Detecting Fraud in Financial Payments

Noël Dutrée[1] & Dirk Hofland[2]

Motivation

Under PwC's Global Economic Crime Survey 36% of the organizations surveyed reported being victimized by economic crime in 2016. As approaches to fraud become increasingly sophisticated, more incidents of economic crime go unnoticed, so that the actual financial impact of economic crime is likely far greater. Most companies still detect fraud predominantly by accident instead of instituting comprehensive and systematic fraud prevention policies. For this project, we applied machine learning techniques to develop a model that can detect fraudulent payments in real-time. To our knowledge, no such model exists as of yet.

Dataset

We worked with a dataset provided by Adyen, a global payments company to which we have a prior affiliation. This set contains data from approximately 300.000 transactions from across the globe, including whether the transaction was deemed fraudulent, non-fraudulent, or undetermined. The dataset is heavily imbalanced with only 0.17% of the transactions being fraudulent. The set of 15 features includes amount, currency, geographical location as well as details on the payment method and relevant timestamps.

Model

The neural network uses a ReLU activation function in the hidden layer and a sigmoid function for the output layer. Batch gradient descent is used for updates. The focal loss is defined as follows:

$$FL(p_t) = -\alpha(1-p_t)^{\gamma}\, y \log(p_t) - (1-\alpha){p_t}^{\gamma}(1-y)\log(1-p_t)$$

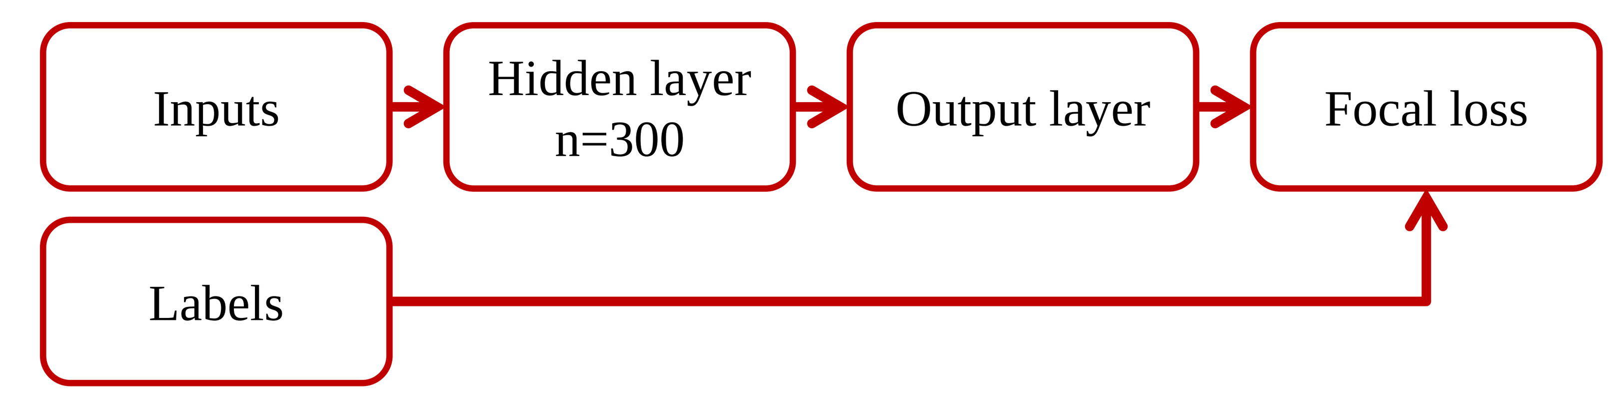

Results

Model	Test accuracy Overall (%)	Fraud (%)	Non-Fraud (%)
Naïve Bayes*	69	90	69
Weighted Logistic Regression	82	80	82
SVM	100	0	100
Plain NN	63	92	63
NN with feature selection	80	78	80
NN with regularization	70	96	70
NN with fraud oversampling	81	70	81
NN with weighted loss function	83	71	83
NN with focal loss, feature selection, regularization, and fraud oversampling	**86**	**76**	**86**

*unstable accuracy

Discussion

The imbalanced nature of the set is the major challenge. As expected, weighted loss functions are essential for decent results. Moreover, several attempts at over- and undersampling yielded improved results over the baseline. We observed significant improvements from a combination of oversampling of the minority class (frauds) and modified α-balancing based on the oversampling bias. The best results however, were obtained by combining the aforementioned with a novel loss function known as 'focal loss', that specifically focuses on minimizing false positives/negatives.

Future

Future work would focus primarily on determining approaches to deal with the set of examples that are classified as undetermined. Moreover, exploring other pertinent features to add, performing a thorough optimization process, and performing real-time tests at Adyen are likely to further improve accuracy. Finally, a financial loss decision system could be looked into, that determines whether or not to follow up on detected frauds based on the confidence and financial loss.

References

T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár. "Focal loss for dense object detection," *in ICCV*, 2017.

[1]ndutree@stanford.edu
[2]dhofland@stanford.edu