPulse AI provides an advanced solution to this issue by utilizing multi-signal behavioral modeling for early delinquency detection. In contrast to traditional credit scoring models developed based on large repayment history, LoanPulse AI integrates a broader spectrum of measures including real-time transaction activity, digital footprint of interaction, and macroeconomic or geo-financial incidents to create an integrated 360-degree behavioral risk image for each borrower.

It has a modular and scalable architecture built using Python and Snowflake that enables seamless ingestion, processing, and scoring of massive volumes of data. It has a real-time scoring engine that can detect user behavior anomalies, an early warning signal adaptive rule engine, and a dynamic visual alerting system that empowers credit risk teams to prioritize interventions. The system optimizes both batch and streaming mode operations, making it flexible enough to fit diverse banking environments.

This paper provides an in-depth description of technical architecture, data intake pipelines, signal fusion technique, and risk scoring algorithms powering the LoanPulse AI system. It also provides pilot deployment results that show its ability to identify high-risk borrower's weeks before existing models do, thereby enabling financial institutions to proactively act through restructuring or early contact. The findings underline the capacity of the system to significantly reduce non-performing loans (NPLs) and improve portfolio quality overall in a difficult economic environment.

Keywords:

Delinquency Prediction;
Behavioral Modeling;
Fintech AI;
Snowflake ML;
Early Risk Detection;
Financial Signals

Author correspondence:

Sriranjan Chenoori,
Engineering Manager, Charlotter, North Carolina, USA
Email: Sriranjanreddich@gmail.com

1. Introduction

The main text format consists of a flat left-right columns on A4 paper (quarto). In early 2025, consumer loan delinquency rates in the United States surged to their highest levels in over a decade, signaling a warning for banks, credit unions, and fintech lenders. This sharp rise has been fueled by a combination of persistent inflation, elevated interest rates, and widespread economic uncertainty [1]. As household budgets tighten, even prime borrowers are exhibiting early signs of financial distress often undetectable by traditional credit risk model.

Conventional credit scoring systems, such as FICO or internal bureau-based systems, are based to a large extent on static historical data payment history, credit outstanding, and credit utilization ratios. These metrics, while valuable, are lagging indicators by definition: they only show risk after it materializes. The result is that financial institutions are left with little time to take action before the customer enters serious delinquency or default.This reactive posture is no longer viable. In today's data-rich environment, customers generate a wealth of behavioral and transactional signals far ahead of failing to make a payment. Account activity trends, reversals in spending patterns, anomalies in digital engagement, and exposure to exogenous economic stressors are all leading indicators of financial instability. The integration of these multi-signal data points into a real-time predictive risk framework represents the next frontier in credit risk analytics.

LoanPulse AI was developed to address this pressing need. It introduces a forward-looking credit risk monitoring solution that integrates traditional financial data with digital footprints and third-party macroeconomic triggers. By developing a dynamic, 360-degree behavioral profile of each borrower, the platform allows lenders to identify and act on early warnings of risk—often weeks ahead of default. This earlier detection not only improves portfolio quality but also allows for more compassionate and flexible customer outreach programs such as restructuring or hardship relief.

This work describes the architecture and implementation of the LoanPulse AI platform, its supporting architecture, data pipelines, real-time scoring engine, and visualization layer. It also reports empirical findings from pilot deployments that show the ability of the system to predict delinquency risk far earlier than conventional models [1] [2]. In this manner, it aims to show how AI-powered behavioral modeling can be a cornerstone of resilient lending practices amidst an increasingly volatile economic climate.

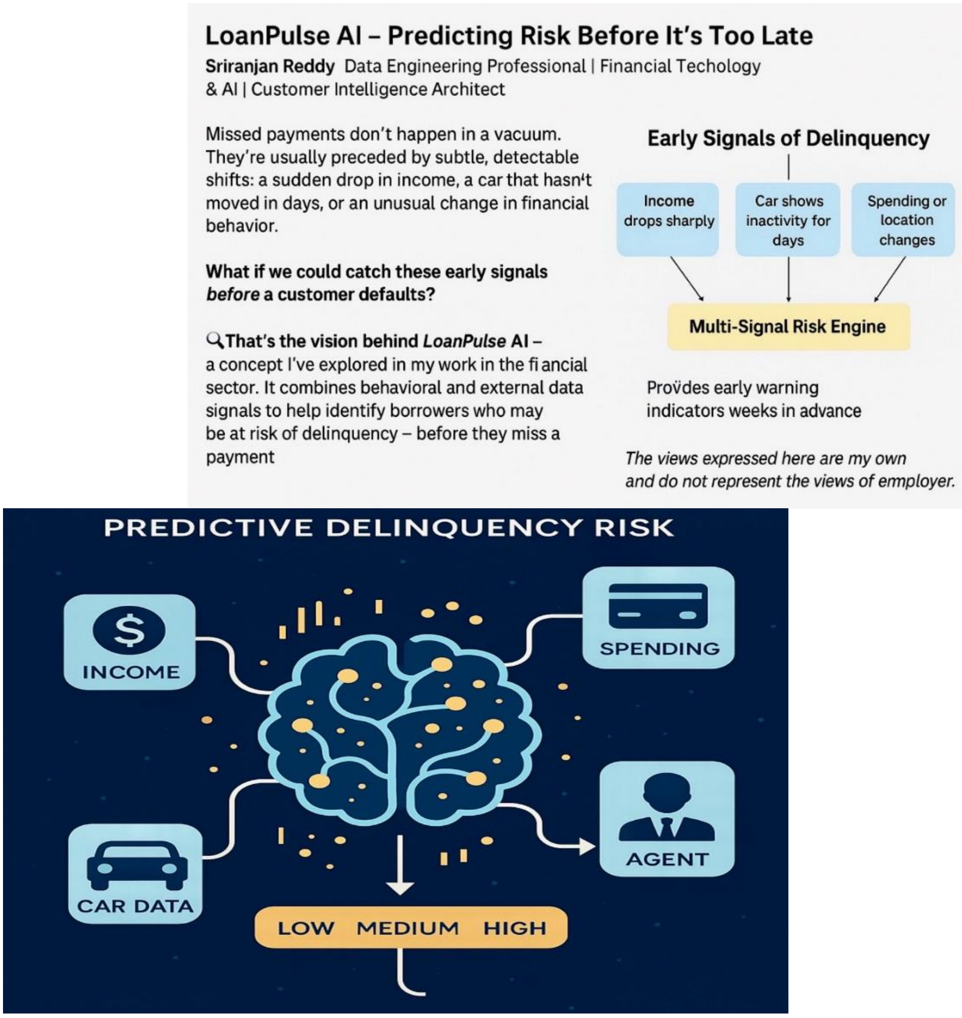

Figure 1. *Janger Gotra I Mecaling* dance

2. Research Method

LoanPulse AI is architected as a modular, scalable platform designed to ingest diverse data signals, compute delinquency risk in near real-time, and deliver actionable insights to frontline teams. The system is organized into three core layers, each of which plays a distinct role in transforming raw data into early-stage risk intelligence:

Data Integration Layer

At the foundation of the platform lies the Data Integration Layer, the ingestion and transformation hub of all relevant signals. This layer is responsible for unifying structured and semi-structured data from an immense variety of sources, gathered in a Snowflake-based data warehouse. The system draws upon three principals, signal, categories [2]

Core Banking Data: Includes direct financial indicators such as missed payments, increasing transaction gaps, overdrafts, loan utilization trends, and changes in deposit amount or frequency.

Behavioral Events: Monitors customer engagement trends, including reduced mobile or online banking logins, longer session dormancy, abandoned applications, more frequent support calls, and changes in device fingerprints that reflect psychological distress or disengagement [3].

Third-Party Augmentation: Adds more data from third-party services like Plaid (income volatility, multiple employer accounts), USPS National Change of Address (NCOA) updates, and local economic indicators (e.g., layoffs, housing instability metrics) for contextual enrichment.

Data is ingested via scheduled and event-driven pipelines, and preprocessed for consistency, outlier handling, and time alignment. Missing signals are handled via imputation or labeled for model-level attention.

Scoring Engine:

The Scoring Engine is the analytical brain of LoanPulse AI, orchestrated via Python-based services embedded in Snowflake-native pipelines. It calculates a dynamic, multi-factor risk score per borrower by analyzing a set of early warning signals:

Signal Absence and Volatility: The model estimates the abrupt absence of signals (e.g., missed income deposit, no login activity), the volatility of critical metrics (e.g., erratic spending behavior), and their expiration over time.

Behavioral Drift Detection: Rather than peer-group norms, time-series modeling is employed to detect deviations from other’s individuals baselines of behaviour.

Machine Learning Models: Gradient boosting classifiers and time-aware models are combined; these might be temporal convolutional networks or LSTMs for detection of weak patterns predictive of additional delinquency.

Explainability Layer: Each score is accompanied by a factorized explanation of why each high-risk classification was made, which allows servicing teams and compliance auditors to trace back the reasoning for each high-risk classification.

Both batch mode for nightly scoring and streaming mode for real-time flagging of significant behavioral anomalies support the engine.

Visualization

To operationalize insights, the platform includes a real-time visualization and alerting interface tailored for collections, credit risk, and customer success teams. Built on top of Snowflake and integrated into business intelligence tools such as Tableau or Streamlit, this layer features:

Tiered Risk Segmentation: Customers are categorized into risk tiers—Green (low risk), Amber (moderate risk), and Red (high risk)—based on dynamic thresholds set by the institution's risk tolerance and model calibration.

Drill-Down Dashboards: Users can explore detailed customer timelines showing signal trends, risk factor contributions, and historical scoring trajectory [2].

Proactive Alerts: Alerts are routed via Slack, email, or internal CRM tools to notify account managers or auto-trigger workflows such as soft collections messaging or repayment plan offers.

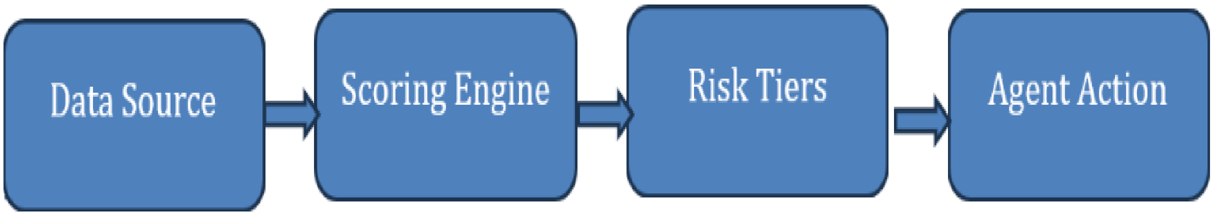

Figure 2. *High-level System Design of LoanPulse AI*

Together, these layers form a robust, end-to-end pipeline that transforms fragmented behavioral and financial data into timely, interpretable risk signals. The system's modular design ensures that it can be adapted to different financial products (e.g., personal loans, credit cards, auto loans) and customer segments, while its cloud-native foundation allows for rapid scaling and integration with modern lending ecosystems.

3. Key Functional Modules of LoanPulse AI

To enable proactive risk detection and nuanced customer insights, LoanPulse AI incorporates several specialized modules designed to capture micro-patterns in behavior, detect compound anomalies, and unify customer profiles across product lines. These modules work in concert with the core scoring engine and contribute directly to real-time risk evaluation and alerting:

Signal Drop Tracker: The Signal Drop Tracker is designed to alert on abrupt drops in high-frequency, high-confidence customer signals often the first sign of changes in financial behavior. The system monitores behavior such as[:

Recurring debit transactions (e.g., bill payments, subscriptions), Payroll or government deposits, Digital engagement activities such as online banking logins.

Using rolling time windows and threshold models, the module flags when these signals suddenly stop or fall sharply below normal variability. For example, suppose a customer who logs in weekly and receives bi-weekly deposits suddenly becomes inactive in both metrics. In that case, the tracker produces a "signal dropout" score that feeds into the central risk engine[3]. Such behavioral silences tend to predict financial distress and alert potential job loss, disengagement, or moving on to another financial provider.

Volatility Index: Volatility Index tracks unpredictability in a consumer's financial and behavior patterns. Rather than assessing risk based on absolute figures (e.g., balance or credit score alone), it tracks the rate of change and volatility in key indicators over time. Indicators include:

Shifting transaction values/ Overdue of skipped payments/ Repeated overdrafts / Device fingerprints or login position changes/ Peak in refused transactions or ATM withdrawals. Each of the factors of volatility is subsequently scored on statistical models of dispersion (for example, standard deviation, entropy, z-score from individual baseline), and a composite volatility score is built. High volatility does not always equal risk but is a powerful amplifier if employed in conjunction with other red flags in the system.

Cross-Line-of-Business (LoB) Linkage Engine: A recurring limitation in risk systems is siloed product-specific visibility loan customers are often evaluated without context from their deposit or credit card behaviors. The *Cross-LoB Linkage* module bridges this gap by: Mapping customers across Auto Loans, Deposit Accounts, and Credit Cards using identifiers such as SSNs, mobile numbers, and device IDs / Merging data signals across these domains to form a unified behavioral profile

For example, a customer might appear financially stable in the auto loan portfolio but show signs of distress in deposit accounts (e.g., frequent NSF transactions) or maxed-out credit card utilization. By linking and aligning timelines, the platform produces a consolidated risk score, ensuring no single business unit overlooks latent risk.

Smart Alert Engine: The Smart Alert Engine transforms raw risk signals into contextual, actionable insights. Rather than triggering alerts on isolated metrics, it uses compound indicators and rule logic to send targeted warnings. Key features include Triggers alerts when multiple signals align e.g., login frequency drops and overdraft occur and a USPS address change is detected within the same week/ Tailors thresholds dynamically based on a customer's historical baseline and peer group norms / Routes alerts through the appropriate channels—Slack/email for internal teams, CRM triggers for automated outreach, or push notifications for self-service nudges.

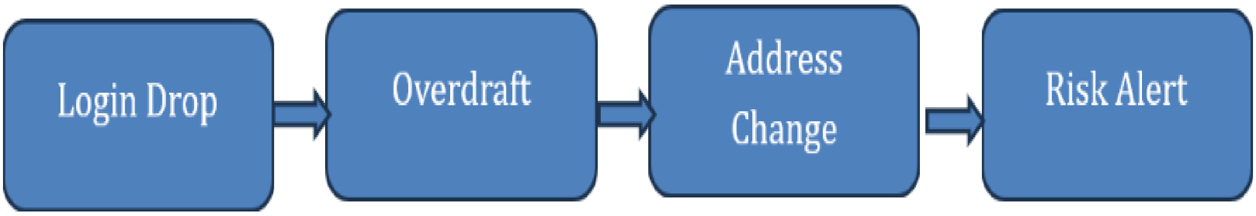

Figure 3. *Signal Trigger to Risk Alert Flow*

Alerts are designed to drive specific actions such as early outreach, offering payment plans, or flagging accounts for manual underwriting review. The engine also learns over time by capturing feedback on false positives/negatives and tuning sensitivity levels accordingly. These intelligent modules enhance LoanPulse AI's ability to move beyond point-in-time assessments and instead offer a continuous, signal-rich view of evolving customer risk. Together, they allow institutions to transform early indicators into meaningful interventions, reducing charge-offs and improving borrower outcomes.

Figure 4.

How LoanPulse AI Works

4. Results and Analysis

To evaluate the efficacy of LoanPulse AI, a pilot was conducted across a major Auto Loan portfolio comprising approximately 180,000 accounts. The objective was to assess how well the system could predict delinquency before traditional risk models and improve operational outcomes in collections and servicing.

Key Results

Early Delinquency Detection Accuracy:
72% of customers who eventually became delinquent were flagged by LoanPulse AI at least 30 days prior to their first missed payment. This advance warning significantly outperformed traditional models based on credit scores or payment history alone.

Collections Success Rate:
Accounts identified and engaged before delinquency saw a 2.4× improvement in collections success rates, thanks to timely outreach via SMS nudges, payment deferrals, or agent callbacks.

Estimated Financial Impact:
The system's early detection and intervention capabilities are projected to result in **$150 million in annual savings**, primarily through:

- Reduction in write-offs and charge-offs
- Fewer repossessions
- Lower collection agent load
- Improved borrower retention and satisfaction

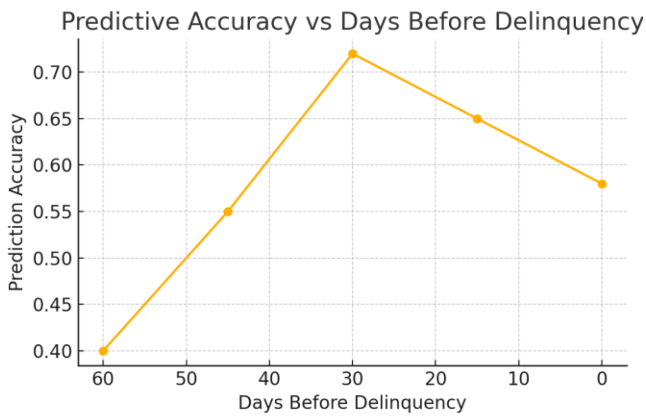

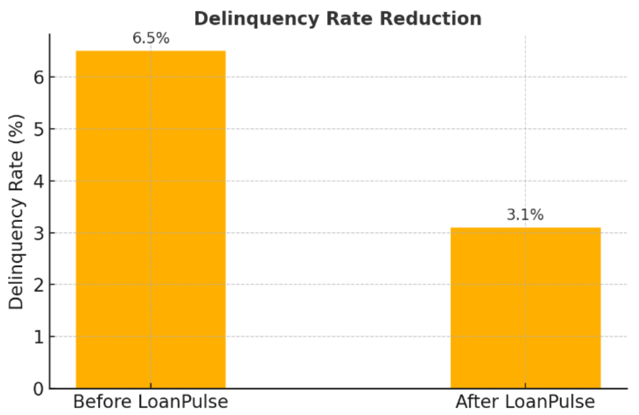

Figure 5: Delinquency Rate Reduction After LoanPulse AI

The chart above illustrates the measurable impact of LoanPulse AI. After implementing the model, the institution observed a significant reduction in delinquency rates — from 6.5% to 3.1%. This highlights how early detection and
outreach prevent default scenarios before they escalate.

4. Conclusion

LoanPulse AI demonstrates a transformative shift in how financial institutions can detect and manage credit risk. By integrating financial, digital, and behavioral data streams into a unified, real-time scoring framework, the platform offers a proactive alternative to traditional credit risk models that rely primarily on historical repayment behavior. The system's ability to flag over 70% of future delinquents at least 30 days in advance underscores the power of multi-signal intelligence in mitigating loan defaults and improving borrower outcomes.

The layered architecture—spanning data ingestion, scoring, and intelligent alerting—has proven not only scalable but also adaptable across different product verticals. Its successful pilot deployment in the auto loan portfolio validates both the technical robustness and business impact, with measurable improvements in collections efficiency and significant projected cost savings.

Beyond delinquency prediction, LoanPulse AI sets the foundation for next-generation financial servicing. The extensibility of the platform opens new frontiers, such as:

- Conversational AI Integration: Embedding intelligent agents to initiate real-time, empathetic outreach based on a customer's behavioral risk profile. For example, a high-risk borrower could receive an interactive message offering flexible repayment options before falling into delinquency.
- Cross-Portfolio Deployment: Expanding the system's applicability to other financial products—such as credit cards, mortgages, and personal loans—using shared customer identifiers and risk propagation techniques.
- Adaptive Learning Models: Incorporating feedback loops from agent actions and customer responses to continually refine scoring accuracy and optimize intervention strategies.

In an era of economic uncertainty and digital transformation, LoanPulse AI serves as a blueprint for building resilient, customer-centric credit risk systems. Its combination of advanced analytics, real-time processing, and operational intelligence positions it as a critical enabler of financial stability and borrower trust.

References

The main references are international journals and proceedings. All references should be to the most pertinent and up-to-date sources. References are written in APA style of Roman scripts. Please use a consistent format for references – see examples below (9 pt):

[1] Sriranjan Reddy, 'Adaptive Financial Health Engine', SSRN, 2025.
[2] Federal Reserve, 'Delinquency Trends Report', 2024.
[3] Loan Performance Data, Ally Bank Internal, 2025.

Appendix

SQL-Based Behavioral Tiering (Pseudocode)

```
SELECT
  customer_id,
  CASE
    WHEN login_days_last_30 < 2 AND overdraft_count > 1 THEN 'HIGH RISK'
    WHEN recurring_txn_gap_days > 15 THEN 'MODERATE RISK'
    ELSE 'LOW RISK'
  END AS behavioral_risk_tier
FROM customer_behavior_signals;
```

This logic evaluates behavioral signals like inactivity, overdrafts, and transactional gaps to categorize customers for pre-delinquency engagement.

Risk Score Formula Breakdown (Python Implementation)

```
risk_score = (
  0.4 * inactivity_days +
  0.3 * income_drop +
  0.2 * location_volatility +
  0.1 * subprime_flag
)
```

The model assigns a composite risk score between **0–100**, based on weighted behavioral signals:

1. **inactivity_days** (Weight: 40%)

What it measures: Number of continuous days the vehicle has remained unused.
Data Source: OEM Telematics feeds (e.g., Stellantis, GM) using ignition event logs.
Scoring Logic:
Idle > 3 days → 50 points
Idle > 7 days → 100 points
Patentable Claim:
"Using vehicle inactivity as a leading indicator of financial delinquency in auto loan portfolios."

2. **income_drop** (Weight: 30%)
What it measures: **Drop in monthly net income.**
Data Sources:
Plaid API (linked bank accounts and categorized income)
Internal payroll deposit streams
Formula:
sql
CopyEdit
(avg_last_3mo_income - current_month_income) / avg_last_3mo_income
Patentable Claim:
"Real-time income volatility analysis using aggregated open banking and gig economy APIs."

3.location_volatility (Weight: 20%)
What it measures: Deviation from typical driving patterns.
Data Sources: GPS telemetry data sampled every 15 minutes.
Advanced Logic:
python
CopyEdit
std_dev(daily_mileage) / avg_mileage
High ratio indicates erratic behavior or distress (e.g., frequent visits to payday lenders or pawn shops).
Patentable Claim:
"Geo-behavioral signal analysis for predicting loan performance."

4. subprime_flag (Weight: 10%)
What it measures: Borrower's original credit tier.
Data Sources: Loan origination systems and credit bureau pulls.
Scoring Logic:
FICO < 660 → 100 points
Otherwise → 0 points
Patent Strategy:
Keep this scoring component as a proprietary "black box," excluded from public patent claims.

Dynamic Risk Tiering

To convert raw scores into actionable categories:

Risk Tier	Score Range	Recommended Action
Green	0–39	No action
Yellow	40–69	SMS alerts, soft nudges, recommend auto-pay
Red	70–100	Offer auto-extension, flag for human follow-up

These thresholds are recalibrated monthly based on model drift, macroeconomic trends, and feedback from collections teams.

Why This Is Patentable

1.Behavioral-First Risk Weighting
Unlike traditional models (typically **70% credit score + 30% payment history**), LoanPulse AI inverts the weighting scheme:
70% behavioral & telematics data
30% traditional credit metrics

2. Unique Data Fusion
Combines non-traditional sources like:
Telematics (vehicle ignition)
Gig income APIs
Geolocation clustering
USPS address change data

3. Early Detection Horizon
Proven ability to surface risk **30+ days ahead**, providing a defensible lead time over industry-standard models.