Real-Time Anomaly Detection with AI

How machine learning helped reveal system failure in a leading payment technology company.

moviri
/ AI LABS

From theory to impact in mission-critical AI

In our work with enterprise leaders, we often discuss the architectural and strategic pillars required to build robust, enterprise-grade AI systems. We cover advanced models, data readiness, and MLOps best practices.

Now, let's move from theory to practice. This case study from the FinTech industry showcases how Moviri experts developed a high-performance, real-time anomaly detection system for a leading payments technology company. This project is a foundational example of the work we do at Moviri AI Labs.

We will explore the complex challenges you likely face in your own mission-critical environments and demonstrate how to overcome them to unlock substantial business value. This is intelligence, engineered for the enterprise.

Please contact us at ailabs@moviri.com to learn more about Moviri AI Labs.

Benedetta Masseroni
Service Line Manager, Analytics at Moviri

The customer and the challenge

About the customer

The customer is a critical infrastructure provider, offering **digital payment services** to a vast portfolio of businesses. Their offerings include issuing, merchant services, and ATM management, the processes that are essential to modern commerce.

Their systems handle **millions of transactions** daily across multiple channels and complex authorization workflows, demanding constant supervision to mitigate risks and technical failures. Given the sheer scale and complexity, a manual approach to oversight is simply impossible.

The customer's primary goal was to instantly identify any **unexpected surges or drops in transaction volumes** to enable rapid fault repair.

Moviri was tasked with engineering an anomaly detection system to monitor transaction volumes in near real-time, catching issues as they happened and preventing downstream impact.

The challenge

Why is this problem so difficult? The data involved single transactions that needed to be aggregated across dozens of different criteria, creating a multitude of distinct time series that each required individual inspection. The core requirements were:

- **Extreme Velocity**: The system had to produce outputs in near real-time while processing massive data volumes without sacrificing accuracy. Fancy, resource-intensive deep learning models were off the table due to long training times and limited computational resources.

- **Complex Seasonality**: Transactional data contains strong seasonal patterns (e.g., holiday sales peaks) that fool simple threshold-based alerting systems. Ad-hoc methods designed to handle these unique business rhythms were essential.

- **Actionable Explainability**: Stakeholders needed to understand why an anomaly was flagged and how critical it was to prioritize interventions. Model explainability wasn't a "nice-to-have"; it was crucial for operational trust and efficiency.

- **Low False Positives**: Alert fatigue is a real problem. The system had to be precise, generating the lowest possible number of false positives to ensure operators could focus on genuine threats.

In short, the challenge was to monitor a high-volume, heterogeneous data environment and determine, every ten minutes, if the transaction count was anomalous.

Solution strategy

Transaction volume as a constraint

The volume of data from the customer's payment infrastructure was immense: millions of daily transactions between consumers, banks, and merchants.

Our goal wasn't to flag every single odd transaction, but rather to **identify meaningful peaks and valleys** in transaction volumes within short timeframes.

A key constraint was that the historical data was unsupervised; we didn't know when or where past anomalies had occurred. We needed an algorithm that could thrive without this prior knowledge.

> Our goal wasn't to flag every single odd transaction, but rather to identify meaningful peaks and valleys in transaction volumes within short timeframes.

Modeling

After a thorough review of machine learning literature, our team selected **Seasonal Extreme Studentized Deviate (SESD)**. Developed by Twitter, SESD is an outlier detection technique ideal for this challenge. It works by decomposing a time series to isolate the residual component and then applies robust statistical tests.

This choice was driven by several factors:

- **Speed**: SESD is incredibly fast, processing thousands of data points in seconds and allowing for parallel execution across many time series to generate near-instantaneous output.
- **Seasonality Handling**: The algorithm intelligently separates and discards trend and seasonality, focusing its analysis only on the irregular, residual signal where true anomalies hide.
- **Unsupervised & Adaptive**: As a training-less, unsupervised method, SESD doesn't require labeled data and is inherently robust against data drift. It automatically adapts to the latest data patterns, ensuring its performance remains high over time.

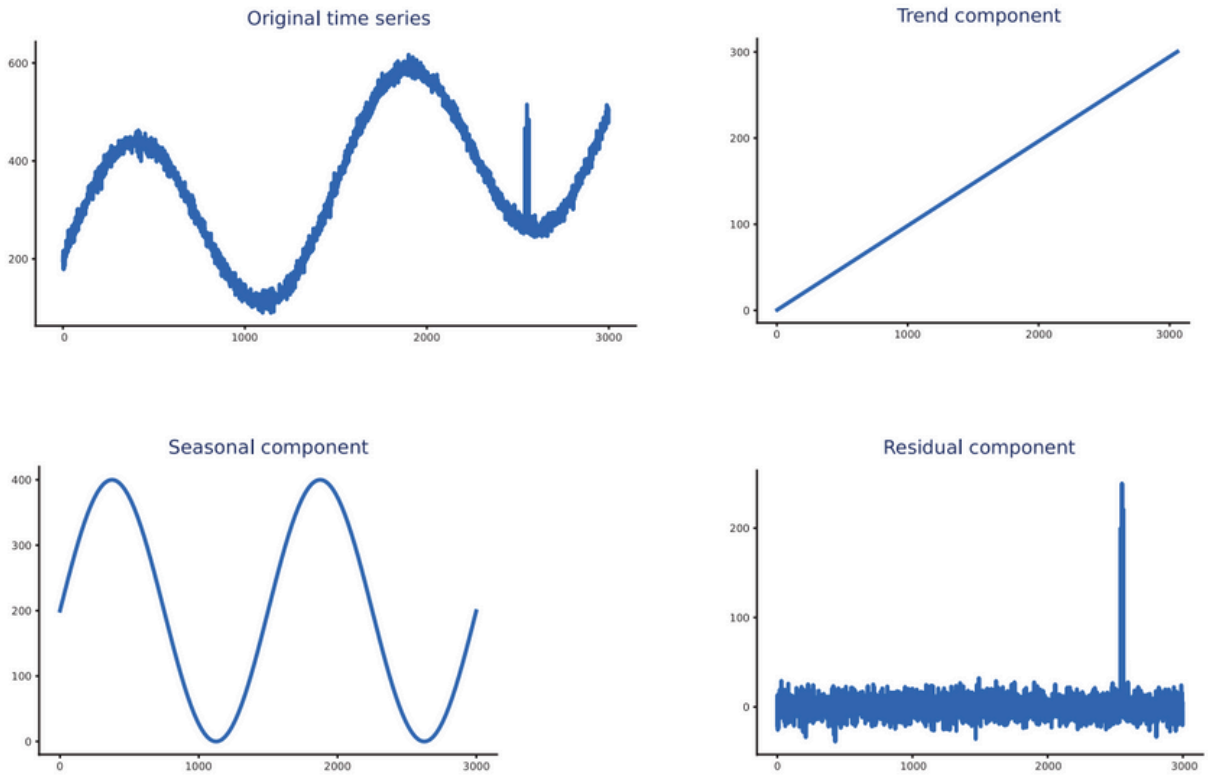

Our approach

We implemented a three-step, continuous pipeline:

- **Aggregation**: The system first collected granular transaction data and aggregated it into 10-minute buckets to form the time series to be analyzed.
- **Detection**: As soon as aggregation was complete, SESD ran in parallel across all time series, identifying and flagging anomalies.
- **Reporting**: The findings were instantly pushed to a user-friendly dashboard that summarized all detected anomalies, sorted by priority. The dashboard was designed for rapid investigation, allowing users to deep-dive into a specific customer, application, or system to understand the root cause.

What is SESD?

Seasonal ESD (Extreme Studentized Deviate) is an anomaly detection technique for time-series data that accounts for seasonality and trend. It combines the Extreme Studentized Deviate (ESD) test, used to find outliers in normally distributed data, with time-series decomposition methods.

Anomaly detection validation

For **unsupervised learning**, traditional validation is tricky. We couldn't use a simple validation set. Instead, we conducted extensive quality assurance tests with the customer's own domain experts. The results were definitive:

- The system independently identified all major anomalies that the customer's team had previously discovered manually.
- It also uncovered new, previously unknown anomalies that the experts later confirmed were genuine issues.

The entire pipeline takes roughly **10 minutes** to process **100,000 data points**. This means that within 10 minutes of an anomaly occurring, stakeholders are notified and can begin mitigation.

A note on MLOps

This project was deployed using MLOps principles to ensure long-term performance and reliability. The choice of a training-less algorithm like SESD was key, as it is naturally resistant to performance degradation and data drift—common challenges in production AI systems. This streamlined the deployment and delivered an efficient, resource-light, and highly effective solution.

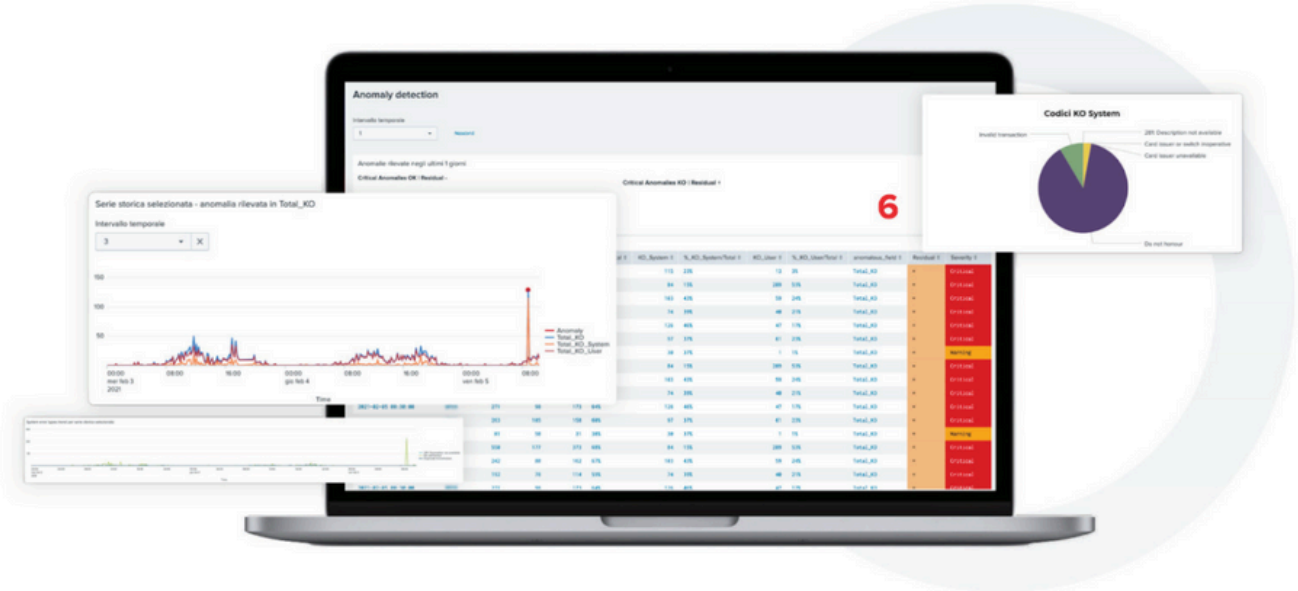

Results and benefits

The benefits of this enterprise-grade AI system were immediate.

1 Mil+	Transactions monitored per hour
600+	Anomalies detected daily
10 Min	Anomaly detection frequency on 100K transactions
100+	Daily person-hours saved

From a qualitative standpoint, the system delivered two transformative benefits:

- **Resilience**: It captures anomalies efficiently even during highly volatile periods, such as when the COVID-19 pandemic caused systematic shifts in consumer behavior.

- **Improved oversight**: It enhanced the automatic detection of underperforming transactional flows from certain banks, an issue previously overlooked by standard monitoring.

This success story demonstrates our ability to translate a critical business need into a fast, reliable, and scalable anomaly detection system.

By designing an ad-hoc solution and leveraging a modern MLOps framework, we delivered a system that had a major impact on core business KPIs.

moviri
/ AI LABS

Agentic AI for the enterprise

Agents are not just tools, but the new way to build IT systems and intelligent applications. At Moviri AI Labs we've created a cross-disciplinary team of analysts, engineers, researchers, and designers who develop and assemble AI agents, by using AI agents in the process.

Connect with Moviri AI Labs at ailabs.moviri.com.

HQ Milan
Via Schiaffino, 11
Milan, Italy 20158

Padova
Via Sant'Andrea, 21
Padova, Italy 35139

Boston
211 Congress Street
Boston, MA 02110

Los Angeles
12130 Millennium Dr
Los Angeles, CA 90094

Singapore
5 Temasek Blvd,
Singapore 038985