

FINANCIAL SERVICES • CASE STUDY

Sequence learning for fraud detection on a crypto exchange

INTRODUCTION

Our client is OVEX, a South African cryptocurrency exchange with a large retail and institutional customer base. Among the financial services offered on the exchange are cross-border payments, wholesale FX, as well as over-the-counter crypto and fiat trading.

Like every exchange, OVEX faces a problem in that its services can be used by criminals for fraudulent activities. Their compliance team handled roughly 1,400 confirmed fraud cases since 2023, the majority of which occurred in the last year.

98%

OF FRAUD CAUGHT IN TESTING

92%

REDUCTION IN UNDETECTED FRAUD

99%

FEWER FALSE ALARMS

Catch all fraud with high precision

At the start of OVEX's engagement with DataProphet, around 15% of retail accounts on the platform had been linked to fraud. Our goal was to replace the hard-coded detection rules with a model that learns how fraudsters and normal customers behave, and to only flag the former.

Existing detection rules, implemented in late 2025, successfully identified roughly 75% of confirmed fraud cases. However, these rules came at the cost of a very high false alarm rate of 93.6%. This highlights a common problem in fraud classification tasks: it is always possible to catch 100% of fraudsters by flagging every customer. The art lies in catching more while calling out fewer.

75%

FRAUD CAUGHT BY PRE-EXISTING RULES

93.6%

PRE-EXISTING FALSE-ALARM RATE


The cost of false alarms

OVEX's compliance team is tasked with investigating suspected fraudulent activity on the exchange. Performing a full investigation of an individual customer can take hours. Considering the 93.6% false alarm rate that their existing rules exhibit, the burden on the team adds up fast. If roughly 800 confirmed cases make up 6% of the total number of alarms raised in one year, then the system would have raised on the order of 10,000 alarms in total in the same period. Even on the conservative side, that amounts to tens of thousands of man-hours per year spent investigating honest customers.

Therefore, a key requirement of the project was to reduce the number of false alarms to an absolute minimum. Or, in technical terms, the model's precision would be maximised.

The cost of missed fraud

On the other hand, not raising the alarm when real fraud is taking place has severe reputational cost. South African banks have their own methods of identifying fraud and will issue return requests to exchanges when they suspect that fraudulent deposits were made from one of their accounts into an exchange. If the request reaches the exchange before the money has been withdrawn, the exchange complies and reverses the deposit. But if the fraudster has already withdrawn the money, the task of recovering the funds becomes much harder. From OVEX's point of view, identifying fraudulent behaviour and preventing the funds from leaving the exchange is critical.



Label pollution and sequence information

Label pollution

The nature of fraud is that it is invisible until discovered. Consequently, there are always accounts on the exchange that are assumed to belong to honest users, but in fact belong to criminals. The net effect is that the not-fraud labels assigned to the training data cannot be fully trusted. This effect does not apply to the fraud labels, since those cases have been vetted and confirmed. To address this concern, a number of mitigation strategies were applied during training.

Sequence information

Often in machine learning, it is sufficient to summarise a sequence of events into a single vector for the purpose of training a model. And while it is possible to capture information about the timing between events in a sequence in this way, the interpretation of this information by the machine learning algorithm often lacks fidelity. When the timing between logins and transactions, as well as the sequence in which certain activities are performed, are what sets normal and criminal behaviour apart, the one-dimensional approach quickly falters.


As a result, it was necessary to employ sequence models. In a sequence model, the notion of time is built into the model's architecture, and the training data changes from a vector to a sequence of vectors, each describing a particular event, rather than a summary of a sequence of events. This model produces a behaviour score, quantifying a particular customer's propensity for fraud, given the sequence of actions they have performed over the lifetime of their account.

The notion of time is built into the model's architecture. The training data is a **sequence of events**, not a summary of them.

An ensemble of behaviour and risk

Some information is not sequential. Financial institutions in South Africa are required to collect personal information about their customers. This data typically consists of identity documents, proof of address, contact details, and the like. While this information may change, it typically does not over short time scales, and almost never between consecutive transactions.

Once confirmed, the information can be used to generate a risk profile that is unlikely to change quickly. Training a separate risk model on static data and combining the risk score with a customer's behaviour score was a crucial component of this project.



Results

The table below shows the results for each of the models individually, as well as the performance of the ensemble, which combines a customer's static risk score with their behaviour score.

MODEL	FRAUD CAUGHT	FALSE ALARM RATE
Existing detection rules	75%	93.6%
Risk model (static KYC data)	88%	2%
Behaviour model (sequence)	99.3%	15%
Ensemble	98%	5%

KEY RESULTS

The key results of DataProphet's fraud detection deployment for our client were as follows:


- 01** Hard-coded detection rules replaced with models that learn how fraudsters and honest customers behave.
- 02** Sequence models that read the order and timing of account activity, not a summary of it.
- 03** A static KYC risk score combined with a behavioural score in a single ensemble model.
- 04** False alarms reduced substantially, returning investigator hours to genuine fraud cases.

98%

FRAUD CAUGHT, UP FROM 75%

5%

FALSE-ALARM RATE, DOWN FROM 93.6%



GET IN TOUCH

**Talk to us about the
fraud in your data**

EMAIL

contact@dataprophet.com

PHONE

+27 21 300 3555

WEB

www.dataprophet.com

