



Towards automated feature engineering for credit card fraud detection using multi-perspective HMMs

Yvan Lucas, Pierre-Edouard Portier, Léa Laporte, Liyun He-Guelton, Olivier Caelen, Michael Granitzer, Sylvie Calabretto

► To cite this version:

Yvan Lucas, Pierre-Edouard Portier, Léa Laporte, Liyun He-Guelton, Olivier Caelen, et al.. Towards automated feature engineering for credit card fraud detection using multi-perspective HMMs. Future Generation Computer Systems, 2020, 102, pp.393-402. 10.1016/j.future.2019.08.029 . hal-02278223

HAL Id: hal-02278223

<https://hal.science/hal-02278223>

Submitted on 20 Dec 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial 4.0 International License

Towards automated feature engineering for credit card fraud detection using multi-perspective HMMs

Yvan Lucas^{1,2}, Pierre-Edouard Portier¹, Léa Laporte¹, Liyun He-Guelton³, Olivier Caelen³, Michael Granitzer², and Sylvie Calabretto¹

¹INSA Lyon

²Universität Passau

³Worldline Lyon

Abstract

Machine learning and data mining techniques have been used extensively in order to detect credit card frauds. However, most studies consider credit card transactions as isolated events and not as a sequence of transactions.

In this framework, we model a sequence of credit card transactions from three different perspectives, namely (i) The sequence contains or doesn't contain a fraud (ii) The sequence is obtained by fixing the cardholder or the payment terminal (iii) It is a sequence of spent amount or of elapsed time between the current and previous transactions. Combinations of the three binary perspectives give eight sets of sequences from the (training) set of transactions. Each one of these sequences is modelled with a Hidden Markov Model (HMM). Each HMM associates a likelihood to a transaction given its sequence of previous transactions. These likelihoods are used as additional features in a Random Forest classifier for fraud detection.

Our multiple perspectives HMM-based approach offers automated feature engineering to model temporal correlations so as to improve the effectiveness of the classification task and allows for an increase in the detection of fraudulent transactions when combined with the state of the art expert based feature engineering strategy for credit card fraud detection.

In extension to previous works, we show that this approach goes beyond ecommerce transactions and provides a robust feature engineering over different datasets, hyperparameters and classifiers. Moreover, we compare strategies to deal with structural missing values.

Introduction

Credit card fraud detection presents several difficulties. One of them is the fact that the feature set describing a credit card transaction usually ignores detailed sequential information. Typical models only use raw transactional features, such as time, amount, merchant category, etc Donato et al. (1999). Bolton and Hand (2001) showed the necessity to use attributes describing the history of the transaction when they used unsupervised methods such as peer group analysis for credit card fraud detection. Consequently, Whitrow et al. (2008) create descriptive statistics as features in order to include historical knowledge. These descriptive features can be for example the number of transactions or the total amount spent by the card-holder in the past 24 hours for a given merchant category or country. Bahnsen et al. (2016) considered Whitrow et al. (2008) strategy to epitomized the state of the art feature engineering technique for credit card fraud detection.

We identified several weaknesses in the construction of these features that motivated our work. First, descriptive statistics provide an aggregated view over a set of transactions. Such aggregated features do not consider fine-grained temporal dependencies between the transactions. For example, a common fraud pattern starts with low amount transactions for testing the card, followed by high amount transaction to empty the account. Second, aggregated features are usually calculated over transactions occurring in a fixed time window (e.g. 24 h). In general, transactions from very different card holders do not follow such a time pattern in general. However, the number of transactions made during such a time period can vary a lot for different card-holders. Fixed size aggregated statistics can't account for that fact. Third, these features consider only the history of the card-holder and do not exploit information of fraudulent transactions for feature engineering. However, a sequence of transactions happening at a fixed terminal can also contain valuable patterns for fraud detection.

In our work we propose to generate history-based features using Hidden Markov Models (HMM). They quantify the similarity between an observed sequence and the sequences of past fraudulent or genuine transactions observed for the cardholders or the terminals.

We have chosen these three perspectives based on the following assumptions: the features made with only genuine historical transactions should model the common behavior of card holders. This is a classic anomaly detection scheme where the likelihood of a new sequence of transactions is measured against historical honest sequences of transactions Chandola et al. (2012). The other four features are made with sequences of trans-

actions with at least one fraudulent transaction. The rationale is that to have a risk of fraud, it is not enough for a new sequence to be far from usual transaction behavior but it is also expected for it to be relatively close to a risky behavior. Thus, these last features should decrease the number of false positive. As stated by Pozzolo et al. (2017), this is a crucial issue since the investigators can only verify a limited number of alerts each day. The second perspective allows the model to take the point of view of the card-holder and the merchant which are the two actors involved in the credit card transactions. The last perspective takes into account two important features for credit card fraud detection: the amount of a transaction and the elapsed time between two transactions. These features are strong indicators for fraud detection.

By using a **real world** credit card transactions dataset provided and labeled by a large European card processing company, we want to assess how much our contributions - namely the addition of the terminal perspectives (among others) in the construction of history features and the construction of multiple perspectives HMM-based features - improve fraud detection.

To quantify the impact of the addition of the HMM-based features, we use the Precision-Recall AUC metric. We fed Random Forest classifiers with transactional data including the state of the art transaction aggregation strategy and measured the increase in PR-AUC due to the HMM-based feature engineering technique we propose. The multiple perspective property of our HMM-based feature engineering allows for an incorporation of a broad spectrum of sequential information that leads to a significant increase of the detection of fraudulent transactions.

This paper consists in an extension to a poster Lucas et al. (2019). As in Lucas et al. (2019), we present the concept of multiperspective HMM-based feature engineering. Moreover, we extend it significantly with additional experiments and evaluations. First of all, the framework is shown to increase the detection of fraudulent transactions not only for e-commerce transactions but also for face-to-face transactions. This result wasn't necessarily a foregone conclusion since e-commerce and face-to-face transactions present very different properties (e.g. merchant not open at night, necessity of a PIN authentication for face-to-face transactions...). Previous work done on the same dataset Jurgovsky et al. (2018) showed that some increase in detection observed on one type of transactions couldn't be extended to all types of transactions. Then, the feature engineering strategy is shown to be relevant for various types of classifiers (random forest, logistic regression and Adaboost) and robust to hyperparameters choices made for constructing the features: The number of hidden state and the length of the sequence have

no strong effect on the quality of the fraud detection. Lastly, the framework suffered from a structural missing value limitation: For some users with few transactions, the HMM-based features couldn't be calculated. From 20% to 40% of the transactions (depending on the choice of the length of the sequences modeled by the HMM) are not associated with HMM-based features. After comparing several solutions to overcome this limitation, we were able to obtain a steady improvement of the detection over all the transactions of the dataset.

In this paper, we present the state-of-the-art approaches for the problems of credit card fraud detection and sequence classification in section 1. Afterwards, we show how the HMM-based features improve on the limitations of the state of the art techniques and the way they are created in section 2. The experimental protocol is described in section 3. In section 4.1 and 4.2, we present the experimental results obtained on face-to-face and e-commerce transactions with different classifiers. The section 4.3 is dedicated to prove the robustness of the method through an hyperparameter study. Finally, in section 5 we compare different solutions to tackle the issue of structural missing values.

The HMM-based features we propose present interesting assets in the context of credit card fraud detection and more generally anomaly detection.

This work opens perspectives for feature engineering in any supervised task with sequential data. In order to ensure reproducibility, the source code of the proposed framework can be found at https://gitlab.com/Yvan_Lucas/hmm-ccfd.

1 Related Work

1.1 Credit card fraud detection

A wide range of machine learning approaches were used in credit card fraud detection. (Maes et al., 2002) evaluated Artificial Neural Network and Bayesian belief network with ROC AUC on Europay International's dataset. (Bhattacharyya et al., 2011) compared Support Vector Machine, Random Forest and logistic regression on a real world dataset using a wide variety of metrics. (Bahnsen et al., 2013) adjusted Bayes Minimum Risk using real financial costs in order to adapt the prediction of Random Forest and Linear Regression classifier. (Pozzolo et al., 2014) tested an architecture to take into account temporal concept drift in the credit card transaction data stream with Random Forest, Support Vector Machine and Neural Network. (Mahmoudi and Duman, 2015) applied a modified Fisher discriminant func-

tion to take into account the higher false negative cost in credit card fraud detection. More recently, (Jurgovsky et al., 2018) used LSTM for sequence classification on the same real world dataset that we use in this article. They showed that, in the case of face-to-face transactions only, sequence modelling with Long Short Term Memory networks (LSTM) improves fraud detection when compared to Random Forest with aggregated features.

Since Random Forest have been shown to perform well for credit card fraud detection in the literature (Bhattacharyya et al., 2011) and in preliminary experiments we have done, we chose to consider them in this work for evaluating the impact of our proposed features on the prediction quality. Moreover, Random Forests offer the possibility to calculate the importance of a feature which is defined as the decrease of gini impurity through a node weighted by the proportion of elements of the dataset passing through this node (Breiman et al., 1984). This property is interesting for studying the impact of a feature engineering strategy.

Feature engineering is critical for credit card fraud detection. Some authors used only raw features in order to detect fraudulent transactions ((Mahmoudi and Duman, 2015), (Minegishi and Niimi, 2011)). (Bolton and Hand, 2001) showed the necessity to use attributes describing the history of the transactions for unsupervised credit card fraud detection (peer group analysis). Lately, Saia and Carta (2019) used Fourier and wavelet transforms in order to move the transaction in a new domain before applying a machine learning algorithm. This allows to raise outliers based on a different view of the dataset (frequential view). This is related to our approach since our approach consists in creating likelihood score for a variety of views on the dataset (sequential views).

(Whitrow et al., 2008) proposed a transaction aggregation strategy to create descriptive features containing information about the past behaviour of the card-holder over a certain period of time. These descriptive features can be for example the number of transactions or the total amount spent by the card-holder in the past 24 h with the same merchant category or country. They showed a 28% increase in the detection of fraudulent transactions by using these aggregated features with a Random Forest as the learning algorithm. (Jha et al., 2012) applied Whitrow’s transaction aggregation strategy to logistic regression in a real world credit card transactions dataset. (Krivko, 2010) presented a rule-based approach using the difference between the recent amount spent by the card holder and either the average amount spent by this card holder or the average amount spent by all the card holders. (Bahnsen et al., 2016) showed that adding periodic features based on the time of the transaction to Whitrow’s aggregated features increases

the savings by an average of 13% with Random Forest, Logistic Regression and Bayes minimum Risk model.

An other feature engineering strategy has been proposed by (Vlasselaer et al., 2015). It consists in using the numbers of transactions between card holders and terminals in the graph of the transactions in order to create a time-dependent suspiciousness score. They showed a 3.4% increase of the ROC AUC by using these network-based features together with Whitrow’s aggregated features and Random Forests.

In this work, similar to other contributors, we consider Whitrow’s card-holder centric aggregated features as the state of the art baseline for comparison with the HMM-based features we propose.

1.2 Sequence Classification

Sequence classification is one of the main machine-learning research field. It aims to consider the sequential properties of the data at the algorithmic level in order to improve the classification of sequential data. Dietterich (2002) reviewed sequential classification based on sliding windows (or recurrent sliding windows). However, sliding windows methods don’t take into account inner dependencies between consecutive events.

Strivastava et al. (2008) tried to overcome this limitation by using generative models such as Hidden Markov Models (HMMs) for credit card fraud detection. They motivated the choice of HMMs by relating the hidden states to the different types of purchase. For this purpose, they created an artificial credit card transactions dataset. In their multinomial HMMs, the transactions were characterized with a symbol (‘big amount’, ‘medium amount’, ‘small amount’) used as the observed variable. After training, the likelihood that the sequence of recent transactions was generated by the HMMs is computed. The decision is taken by comparing the likelihood to a threshold value.

Graves (2012) observed that Long Short Term Memory networks are better than other sequential algorithms such as HMMs for speech recognition and handwriting recognition tasks since they allow to learn long term dependencies in sequences. Wiese and Omlin (2009) compared feed forward neural network combined with Support Vector Machines and LSTM for credit card fraud detection and showed that LSTM are relevant in the credit card fraud detection context because they can model time series of different length for each card holder. However, Jurgovsky et al. (2018) showed recently on the real world dataset that we use that LSTM offers a small improvement over Random Forest in the case of face-to-face transactions. There is no

improvement for e-commerce transactions.

2 Multiple perspectives Hidden Markov Model based feature engineering

The state of the art feature engineering techniques for credit card fraud detection creates descriptive features using the history of the card-holder (such as: "amount spent by the card-holder in shops from a given country in the last 24h" (Whitrow et al. (2008), Bahnsen et al. (2016))). These descriptive features present several limits we aim to overcome. First, they do not take into account the history of the seller even if it is clearly identified in most credit card transactions dataset. Moreover these descriptive features do not consider dependencies between transactions of a same sequence. Therefore we use Hidden Markov Models which are generative probabilistic models and a common choice for sequence modelling Rabiner and Juang (1991). Finally, the choice of the descriptive feature created using the transaction aggregation strategy (Whitrow et al. (2008), Bahnsen et al. (2016))) is guided by expert knowledge. In order not to depend on expert knowledge, we favor automated feature engineering in a supervised context.

In addition to the descriptive aggregated features created by Whitrow et al. (2008), we propose to create eight new HMM-based features. They quantify the similarity between the history of a transaction and eight distributions learned previously on set of sequences selected in a supervised way in order to model different perspectives.

We model the sequence of transactions from the combinations of three binary perspectives (genuine / fraudulent, card-holder / merchant, amount / timing) and therefore learn eight (2^3) different HMMs. At the end the set of 8 HMM-based features will provide information about the genuineness and the fraudulence of both terminal and card holder histories.

In particular, we select three perspectives for modelling a sequence of transactions (see figure 1). A sequence (i) can be made only of genuine historical transactions or can include at least one fraudulent transaction in the history, (ii) can come from a fixed card-holder or from a fixed terminal, and (iii) can consist of amount values or of time-delta values (i.e. the difference in time between the current transaction and the previous one). We optimized the parameters of eight HMMs using all eight possible combinations (i-iii).

To learn the HMM parameters on observed data, we create 4 datasets:

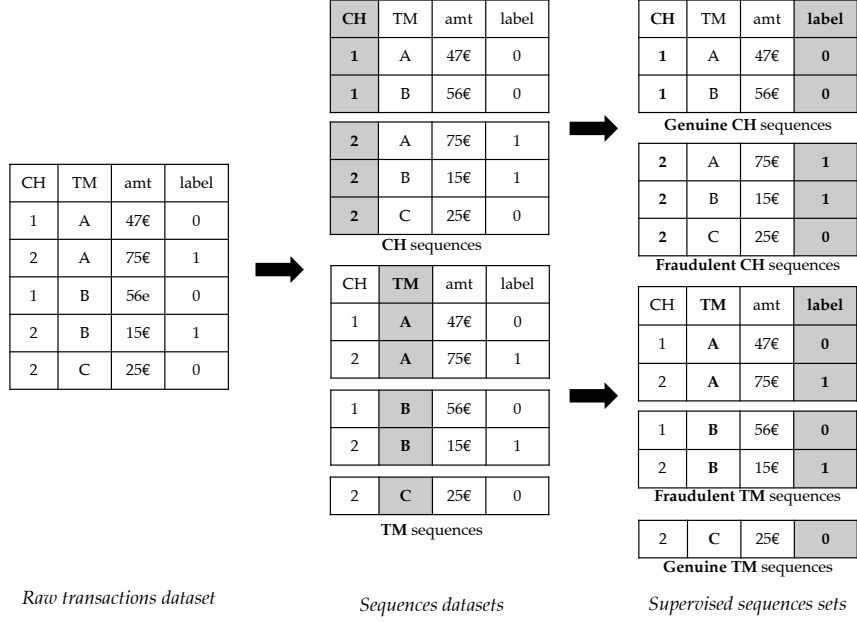


Figure 1: Supervised selection of sequences for the training sets of the multiple perspectives Hidden Markov Models ($CH=Card-holder$, $TM=Terminal$)

- sequences of transactions from **genuine credit cards** (without fraudulent transactions in their history).
- sequences of transactions from **compromised credit cards** (with at least one fraudulent transaction)
- sequences of transactions from **genuine terminals** (without fraudulent transactions in their history)
- sequences of transactions from **compromised terminals** (with at least one fraudulent transaction)

We then extract from these sequences of transactions the symbols that will be the observed variable for the HMMs. In our experiments, the observed variable can be either:

- the amount of a transaction.
- the amount of time elapsed between two consecutive transactions of a card-holder (time-delta).

Hidden Markov Models main hypothesis is that behind the observed distribution, there is a simpler latent model that rules the sequential distribution. Hypothetically, the complexity of the observed distribution comes partly from additive Gaussian noise corrupting both the state evolution process and the emission process. Hidden Markov Models allow for a compression of the observed sequence in order to generalize the observed behaviour into an abstracted latent behaviour. It is comprised of two processes represented by matrices (see figure 2):

- The *transition matrix* describes the evolution of the hidden states. Each row i of the transition matrix is a multinomial distribution of the next state given that the current state is i . The hidden states obey to the Markov property (i.e. given the present the future does not depend on the past).
- The *emission matrix* describes the conditional distribution of the observed variables given the current hidden state. Usually, the distribution is considered multinomial for categorical observed variables or gaussian for continuous observed variables.

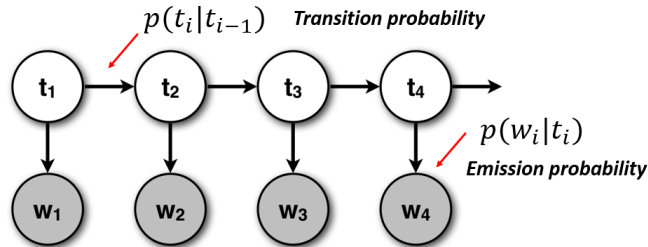


Figure 2: Hidden markov model architecture.

The transition and emission conditional probability matrices of the HMMs are optimized by an iterative Expectation-Maximisation algorithm known as the Baum-Welch algorithm (Baum (1972), Rabiner and Juang (1991)). EM optimization of a model with latent (hidden) parameters consists in (for parameters initialized with a value):

Expectation: Find the latent states distributions that correspond the most to the sequences of observed data. This is usually done with the help of the Viterbi algorithm which recursively leverage the Markov property in order to simplify the calculations of the conditional probabilities to observe a sequence of event given the parameters of the transition

and emission matrices (also referred to as forward-backward algorithm Viterbi (1967)).

Maximisation: Maximise the correspondence between the latent distribution inferred during the expectation step and the parameters of the transition and emission matrices by adjusting the parameters.

The expectation and maximization steps are repeated until convergence of the graphical model to the observed data. The convergence can be monitored by observing the increase of the value of the likelihood that the set of observed sequences has been generated by the model. This likelihood increases over the iterations until it reaches a ceiling when the hyper parameters ruling the architecture of the generative model don't allow it to fit more to the set of observed sequences.

At the end, we obtain 8 trained HMMs modeling 4 types of behaviour (genuine terminal behaviour, fraudulent terminal behaviour, genuine card-holder behaviour and fraudulent card-holder behaviour) for both observed variables (amount and time-delta).

Algorithm 1 Online: calculate likelihood of sequences of observed events

```

for  $tx_i$  in all transactions do
  for  $perspective_j$  in perspectives combinations do
     $[tx_i, tx_{i-1}, tx_{i-2}] \leftarrow usersequence_{i,j}$ 
     $HMM_j \leftarrow HMM(usertype, signaltype, sequencetype)$ 
     $AnomalyScore \leftarrow \log(P([tx_i, tx_{i-1}, tx_{i-2}] | HMM_j))$ 
  end for
end for

```

The HMM-based features proposed in this paper (table 1) are the likelihoods that a sequence made of the current transaction and the two previous ones from this terminal/card holder is generated by each of these models. In order to calculate their value, the most probable sequence of hidden states for each observed sequence has to be computed. This is usually done with the help of the Viterbi algorithm Viterbi (1967) which also leverages the Markov property in order to simplify the calculations of the conditional probabilities to observe a sequence of hidden states given the parameters of the HMM (initial probabilities and transition and emission matrices) and an observed sequence of events.

User	Feature	Genuine	Fraudulent
Card Holder	Amount	<i>HMM1</i>	<i>HMM5</i>
	Tdelta	<i>HMM2</i>	<i>HMM6</i>
Terminal	Amount	<i>HMM3</i>	<i>HMM7</i>
	Tdelta	<i>HMM4</i>	<i>HMM8</i>

Table 1: Set of 8 HMM-based features describing 8 combinations of perspectives

3 Experimental Setup

3.1 Dataset Description

We used a credit card transactions dataset provided by our industrial partner in order to quantify the increase in detection when adding HMM-based features. This dataset contains the anonymized transactions from all the belgian credit cards between 01.03.2015 and 31.05.2015.

The classification task is to predict the class of the transactions (genuine or fraudulent).

Transactions are represented by vectors of continuous, categorical and binary features that characterize the card-holder, the transaction and the terminal. The card-holder is characterized by a unique card-holder ID, its age, its gender, etc. The transaction is characterized by variables like the date-time, the amount and other confidential features. The terminal is characterized by a unique terminal-ID, a merchant category code, a country.

In addition to the work already presented as a poster Lucas et al. (2019), we studied the impact of the proposed multiple perspective HMM-based feature engineering strategy for face-to-face transactions which present very different properties than e-commerce transaction: The merchant is usually closed at night, the necessity of a PIN authentication decreases significantly the number of fraudulent transactions Ali et al. (2019), etc. For comparison, Jurgovsky et al. (2018) have shown on the same belgium transactions dataset that some machine learning approaches (LSTM) gave better results on face-to-face transactions than on e-commerce transactions.

3.2 Feature Engineering and Dataset Partitioning

In order for the HMM-based features and the aggregated features to be comparable, we calculate terminal-centered aggregated features in addition to (Whitrow et al., 2008) card-holder centered aggregated features (table 2)

Feature	Signification
AGGCH1	# transactions issued by user in 24h.
AGGCH2	Amount spent by user in 24h.
AGGCH3	# transactions in the country in 24h.
AGGCH4	Amount spent in the country in 24h.
AGGTM1	# transactions in terminal in 24h.
AGGTM2	Amount spent in terminal in 24h.
AGGTM3	# transactions with this card type in 24h.
AGGTM4	Amount spent with this card type in 24h.

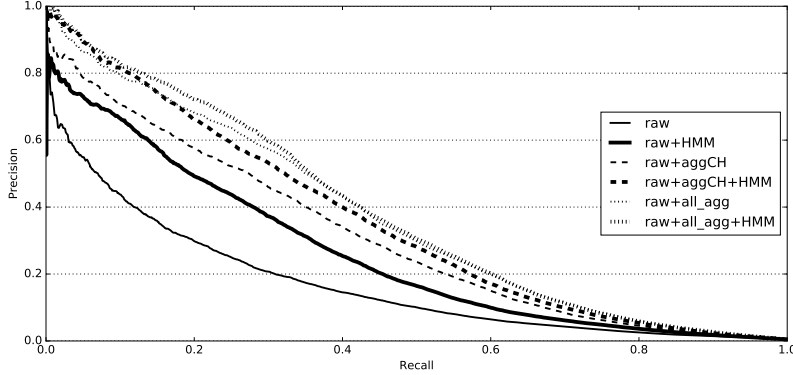
Table 2: Aggregated features centered on the card holders and the terminal

We split temporally the dataset in three different parts: the training set, the validation set and the testing set. We chose to separate the time period corresponding to the training and validation set and the time period corresponding to the testing set with a gap of 7 days. The reason is that for real world fraud detection systems, human investigators have to verify the alerts generated by the classifiers. Since this process takes time, the ground truth is delayed by about one week. The transactions appearing during this gap of 7 days before the testing set are used in order to calculate the value of the aggregated and HMM-based features in the testing set but not for training the classifiers. Therefore, the training set contains all the transactions between 01.03.2015 to 26.04.2015, the validation set used for the tuning of the random-forests contains all the transactions between 27.04.2015 to 30.04.2015 and the testing set contains all the transactions between 08.05.2015 to 31.05.2015.

4 Results

4.1 Improvement in fraud detection when using HMM-based features

We train Random Forest Classifiers using different feature sets in order to compare the efficiency of prediction when we add HMM-based features to the classification task for the face-to-face and e-commerce transactions. The difference in term of raw AUC between the e-commerce and the face-to-face transactions is due to the difference in imbalance between these datasets: the imbalance is 17 times stronger for the face-to-face transactions than for the e-commerce transactions. For face-to-face, there are 0.2 frauds per 1000 transactions whereas for e-commerce there are 3.7 frauds per 1000



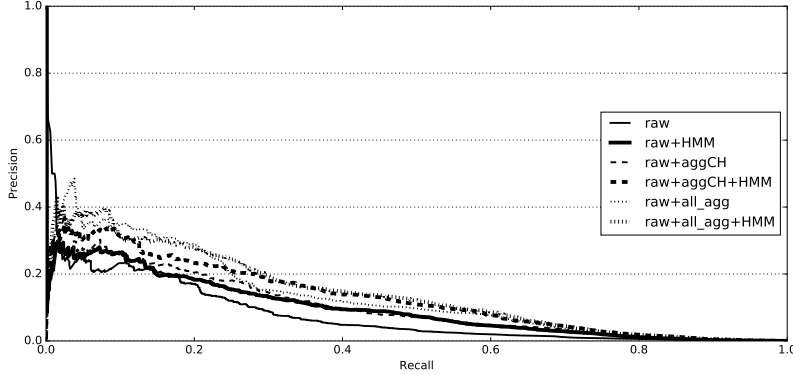
Feature set	no HMM-features	HMM-features	increase through HMMs
raw	0.212 ± 0.009	0.298 ± 0.0006	40.5%
raw+aggCH	0.343 ± 0.001	0.375 ± 0.0006	9.3%
raw+all agg	0.383 ± 0.002	0.397 ± 0.003	3.6%

Figure 3: Precision-recall curves for e-commerce transactions. (Each color corresponds to a specific feature set, the line style corresponds to the presence or not of HMM-based features. The addition of HMM-based features (bold lines) to each feature set considered - even the most informative ones - allows for an increase in the detection of fraudulent transactions when compared to the same prediction without HMM-based features (thin lines).)

transactions.

We tested the addition of our HMM-based features to several feature sets. We refer to the feature set **raw+aggCH** as the state of the art feature engineering strategy since it contains all the raw features with the addition of Whitrows aggregated features Whitrow et al. (2008). The feature groups we refer to are: the raw features (raw), the features based on the aggregations of card-holders transactions (aggCH), the features based on the aggregation of terminal transactions (aggTM), the proposed HMM-based features (HMM features).

In this section, the HMMs were created with 5 hidden states and the HMM-based features were calculated with a window-size of 3 (actual transaction + 2 past transactions of the card-holder and of the terminal). We showed in section 4.3 that the HMM hyperparameters (number of hidden states and size of the window considered for the calculation of HMM-based features) did not change significantly the increase in Precision-Recall AUC



Feature set	no HMM features	HMM features	increase through HMMs
raw	0.082 ± 0.001	0.152 ± 0.001	85.4%
raw+aggCH	0.144 ± 0.001	0.170 ± 0.001	18.1%
raw+all agg	0.167 ± 0.001	0.177 ± 0.0006	6.0%

Figure 4: Precision-recall curves for face-to-face transactions.

observed.

We tuned the Random Forest hyperparameters (table 3) through a grid search that optimizes the Precision-Recall Area under the Curve on the validation set. The choice of the Precision-Recall AUC for imbalanced dataset was motivated by the work of Davis and Goadrich (2006).

n-trees	n-features	min-samples-leaf	max depth
{300}	{1, 7, 13}	{1, 20, 40}	{4, <i>None</i> }

Table 3: Random Forest grid search

Figures 3 and 4 show the precision recall curves and their AUC obtained by testing the efficiency of Random Forests trained with several feature sets on the transactions of the testing set. The AUC numbers correspond to the average \pm standard deviation on 3 different runs. The AUCs are stable over the different runs and the standard deviation numbers are low.

We can observe that the face-to-face results (figure 4) and the e-commerce results (figure 3) both show a significant improvement in precision-recall AUC when adding sequence descriptors such as the proposed HMM-based features or Whitrows aggregated features to the raw feature set. For the

e-commerce transactions, this improvement ranges from 3.6% for the best feature set without HMM-based features (raw+all-agg for the e-commerce transactions) to 40.5% for the worst feature set (raw for the e-commerce transactions). For the face-to-face transactions, this improvement ranges from 6.0% for the best feature set without HMM-based features (raw+all-agg for the face-to-face transactions) to 85.4% for the worst feature set (raw for the face-to-face transactions).

By comparing the AUC of the curves raw+aggCH and raw+aggCH+HMM, we observe that adding HMM-based features to the state of the art feature engineering strategy introduced in the work of Whitrow et al. (2008) leads to an increase of 18.1% of the PR-AUC for the face-to-face transactions and to an increase of 9.3% of the PR-AUC for the e-commerce transactions.

The relative increase in PR-AUC when adding terminal centered aggregated features to the feature set is of 16.0% for the face-to-face dataset. The relative increase in PR-AUC when adding terminal centered aggregated features to the feature set is of 11.7% for the e-commerce transactions.

Overall, we can observe that the addition of features that describe the sequence of transactions, be it HMM-based features or Whitrows aggregated features, increases a lot the Precision-Recall AUC on both e-commerce and face-to-face transactions. The addition of HMM-based features improves the prediction on both e-commerce and face-to-face transactions and allows the classifiers to reach the best levels of accuracy on both e-commerce and face-to-face transactions.

4.2 Robustness over different machine learning algorithms

We have shown in section 4.1 that the addition of HMM-based features to the existing transaction aggregation strategy allows for a consistent and significant increase in the precision-recall AUC of random forest classifiers.

In order to make sure that this increase is stable over different types of classifiers, we did the same experiments with adaboost and logistic regression classifier.

We tuned the gradient boosting and logistic regression hyperparameters (table 5) through a grid search that optimizes the Precision-Recall Area under the Curve on the validation set.

By comparing the AUCs reported on table 5 and 6, we can conclude that the improvement observed when integrating the proposed multiple perspectives HMM-based feature engineering in addition to the state of the art transaction aggregation strategy is significant and reliable over different classifiers and datasets. The relative improvement is bigger for weaker clas-

tree numbers	learning rate	tolerance for stopping	max tree depth
$\{100, 400\}$	$\{0.1, 1, 100\}$	$\{10, 100\}$	$\{1, 4\}$

Table 4: Adaboost grid search

C parameter	penalty	tolerance for stopping
$\{1, 10, 100\}$	$\{l1, l2\}$	$\{10, 100\}$

Table 5: Logistic regression grid search

E-commerce	no HMM-features	HMM-features	increase through HMMs
raw	0.015 ± 0.0001	0.049 ± 0.001	267%
raw+aggCH	0.092 ± 0.005	0.121 ± 0.003	31.5%
raw+all agg	0.096 ± 0.004	0.124 ± 0.005	29.2%
Face-to-face	no HMM-features	HMM-features	increase through HMMs
raw	0.0094 ± 0.0001	0.030 ± 0.002	219%
raw+aggCH	0.014 ± 0.001	0.0247 ± 0.02	76.4%
raw+all agg	0.0161 ± 0.0003	0.0315 ± 0.0007	95.6%

Figure 5: Precision-recall AUC for logistic regression classifiers.

sifiers such as logistic regression classifiers than for strong credit card fraud detection classifiers such as random forest classifiers.

We observe a significant decrease in detection efficiency with logistic regression classifier compared to random forest classifier. Logistic regression classifier is known to be weak to categorical features encoded with a label encoder (category transformed to numbers) since the label encoding implies an order between the categories that is irrelevant. The results obtained with logistic regression classifiers may be improved using one-hot encoding or frequency encoding like advised in some other credit card fraud detection project (Pozzolo et al. (2015), Jurgovsky et al. (2018)), however the evolution observed when including the proposed feature engineering strategy to logistic regression classifier is consistent with what we observed using random forest classifier.

We observed that shallow adaboost classifiers with a tree depth restricted to 1 as advised in the literature Freund and Schapire (1997) lead to significantly worse detection of fraudulent transactions than deeper adaboost classifiers using trees of depth 4. The downside of deeper trees is the increased risk of overfitting of the adaboost classifiers. Monitoring the number of iterations (tree added) after which the classifier starts to overfit could be done by tracking the PR-AUC on the validation set and on the training set.

E-commerce	no HMM-features	HMM-features	increase through HMMs
raw	0.103 ± 0.001	0.163 ± 0.007	58.3%
raw+aggCH	0.238 ± 0.009	0.247 ± 0.011	3.8%
raw+all agg	0.263 ± 0.004	0.264 ± 0.006	0.4%
Face-to-face	no HMM-features	HMM-features	increase through HMMs
raw	0.034 ± 0.003	0.058 ± 0.009	70.6%
raw+aggCH	0.048 ± 0.002	0.079 ± 0.006	64.6%
raw+all agg	0.053 ± 0.003	0.061 ± 0.004	15.1%

Figure 6: Precision-recall AUC for Adaboost classifiers.

Due to adaboost design, the PR-AUC on the training set should increase for each tree added, however when the classifier overfits, the PR-AUC on the validation set would decrease: the classifier will start to learn particularities of the training set that are not relevant in other sets.

4.3 Robustness to hyperparameters changes

In order to understand if the feature engineering strategy is sensitive to the hyperparameters used for the construction of the HMM-based features, we constructed 9 sets of HMM-based features with different combinations of the HMM-based features hyperparameters. The hyperparameters considered are the number of hidden states in the HMMs and the size of the sequence of past transactions used for the calculation of the likelihood score.

We measure the AUC obtained on the test set when adding different set of HMM-based features obtained with different combinations of hyperparameters to the raw feature set.

		Window size		
		3	5	7
Hidden states	3	0.280 ± 0.013	0.292 ± 0.015	0.295 ± 0.019
	5	0.315 ± 0.008	0.310 ± 0.006	0.297 ± 0.007
	7	0.307 ± 0.003	0.307 ± 0.004	0.296 ± 0.011

Table 6: PR-AUCs for E-commerce HMM hyperparameters ($raw = 0.203 \pm 0.005$)

We observe that, the combination of parameter {window size: 3, hidden states: 5} gives the best AUCs on average on 3 runs. However the standard deviation values (\pm) are too high to confidently say that a hyperparameters

		Window size		
		3	5	7
Hidden states	3	0.159 ± 0.020	0.135 ± 0.007	0.150 ± 0.013
	5	0.139 ± 0.006	0.135 ± 0.012	0.123 ± 0.009
	7	0.129 ± 0.008	0.131 ± 0.009	0.124 ± 0.007

Table 7: PR-AUCs for Face-to-face HMM hyperparameters ($raw = 0.089 \pm 0.011$)

combination is better than all the others.

A window size of 3 and 5 hidden states was the combination of hyperparameter chosen for the construction of the HMM-based features of the section 4.1.

In a real world case, doing the hyperparameter search can be worth it but requires to calculate several set of HMM-based features. Overall we can conclude that the HMM-based feature engineering strategy is relatively stable over the hyperparameters.

We also observed that adding all 9 newly calculated HMM-based features sets provided a slight improvement on the face-to-face transactions (13.3%). We thought that adding HMM-based features with different hyperparameters could possibly bring different type of information. For example the window-size parameter could be seen as: description of short-term/mid-term/long-term history of the transaction. This small improvement wasn't observed on the e-commerce transactions (8.5% decrease).

5 Handling the missing value limitations

	e-commerce		face-to-face	
	# transactions	# frauds	# transactions	# frauds
<i>All transactions</i>	$5.4 * 10^6$	18472	$6.8 * 10^6$	2764
<i>History</i> ≥ 3	$4.5 * 10^6$	15650	$5.3 * 10^6$	1305
<i>History</i> ≥ 7	$3.4 * 10^6$	12493	$4.1 * 10^6$	859

Table 8: Number of transactions in the test set with different history constraints (*History* ≥ 3 means that all the transactions have at least 2 transactions in the past for the card-holder **and** for the terminal: we can build sequences of 3 transactions for both perspectives)

With the HMM feature engineering framework, we can calculate sets of

HMM features for different window sizes. However, when the transaction history is not big enough we can't calculate the HMM-based features for it. Because of this limitation we had to dismiss about 20% of the transactions for the experiments described in section 4.1 and around 40% of the transactions for the experiments of the hyperparameter section 4.3. There is a strong need to tackle the issue of structural missing values depending on the length of users sequences in order to integrate transactions with short history where part or all of the HMM-based features couldn't have been calculated.

In this section we consider 9 sets of HMM based features obtained with a window size of 3, 5 and 7 for the card-holder or for the terminal. We therefore have 16 sets of transactions with different history constraints (terminal history: $[0, 3, 5, 7]$ * card-holder history: $[0, 3, 5, 7]$).

We consider 3 missing values strategies:

default0 a genuine default value solution where the HMM-based features that couldn't have been calculated for the current transactions are replaced by 0.

weighted PR a weighted sum of the predictions of the Random Forests specialized on the history constraints of the transaction, among 16 Random Forests each trained on one of the 16 possible history constraints. For example, for a transaction with a terminal history of 6 and a card-holder history of 3 we will sum the predictions of the Random Forests $[0,0]$, $[0, 3]$, $[3, 0]$, $[3, 3]$, $[5, 0]$, $[5, 3]$ ¹. Each Random Forest is weighted by its efficiency on the validation set. We used PR AUCs values as weights for the Random Forests.

stacked RF a stacking approach where a Random Forest classifier is trained on the predictions of the 16 Random Forests specialized on the constraints (0 for missing prediction, when the transaction doesn't satisfy the constraints of the considered Random Forest). This approach has the second benefit that we might stack individual classifiers thereby creating a more accurate new one.

Other approaches to handle missing values advise to generate them by modelling the distribution of the corresponding features. We considered that these solutions don't apply in our case since the missing values appear because there wasn't enough historical information to calculate the corresponding feature. We thought that replacing the value of a model based feature

¹The first number describes the terminal history constraint, the second number describes the card-holder history constraint

(HMM based features) using an other model (kNN or Random Forest) wasn't an appropriate solution.

	e-commerce PR-AUC	face-to-face PR-AUC
<i>raw (no HMM-based features)</i>	0.264 ± 0.002	0.058 ± 0.012
<i>default0</i>	0.320 ± 0.001	0.145 ± 0.004
<i>weighted PR</i>	0.275 ± 0	0.040 ± 0.002
<i>stacked RF</i>	0.317 ± 0.002	0.133 ± 0.021

Table 9: Fraud detection on the whole test set with different missing values strategies

We can observe in table 9 that the simplest solution (*default0*) allows for the best PR-AUCs with the best stability for both face-to-face and e-commerce transactions. This is also by far the fastest method since it needs to only train one Random Forest instead of 16 (resp. 17) for the *weighted PR* approach (resp. *stacked RF* approach).

The *weighted PR* solution doesn't allow for satisfying results. However the *stacked RF* solution gives good PR-AUCs and presents interesting properties in order to combine different types of classifiers.

Finding satisfying solutions to integrate transactions with structural missing value increases drastically the range of application of the proposed framework (see table 8). Moreover it adds an other perspective to the framework: HMM-based features calculated for a small (resp. big) window-size will characterize the short (resp. long) term history.

Conclusion

In this work, we propose an HMM-based feature engineering strategy that allows us to incorporate sequential knowledge in the transactions in the form of HMM-based features. These HMM-based features enable a non sequential classifier (Random Forest) to use sequential information for the classification.

The multiple perspective property of our HMM-based automated feature engineering strategy gives us the possibility to incorporate a broad spectrum of sequential information. In fact, we model the genuine and fraudulent behaviours of the merchants and the card-holders according to two features: the timing and the amount of the transactions. Moreover, the HMM-based features are created in a supervised way and therefore lower the need of ex-

pert knowledge for the creation of the fraud detection system. The terminal perspective is usually not used in credit card fraud detection and is shown in this paper to greatly help the detection for face-to-face and e-commerce transactions.

This extension to Lucas et al. (2019) consolidates the claims already made with additional experiments and evaluations. More precisely:

- The feature engineering strategy is shown to perform well for e-commerce and face-to-face credit card fraud detection: the results show an increase in the precision-recall AUC of 18.1% for the face-to-face transactions and 9.3% for the e-commerce ones.
- Then, the feature engineering strategy is shown to be relevant for various types of classifiers (random forest, logistic regression and Adaboost) and robust to hyperparameters choices made for constructing the features.
- Finally, the structural missing values limitation of the framework is looked at and several solutions are benchmarked.

HMM-based feature engineering strategy is a powerful tool that is shown to present interesting properties for fraud detection. We can imagine building similar HMM-based features in any supervised task that involve a sequential dataset.

To ensure reproducibility, a source code for calculating and evaluating the proposed HMM-based features can be found at https://gitlab.com/Yvan_Lucas/hmm-ccfd.

As a future work, it would be interesting to combine the predictions of an LSTM with the prediction of some HMM-based features enhanced Random Forest since these classifiers have been shown to not detect the same frauds in face-to-face transactions by (Jurgovsky et al., 2018).

5.1 Acknowledgement:

The work has been funded partially by the Bavarian Ministry of Economic Affairs, Regional Development and Energy in the project Internetkompetenzzentrum Ostbayern”.

References

- Ali, M.A., Azad, M.A., Centeno, M.P., Hao, F., van Moorsel, A., 2019. Consumer-facing technology fraud: Economics, attack methods and potential solutions. *Future Generation Computer Systems* (volume 100) .
- Bahnsen, A.C., Aouada, D., Stojanovic, A., Ottersten, B., 2016. Feature engineering strategies for credit card fraud detection. *Expert Systems With Applications* .
- Bahnsen, A.C., Stojanovic, A., Aouada, D., Ottersten, B., 2013. Cost sensitive credit card fraud detection using bayes minimum risk. *Proceedings of the 2013 12th international conference on machine learning and applications* .
- Baum, L.E., 1972. An inequality and associated maximization technique. *Statistical Estimation for Probabilistic Functions of Markov Processes* .
- Bhattacharyya, S., Jha, S., Tharakunnel, K., Westland, J.C., 2011. Data mining for credit card fraud: A comparative study. *Decision Support Systems*, 50(3) .
- Bolton, R., Hand, D.J., 2001. Unsupervised profiling methods for fraud detection. *Credit scoring and credit control VII* .
- Breiman, L., Friedman, J., Olshen, R.A., Stone, C.J., 1984. Classification and regression trees (p. 147). *The Wadsworth statistic / probability series* .
- Chandola, V., Banerjee, A., Kumar, V., 2012. Anomaly detection for discrete sequences: A survey. *IEEE Transactions on Knowledge and Data Engineering* .
- Davis, J., Goadrich, M., 2006. The relationship between precision-recall and roc curves. *ICML '06 Proceedings of the 23rd international conference on Machine learning* .
- Dietterich, T., 2002. Machine learning for sequential data: A review. *Structural, syntactic, and statistical pattern recognition* .
- Donato, J.M., Schryver, J.C., Hinkel, G.C., Schmoyer, R.L., Leuze, M.R., Grandy, N.W., 1999. Mining multi-dimensional data for decision support. *Future Generation Computer Systems* (volume 15) .

- Freund, Y., Schapire, R., 1997. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of computer and system sciences* .
- Graves, A., 2012. Supervised sequence labelling with recurrent neural networks. Springer vol. 385 .
- Jha, S., Guillen, M., Westland, J., 2012. Employing transaction aggregation strategy to detect credit card fraud. *Expert Systems with Applications*, 39(16) .
- Jurgovsky, J., Granitzer, M., Ziegler, K., Calabretto, S., Portier, P., He-Guelton, L., Caelen, O., 2018. Sequence classification for credit-card fraud detection. *Expert Systems with Applications* .
- Krivko, M., 2010. A hybrid model for plastic card fraud detection systems. *Expert Systems with Applications*, 37(8) .
- Lucas, Y., Portier, P.E., Laporte, L., Calabretto, S., Caelen, O., He-Guelton, L., Granitzer, M., 2019. Multiple perspectives hmm-based feature engineering for credit card fraud detection. 34th ACM/SIGAPP Symposium on Applied Computing (SAC2019) .
- Maes, S., Tuyls, K., Vanschoenwinkel, B., Manderick, B., 2002. Credit card fraud detection using bayesian and neural networks. *Proceedings of NF* .
- Mahmoudi, N., Duman, E., 2015. Detecting credit card fraud by modified fisher discriminant analysis. *Expert Systems with Applications*, 42(5) .
- Minegishi, T., Niimi, A., 2011. Proposal of credit card fraudulent use detection by online-type decision tree construction and verification of generality. *International Journal for Information Security Research (IJISR)*, 1(4) .
- Pozzolo, A.D., Boracchi, G., Caelen, O., Alippi, C., Bontempi, G., 2015. Credit card fraud detection and concept-drift adaptation with delayed supervised information. 2015 International Joint Conference on Neural Networks (IJCNN) .
- Pozzolo, A.D., Boracchi, G., Caelen, O., Alippi, C., Bontempi, G., 2017. Credit card fraud detection: a realistic modeling and a novel learning strategy. *IEEE transactions on neural networks and learning systems* .

- Pozzolo, A.D., Caelen, O., Borgne, Y.L., Waterschoot, S., Bontempi, G., 2014. Learned lessons in credit card fraud detection from a practitioner perspective. *Expert Systems with Applications* .
- Rabiner, L.R., Juang, B.H., 1991. Hidden markov models for speech recognition. *Technometrics* .
- Saia, R., Carta, S., 2019. Evaluating the benefits of using proactive transformed-domain-based techniques in fraud detection tasks. *Future Generation Computer Systems* (volume 93) .
- Strivastava, A., Kundu, A., Sural, S., Majundar, A.K., 2008. Credit card fraud detection using hidden markov model. *IEEE Transactions on dependable and secure computing* .
- Viterbi, A.J., 1967. Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE Transactions Information Theory*, vol. IT-13 .
- Vlasselaer, V.V., Bravo, C., Caelen, O., Eliassi-Rad, T., Akoglu, L., Snoeck, M., Baesens, B., 2015. Apate: A novel approach for automated credit card transaction fraud detection using network-based extensions. *Decision Support Systems* .
- Whitrow, C., Hand, D.J., Juszczak, P., Weston, D.J., Adams, N.M., 2008. Transaction aggregation strategy for credit card fraud detection. *Data Mining and Knowledge Discovery* 18(1) .
- Wiese, B., Omlin, C., 2009. Credit card transactions, fraud detection and machine learning: Modelling time with lstm recurrent neural networks. *Innovations in neural information paradigms and applications*, Springer .