Classifying payment patterns with artificial neural networks: an autoencoder approach

Jeniffer Rubio, Paolo Barucca, Gerardo Gage, John Arroyo and Raúl Morales-Resendiz
Agradecimiento especial: Dr. Serafín Martínez
BCE-CEMLA – UCL

INTRODUCTION

- ❖ Monitoring payments and Financial Market Infrastructures is one of the primary objectives for central banking.

- ❖ Focus on the major FMI in Ecuador, the *Sistema de Pagos Interbancarios* (SPI)

- ❖ In Ecuador, the SPI annually channels more than USD 107 billion (equivalent to 1 times the country's GDP) through 69 million electronic transactions. It has more than 300 participating financial institutions, but only 24 banks channel the 90% of the total amount.

- ❖ Introduce and test a pattern recognition tool for anomaly detection based on an autoencoder architecture.

- ❖ It´s a different approach to measure systemic risk. Payment data provide an accurate and system-wide overview of how banks manage their liquidity over time

MAIN CONTRIBUTION

Our work is related to Triepels-Daniels-Heijmans (2018), Triepels-Heuver (2019), and Sabetti-Heijmans (2020).

Training autoencoders able to detect a wide range of anomalies in the SPI, ranging from spotting the anomalous behavior of individual banks to detecting changes in the overall activity of the payments network.

Novel techniques are robust enough to support payments' and market infrastructures' oversight, and ultimately to monitor financial stability

Autoencoder Model

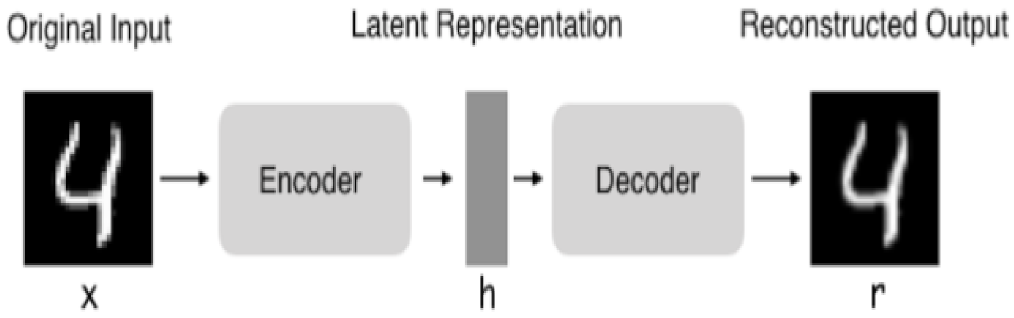

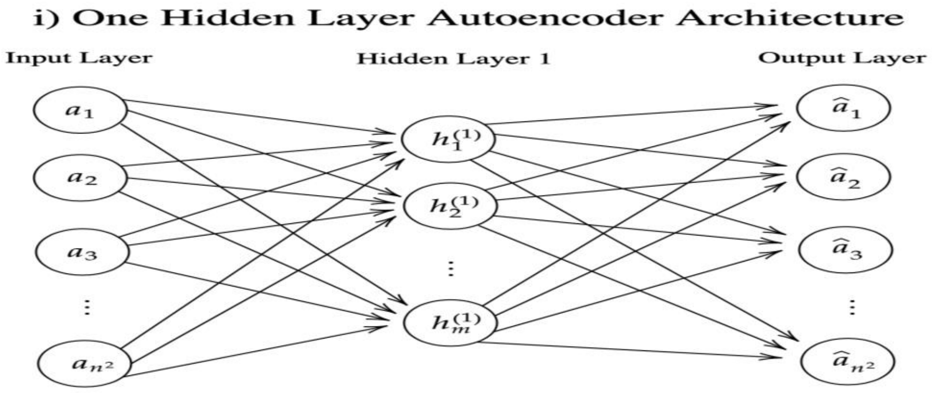

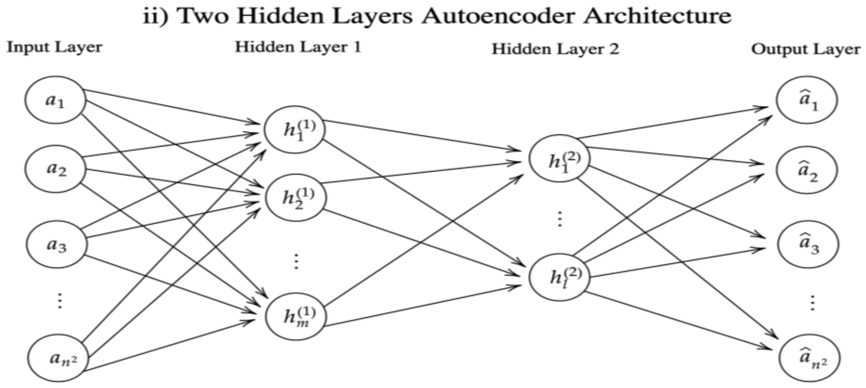

- *Unsupervised Neural Network (Unlabeled Data)*
- *It is made up of the input layer, hidden layer/s and output layer.*
- *The autoencoder is trained to reconstruct the input data in the output, based on the compression in the bottleneck layer, with them the most relevant characteristics (patterns) of the input are obtained.*
- *Input layer dimension is less than hidden layer / s (x <h) which allows compressing input data.*
- *Activation function*
- *Performance metric of the the mean reconstruction error (MSE) that measures how well the neural network reproduces the input layer.*
- ***Anomaly detection:*** *MSE > umbral (ε)= anomaly.* If the reconstruction error (MSE), it is a frequently recurring pattern that the compression model has learned to compress well. If it´s Is large, then the model does not recognize the liquidity flows and fails to reconstruct their values. **IT IS AN ANOMALY!**

Data

$B = \{b_1, \ldots b_n\}$ is a set of n banks participating in a interbank payment system (SPI)

$T = \{t_1, \ldots, t_m\}$ is an ordered set of **m** time intervals

We extract $D = \{A_1, \ldots A_m\}$ a set of **m** liquidity matrices from a **SPI** system where each $A^{(k)} \in D$ is the $n * n$: matrix:

$$A^{(k)} = \begin{bmatrix} a_{11}^{(k)} & \cdots & a_{1n}^{(k)} \\ \vdots & \ddots & \vdots \\ a_{n1}^{(k)} & \cdots & a_{nn}^{(k)} \end{bmatrix} \qquad (1)$$

Each element $a_{ij}^{(k)} \in [0, +\infty)$ is the liquidity flow between b_i and b_j at t_k. $A^{(k)}$ is a liquidity matrix

For analysis purposes, Liquidity Vector:

$$a^{(k)} = \left[a_{11}^{(k)}, \ldots, a_{n1}^{(k)}, \ldots, a_{1n}^{(k)}, \ldots, a_{nn}^{(k)}\right]^T \qquad (2)$$

Data used:

- 24 banks (24x24 = 576 flows of payments)
- 741 times intervals in a year

Model fitting

Data Partition

80% of the original data (to perform cross-validation) and 20% (to test)

4 Models

re-trained autoencoder (80% data) hyper-parameters

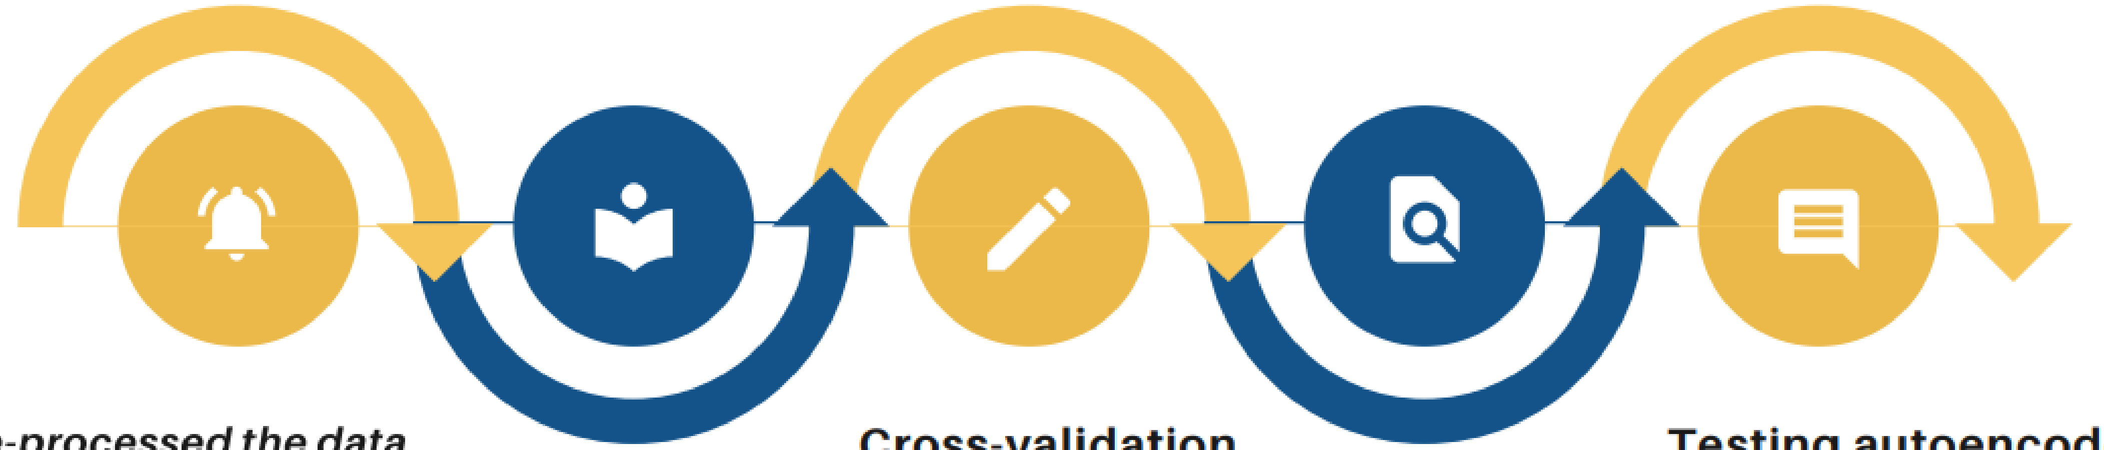

Pre-processed the data

log-transformation and that was followed by a min-max standardization

Cross-validation

5-fold cross-validation

Testing autoencoder

autoencoder with the unobserved data (20%)

Set the best model configuration
Número optimo de neuronas
80% data

Cross-validation (neuron numbers optimization)

5-fold cross-validation to determine the optimal number of neurons for the hidden layers

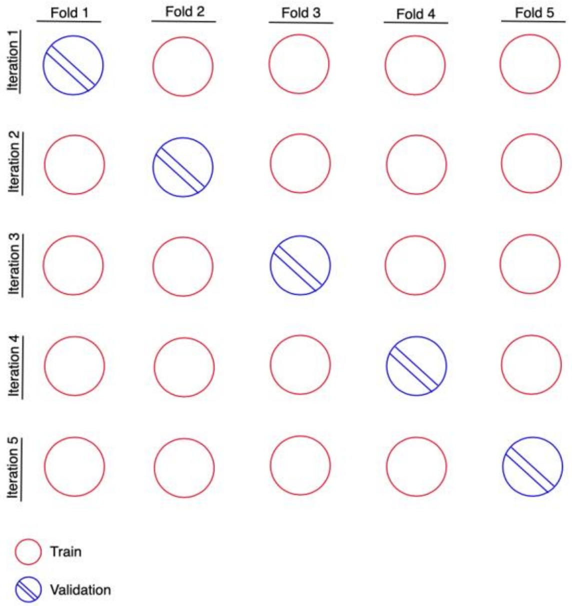

	Activation Function	Neurons in input layer	Neurons in first hidden layer	Neurons in second hidden layer	Neurons in output layer
Model 1	Tanh	576	(10, 20, 30,..., 450)	----------------	576
Model 2	Tanh	576	(10, 20, 30,..., 400)	(8, 16, 32)	576
Model 3	ReLU	576	(10, 20, 30,..., 450)	----------------	576
Model 4	ReLU	576	(10, 20, 30,..., 400)	(8, 16, 32)	576

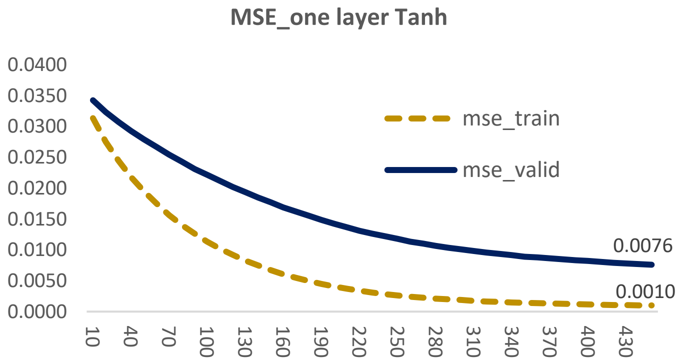

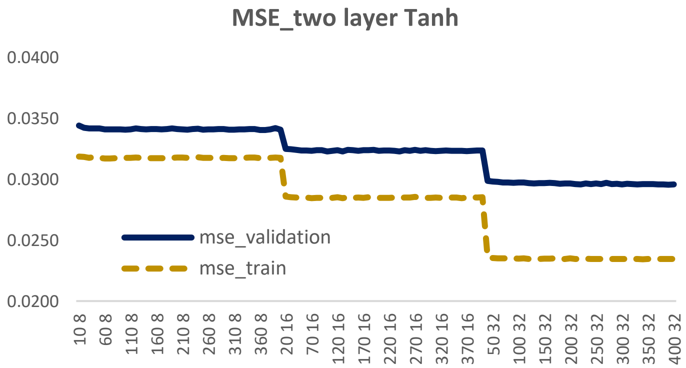

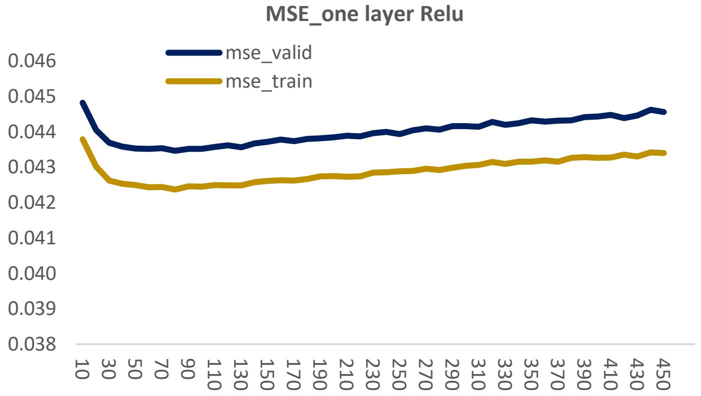

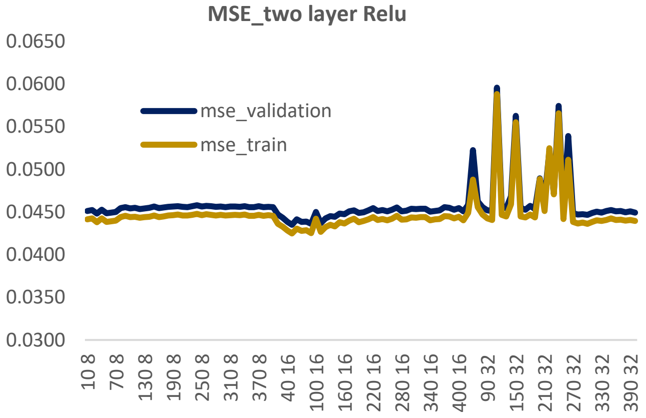

Optimal number of neurons when MSE reaches its minimum or does not have greater variation with increases in neurons

Validation Autoencoder

Re-train the models with unseen data (Test data set)

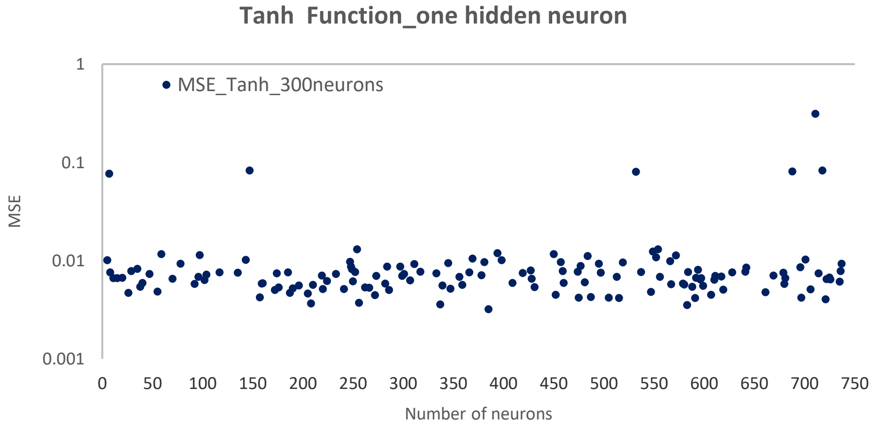

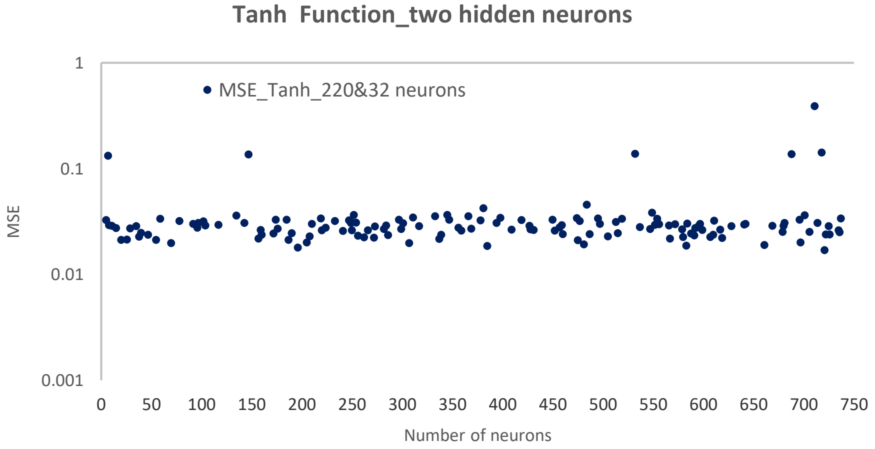

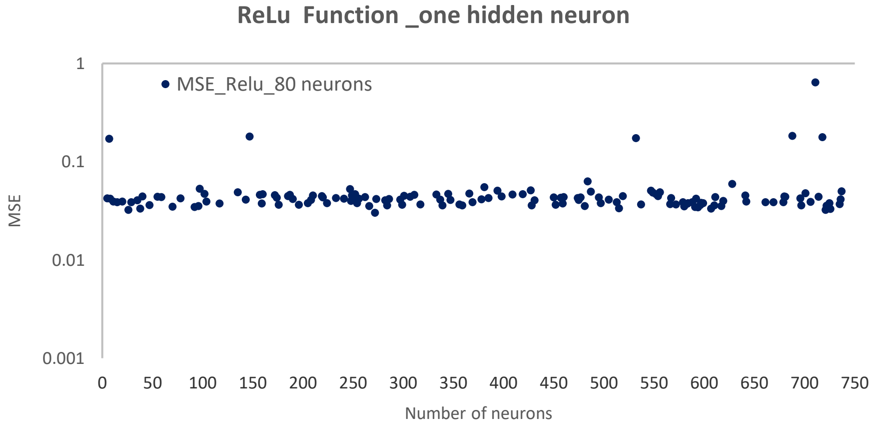

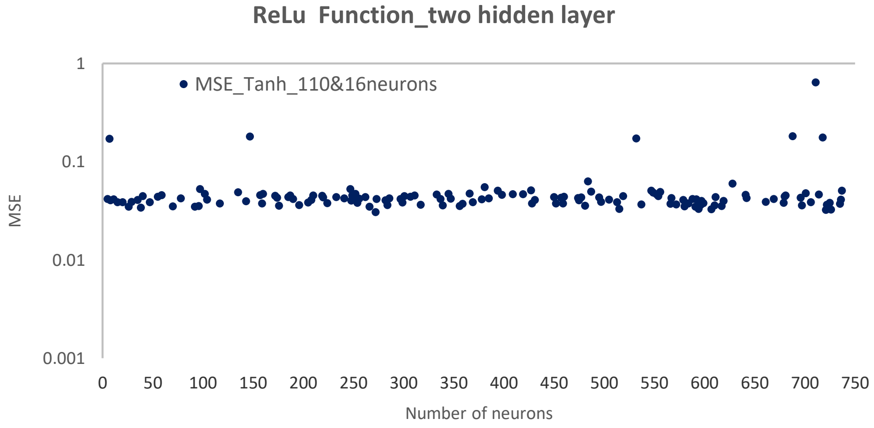

Alert analysis by a supervisor

Investigamos cuál de los flujos (uno o varios) causó estas anomalías?

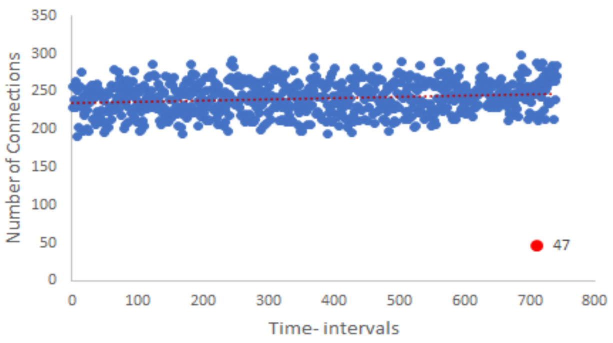

Systemic alert

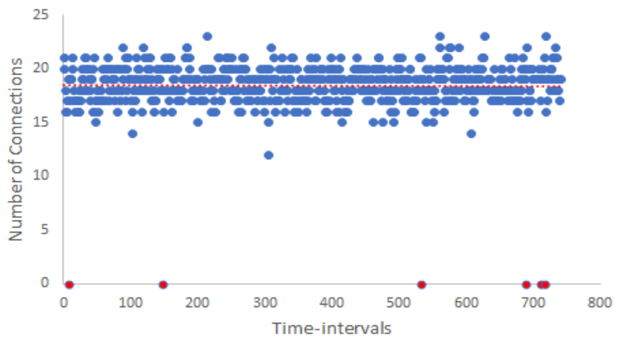

Individual alert for a large participant

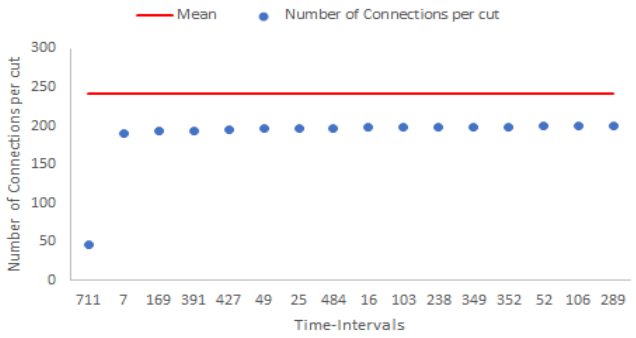

Alerts of low number of participants and connections

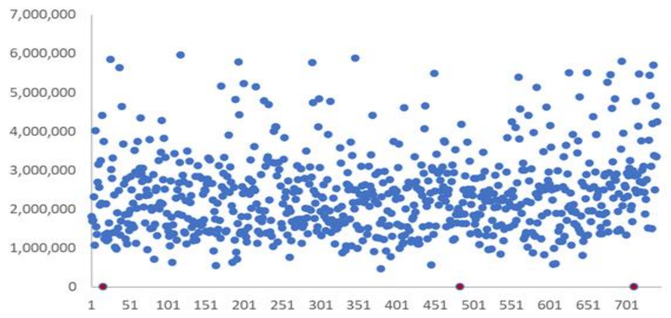

Problem with communication provider and impact on a large bank

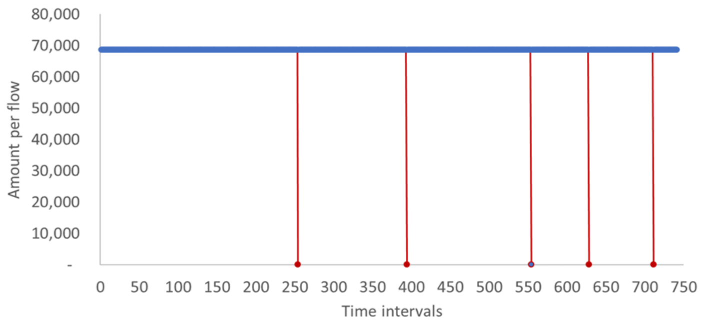

Individual alert for a SPI "average Joe" participant

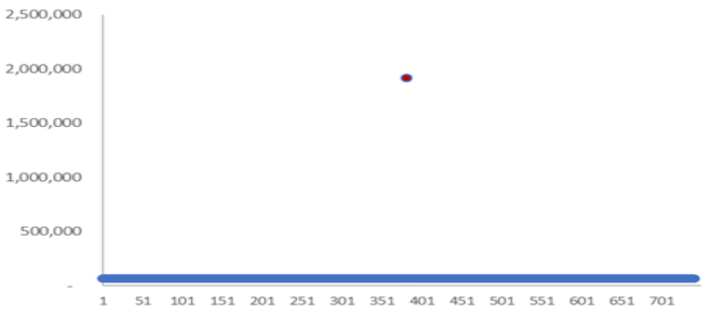

Unusual payment amounts

Bank runs simulations testing

The simulation* consisted of stressing institution outflows toward the rest of participants.

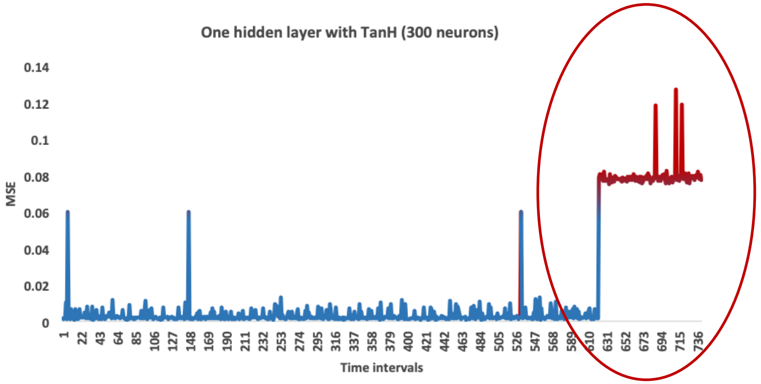

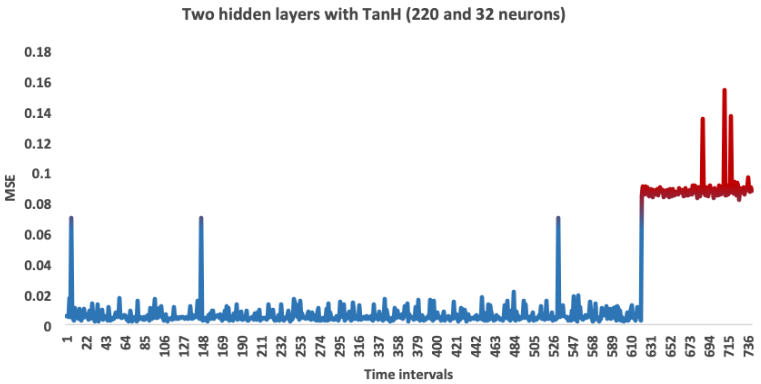

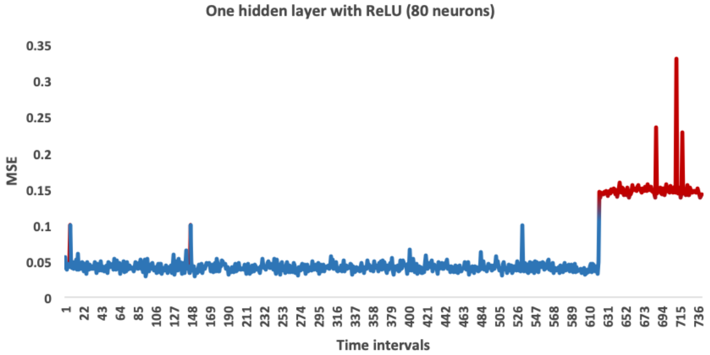

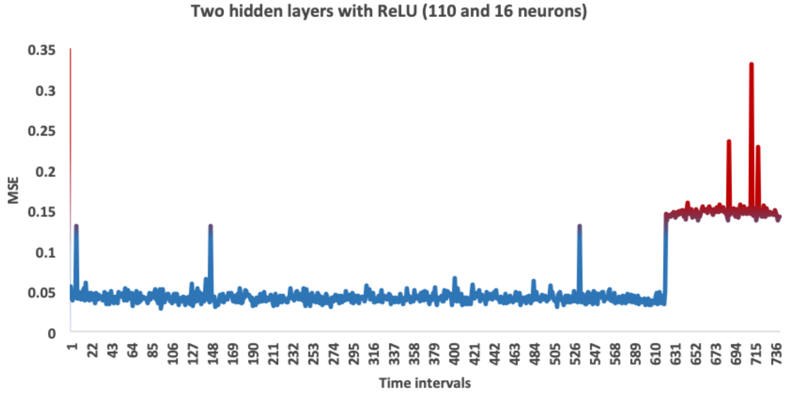

MSE rapidly changed as the payment network unexpectedly began to change as well, autoencoder was able to flag them as anomalies

*We introduce the bank run as a random value with exponential distribution and a probability of occurrence for the selected *bi* outflows, sampled from a Bernoulli distribution.

Conclusions

- The autoencoder needs to be applied in tandem with the thoughtful review of a payment systems oversight team, to verify the real causes of the alert.

- The construction and training of the model is a careful process involving numerous validations and tests that must be carried out, but once the model has been trained its daily application in detecting anomalies in SIPS and FMI operations can take only minutes.

- This document contributes to identify evidence of incidents in the payments flow of a respective system, and thereby provides new tools for payments oversight, and it ultimately sets the basis for early warning tools.

- Identify alerts that could affect the system, such as: 1) non-participation of systemically important participants, 2) a low number of connections (payment flows), 3) medium-size banks not sending payments to systemically important participants, among others.

- Studies in machine learning for anomaly detection in payment systems can improve the accuracy, reliability, and speed of the methodology.