Université
franco-allemande
Deutsch-Französische
Hochschule

Context-Aware Credit Card Fraud Detection

Credit card fraud has emerged as major problem in the electronic payment sector. In this thesis, we study data-driven fraud detection and address several of its intricate challenges by means of machine learning methods with the goal to identify fraudulent transactions that have been issued illegitimately on behalf of the rightful card owner. In particular, we explore several means to leverage contextual information beyond a transaction's basic attributes on the transaction level, sequence level and user level.

On the transaction level, we aim to identify fraudulent transactions which, in terms of their attribute values, are globally distinguishable from genuine transactions. We provide an empirical study of the influence of class imbalance and forecasting horizons on the classification performance of a random forest classifier. We augment transactions with additional features extracted from external knowledge sources and show that external information about countries and calendar events improves classification performance most noticeably on card-not-present transaction.

On the sequence level, we aim to detect frauds that are inconspicuous in the background of all transactions but peculiar with respect to the short-term sequence they appear in. We use a Long Short-term Memory network (LSTM) for modeling the sequential succession of transactions. Our results suggest that LSTM-based modeling is a promising strategy for characterizing sequences of card-present transactions but it is not adequate for card-not-present transactions.

On the user level, we elaborate on feature aggregations and propose a flexible concept allowing us to define numerous features by means of a simple syntax. We provide a CUDA-based implementation for the computationally expensive extraction with a speed-up of two orders of magnitude compared to a single core CPU-based implementation. Our feature selection study reveals that aggregates extracted from users' transaction sequences are more useful than those extracted from merchant sequences. Moreover, we discover multiple sets of candidate features with equivalent performance as manually engineered aggregates while being vastly different in terms of their structure.

Regarding future work, we motivate the usage of simple and transparent machine learning methods for credit card fraud detection and we sketch a simple user-focused modeling approach.