DIGITAL PAYMENT CARD FRAUD: NEW VECTORS AND DETECTION

Aleksey MINTS
Pryazovskyi State Technical University
Universytetska str. 7, Mariupol 87555, Ukraine,
E-mail mints.alexey@gmail.com
ORCID ID: 0000-0002-8032-005X

Pavlo SIDELOV
Pryazovskyi State Technical University
Universytetska str. 7, Mariupol 87555, Ukraine,
E-mail pavlo@sidelov.com
ORCID ID: 0000-0001-5672-8189

Abstract: *This research explores gross losses from payment card fraud worldwide and considers the main modern methods of fraud. The issue of payment fraud will be exacerbated by the digitalization of economic relations, in particular the introduction by banks of the concept of "Bank-as-a-Service," which will increase the burden on payment services.*
The aim of this study is to show effective methods for detecting fraud in digital payment systems using automated machine learning and big data analysis algorithms.
Approaches to expanding the information base to detect fraudulent transactions are proposed and systematized. The choice of performance metrics for building and comparing models is substantiated.
The use of automatic machine learning algorithms is proposed to resolve the issue, which makes it possible in a short time to go through a large number of variants of models, their ensembles, and input data sets. As a result, our experiments allow us to obtain a quality of classification based on the AUC metric at a level that exceeds the effectiveness of the classifiers developed by traditional methods, even as the time spent on the synthesis of the models is much less and is measured in hours. The ensemble of models makes it possible to detect up to 85.7% of fraudulent transactions in the sample. The accuracy of fraud detection is also high.
The results of our study confirm the effectiveness of using automatic machine learning algorithms to synthesize fraud detection models in digital payment systems. In this case, efficiency is manifested not only by the resulting classifiers' quality but also by the reduction in the cost of their development, as well as by the high potential of interpretability. Implementing the study results could enable financial institutions to reduce the financial and temporal costs of developing and updating active systems against payment fraud, as well as improve the effectiveness of monitoring financial transactions.

Keywords: *payment card fraud, fraud prevention, PSD2, spotting fraudulent transactions, artificial intelligence, automatic machine learning, AutoML solutions.*

Introduction

Payment card fraud affects everyone. Almost $50 billion are lost yearly to card fraud (Nilson Report, 2020) and identity theft (Tedder & Buzzard, 2020) worldwide. Although financial institutions are locked in an escalating arms race against cybercriminals and scammers, losses still have to be accounted for. Consumers end up paying for money lost to fraud out of pocket, in the form of vendor and transaction fees, while corporations and governments spend billions more investigating and handling fraud cases.

The main body of the paper

Modern fraud prevention is expensive. Digital ID checks cost around $2 per document, companies spend millions on KYC (know your customer) and AML (anti money laundering), and still the number of fraudulent transactions is growing. Banks have been relying on passive measures to counteract fraud based on past breaches or fraud behavior history, and only some have invested in proactive or predictive fraud prevention.

Card-based payment systems worldwide generated gross fraud losses of $28.65 billion in 2019, amounting to 6.8¢ for every $100 of total volume (Nilson Report, 2020). To understand what financial institutions can do to improve their fraud prevention efforts, the current protection mechanisms need to be examined. Payment cards that hold a set of credentials or cardholder data act as keys to a customer's bank account and enable two types of transactions: card-present and card-not-present.

Payment card fraud basics

Card-present fraud is when a payment card is physically used to make purchases or withdraw money from ATMs by entering a PIN. For decades, scammers have been using cameras, sensors, ATM skimmers, and other devices to make copies of cards and extract PINs.

In one case, a waiter was discovered using a portable magnetic stripe reader in his shoe to copy customers' cards while walking to the register. In another scam, criminals used NFC readers to steal small amounts from people's cards on the subway. By coming close to pockets and bags, they were able to charge cards without people noticing.

Phone confirmations for larger amounts and RFID blocking wallets can partially counter card-present fraud. Card-not-present transactions are more complex as they happen remotely, where a cardholder does not present a card to a merchant in person. The CVV code on the back of the card is most often used to confirm that the person paying has physical access to the payment card, but 2FA methods via SMS OTP (one-time password) and in-app authentications are becoming more widespread.

Scammers can intercept OTPs, consumer sessions, cardholder data (PAN, EXP, NAME, CVV), and even steal app credentials. The Lazarus Group from North Korea is notorious for using military-grade cyber expertise to steal money using man-in-the-middle software and cloned credit cards to withdraw cash from ATMs – so much so that an estimated 75% to 80% of all ATM cash-out losses get repatriated to North Korea (Nilson Report, 2020).

Banking: Old vs. new

PSD2 – the revised European Payment Service Directive that covers the whole of the EU, brought into law in 2018 – aimed to fix the lack of an open banking regulatory environment, improve security, and protect customers, among other goals. Before banks started to adopt OpenAPI, companies found it extremely difficult to integrate with banks using ancient file exchange systems. PSD2 standardized how payment and financial institutions interact with each other and with third-party providers. The directive enabled AISPs (Account Information Service Providers) to access information from multiple financial institutions with a customer's permission. AISP services, for example, can aggregate data from different accounts in different banks and show it to a consumer in one place or application.

PISPs (Payment Initiation Service Providers) can go a step further and make payments on behalf of consumers. PISPs can pay incoming utilities, Internet, and service bills that a consumer receives automatically. Although AISPs and PISPs are still in the early stages of development and adoption,

similar initiatives are already being implemented worldwide.
In 5–10 years, OpenAPI initiatives will reach their potential and unlock the benefits of digital banking. Truly interactive banking experiences are a significant benefit for consumers, but these changes open the industry to completely new attack vectors that need to be accounted and prepared for.

Spotting fraudulent transactions using AI & ML
Global cataclysms of the 21st century quickly changed consumers' behavior, and this led to a crash in financial anti-fraud systems. Models used to predict consumer behavior, supply, and demand had to be retaught to account for new patterns and spikes.
Let us assume that, as a financial institution, a customer's payment card was compromised during the COVID-19 pandemic. What can be done to spot fraudulent transactions early on? One course of action is to take a data set, mark confirmed fraudulent transactions with a chargeback or other documented problem, and analyze it to determine correlations.
For most areas, obtaining a comprehensive data set is not a problem. However, privacy laws protect banking and transaction data from being disclosed. General Data Protection Regulation (GDPR) in the EU provides customers with the right to be forgotten, and Big Tech companies are already being sued for billions for breaching privacy laws (Brando, 2018). In terms of machine learning, if a consumer asks for their data to be deleted, does that request apply to the results of calculations based on their data? How far the law reaches will be discussed for years to come.

Raw data
As a result, there are very few data sets with real customer data in the public domain. A relatively large 150 MB data set from Kaggle – with hundreds of thousands of anonymized transactions from European credit card users recorded in 2013 (Machine Learning Group, 2013) – is used to research how to prevent payment card fraud in this research. Locating useful information in a raw data set is a very resource-intensive task that usually requires multiple data scientists and analysts.
A technology executive with a heavy managerial workload cannot spare weeks to clean, spot anomalies, and balance the data set. Under this scenario, a relatively new Automatic Machine Learning (AutoML) approach could take on all of the routine and repetitive tasks that come with in-depth data analysis and extract insights from raw data (Heller, 2019).
There are many AutoML solutions to choose from today. Giant AWS SageMaker, Google AutoML, AutoAI with IBM Watson Studio, Microsoft Azure ML, and Oracle AutoML are complemented by the smaller, but no less interesting, DataRobot, Auto Weka, AutoML-Freiburg-Hannover, and H2O Driverless AI. The latter of these solutions is preferable because it can run on a local server or even a laptop instead of relying on the cloud (H2O Driverless AI, 2021).
Whenever financial data is involved, most regulators restrict its movement to prevent data transfers outside the country or into the cloud. Another important point is that cloud-based solutions are often limited by the amount of processable data. Some products restrict tables to one million rows and add other restrictions to encourage users to purchase expensive enterprise-level licenses.
Even though H2O is a commercial project, there is also a free version. This does not have a helpful GUI, but that should not be a problem for skilled hands. As a result, H2O Driverless AI with an educational license was used for this research.
The data set itself was a CSV file with only a few readable variables: time, amount, and class – whether a transaction was fraudulent or not. The rest were anonymized to protect the privacy of consumers (Figure 1). This makes it more interesting as one can observe how the system behaves with many unknown variables.

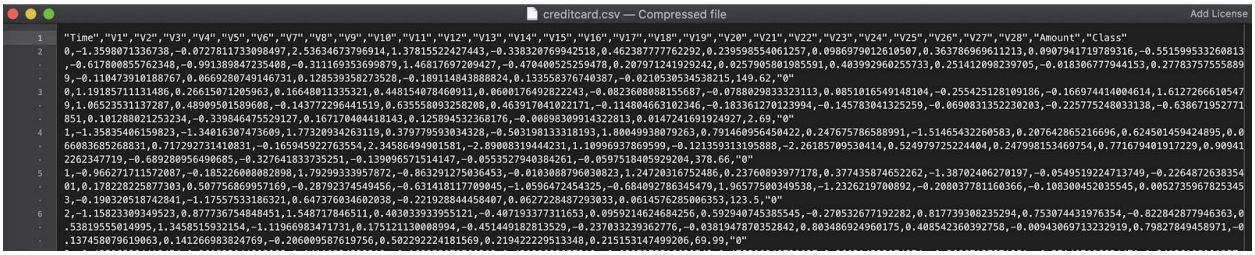

Figure 1. Raw credit card fraud detection data set

Source: Machine Learning Group (2013)

Data preparation

Raw data from the information systems of banks or payment organizations contain a technical minimum of information about transactions, which includes a relatively small set of parameters, including:

– date and time;

– the amount of the transaction (in the transaction currency and the currency of the account);

– the currency of the transaction;

– customer ID;

– the merchant/terminal ID through which the transaction was made;

– the type of identification procedure carried out;

– the authorization code, or the issuer's bank response if the transaction is rejected.

For a series of reasons that are detailed below, this set of parameters is not sufficient to effectively identify fraudulent transactions. This necessitates additional data preparation procedures. Thus, when using machine learning methods, additional information, which is usually contained in related relational databases, should be immediately added to the input sample. This is because machine learning methods are focused on analyzing data arranged in flat tables.

In order to effectively detect fraud, in addition to the information available in the databases of the bank or payment organization, it is necessary to add to the data external information (for example, a posteriori assessment of various risks associated with the transaction). The process of adding such information is termed feature augmentation. Its necessity is due to the fundamental limitations inherent in methods focused on line-by-line data analysis, which were formulated for the first time by Minsky and Papert (1969) for perceptron neural networks, but are also true for most other machine learning methods:

– the inability to generalize their characteristics to new stimuli or new situations;

– the inability to analyze complex situations in an external environment by decomposing them into simpler ones;

– the limitations in problems related to the invariant representation of images.

Expanding the feature augmentation by adding information that clarifies the current situation reduces the impact of the specified restrictions. This information can be of several kinds:

– additional information that is not contained in the transaction's original data, but expands knowledge about it (for example, associated risks);

– information that can be obtained as a result of vertical data analysis (for example, the average, maximum, minimum parameter values);

– information obtained from empirical analysis models, which are usually formulated as a condition or consequence (for example, signs of transactions subject to mandatory financial monitoring). In fact, such models can be considered micro-expert systems.

An analysis of research in the field of monitoring and detection of fraudulent transactions (Dal Pozzolo et al, 2013; Fu et al., 2016) has made it possible to formulate the following list of additional parameters for building a machine learning information base:
– the assessment of the risk of the terminal/merchant in which the transaction took place;
– the assessment of the merchant's risk in which the previous transaction was made;
– the assessment of the risk of the merchant category;
– the estimates of geographic risk (continental, country, regional);
– assessing the risk of the card issuer;
– assessments of other risks (related to the age and gender of the client, language group, place of the previous transaction, transaction amount, etc.);
– the total amount of customer transactions for the period under review;
– the minimum amount of a customer's transaction for the period under review;
– customer data (age, gender, etc.);
– additional transaction parameters (the use of special identification technologies, time, information about accepting, or rejecting the transaction, etc.).
Thus, the original data set can be expanded from 8 parameters to 25–30 (depending on the availability of information in the bank's databases and the available risk assessment capabilities).
In practice, building a representative information base for investigating fraudulent transactions involves solving the problem of finding a large enough sample of reliable evidence for analysis. The issue is due to the fact that employees of banks and payment organizations consider the transaction data of their customers as confidential, and provide them for research only to commercial developers, subject to contractual terms.
The only data set with real-world data containing parameters similar to the above is the Credit Card Fraud Detection data set (Machine Learning Group, 2013). This data set contains actual credit card payments from customers of Western European banks, and has the following parameters:
– the data set contains 284,807 transactions, including 492 cases of fraud identified post-factum (Figure 1). This represents only 0.172% of the total sample. Thus, the data sample is heavily unbalanced. This skew is typical of real data of digital payment systems and is due to the fact that the actual number of fraudulent transactions is relatively small;
– the sample contains 30 inputs and 1 output variable. Of these, only 3 variables (Time, Amount, and Class) contain name-appropriate transaction data that are not subject to additional conversions. The remaining variables are converted into dimensionless values for the preservation of confidentiality, as demanded by European law, and their titles are replaced with conditionals (V1, V2, ..., V28);
– there are no missed values in the data set;
– there are no cardholder identifiers, so all transactions can be considered independent of each other.
It should also be noted that, in accordance with the 2016 ruling of the European Union (GDPR), personal data can be provided only in a fully anonymized form, which does not allow their deanonymization. Therefore, the data set used is processed using the main component method; most of the variable names have been replaced with conditional ones. Despite some discomfort in interpreting the results, this can be interpreted as positive as part of the study's goal, as it reduces the subjectivity of assessments.

Setting up AutoML in H2O Driverless AI
For the experiments, the following configuration was used: IBM System X 3300 M Server with 12 Cores, 32 GB RAM, and Ubuntu Linux 18.04 LTS. This system is somewhat of an old workhorse,

without a GPU, but this provides a clearer picture of the performance of the process. After importing the data set into H2O, the system automatically analyzed the type and structure of data and suggested the best preliminary models, classificators, and analysis tools based on what was inside the data set. In this case, the data set was highly unbalanced, so H2O recommended the Log Loss scorer (Wikiwand, 2018).

Immediately after importing the data set, H2O quickly showed the problem and unbalanced areas. After confirming a wide variety of settings, the system began to analyze the data. The GUI showed preliminary results during the process, which could be explored and changed before full analysis was completed (Figure 2). Overall, it took around five days to process the data in this setup.

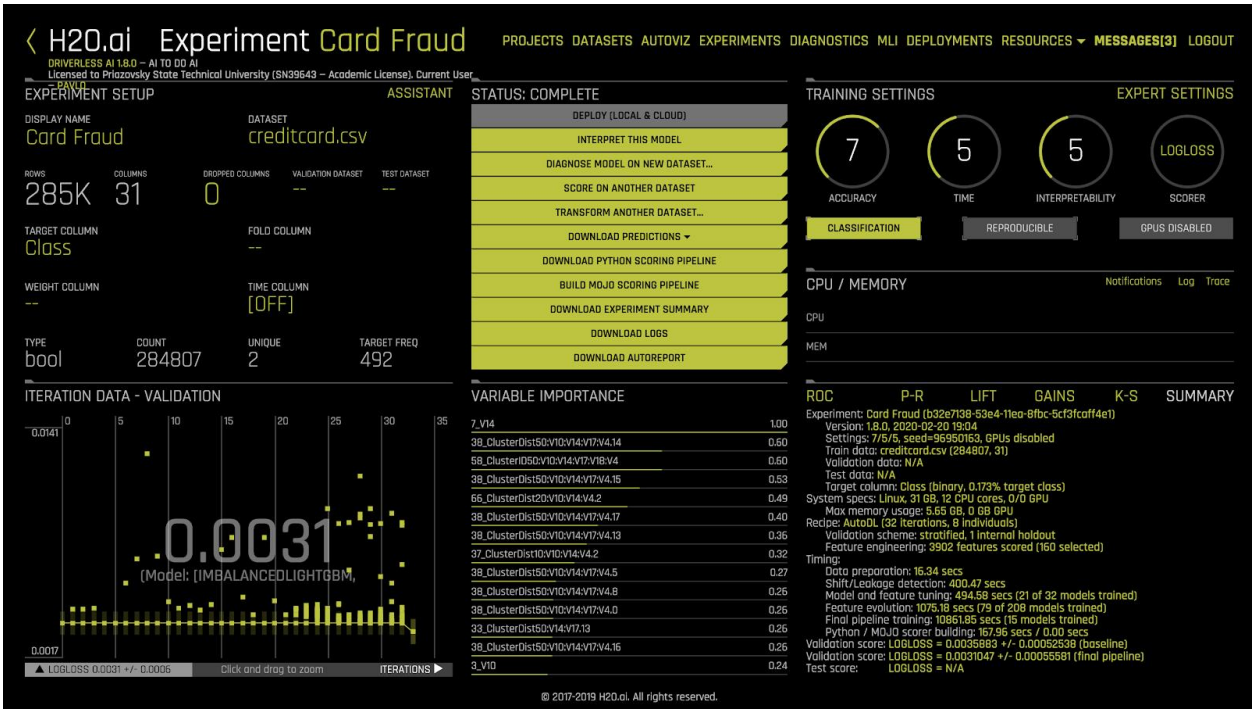

Figure 2. The results of the experiment in H2O Driverless AI
Source: Authors' own experiments

After completing the experiment, H2O offered a choice of models ready to be deployed on the cloud, servers, or datacenters. This enables almost seamless continuous delivery, or delivery after pressing a single button. Both options are very beneficial because updating such systems is a complex process that requires specialist skills.

Model interpretation

In this step, the selected system automatically processes the models that were built by themselves. In this experiment, another day and a half was required, and the system returned results that showed the influence, dependence, importance, and weight of different variables in the data set (Figure 3 and 4).

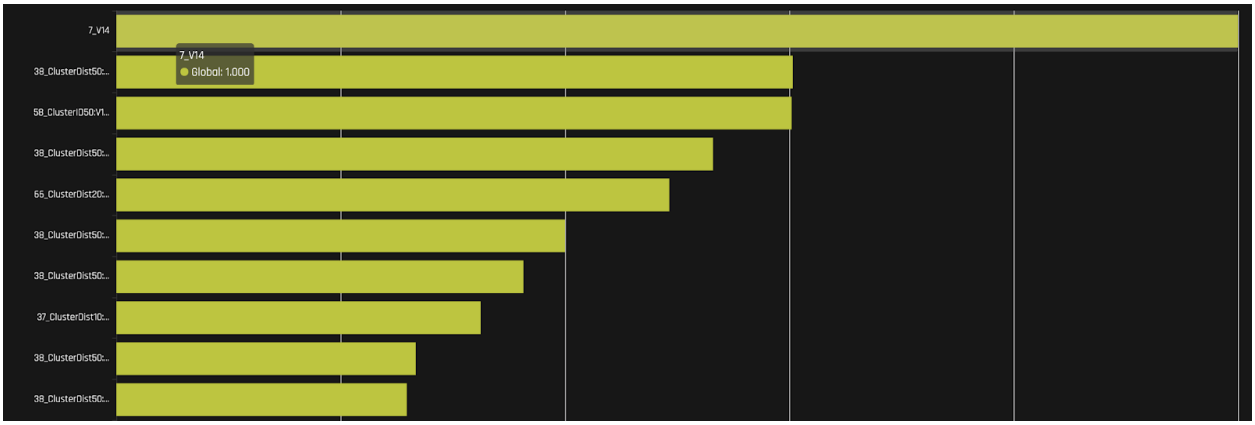

Figure 3. The results of model interpretation
Source: Authors' own experiments

Figure 3 demonstrates the importance of variable V14, which needs to be examined further. The rest of the results (Figure 4) consisted of other synthesized cluster functions. Using these results, it is possible to go through each function separately and analyze whether it is essential or not.

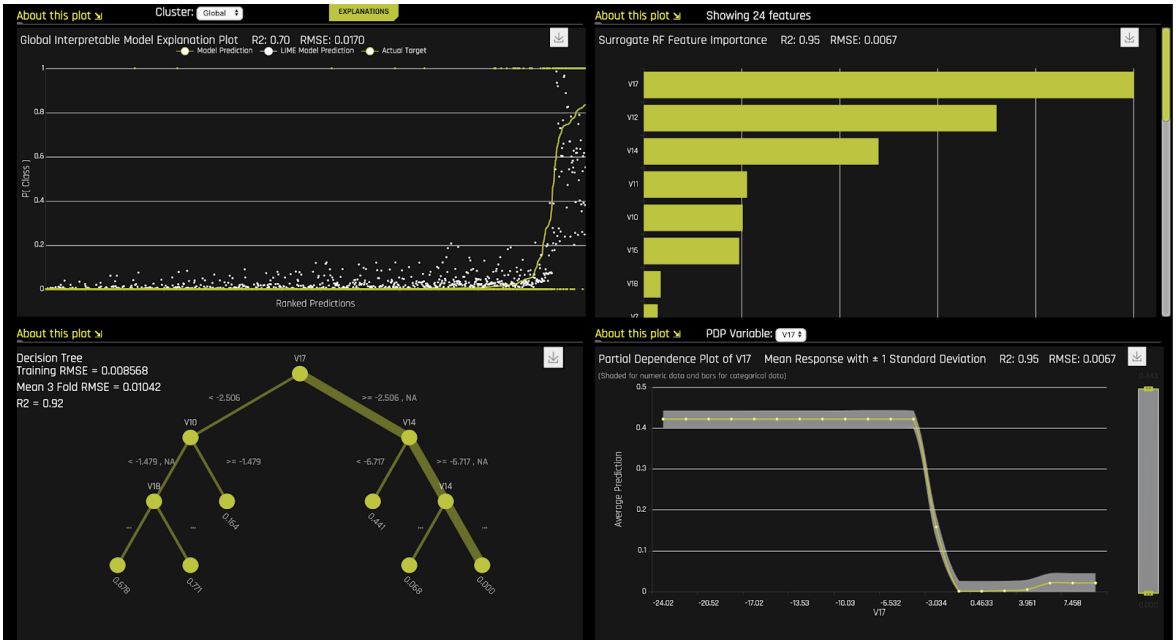

Figure 4. The results of model interpretation
Source: Authors' own experiments

Through these results, it can be observed how the system taught itself, made decisions, and approached maximum results – and where it was incorrect. H2O shows different patterns and possible interpretations of different values and how the system sees relationships between other variables and results in the data set. The analysis of peaked variables then showed a possible relationship between time, amount, and some merchant attributes.

There were two additional experiments where fields that came up in the first experiment were excluded to observe and ensure whether the result would remain the same without them. Overall, the new experiments were successful. Sometimes, there were differences in the variables' influence, and in other cases, H2O synthesized new functions.

AutoML results

We conducted several experiments on the automatic synthesis of the ensembles of models to identify fraudulent transactions in digital payment systems. The main differences between the conditions of the experiments were the use of different metrics for determining the quality of classification, as well as in the volumes of data analyzed. The methodology of these experiments is described by Kolodiziev et al. (2020).

The confusion matrix of the results of the first experiment is given in Table 1.

Table 1. Confusion matrix for Experiment 1

Source: Authors' own experiments

		Transaction genuine class	
		P	*N*
Predicted transaction class	*P*	388	67
	N	104	284,248

It follows from the analysis of Table 1 that the given experiment analyzed 284,807 transactions, of which 492 were fraudulent in reality. At the same time, the model correctly identified 388 fraudulent transactions (78.9% of the total amount). The remaining 104 fraud transactions (21.1%) were interpreted by the system as genuine. In addition, 67 genuine transactions were interpreted by the system as fraudulent.

The Receiver Operating Characteristic Curve chart and the appropriate Area Under Curve (AUC) value are shown in Figure 5.

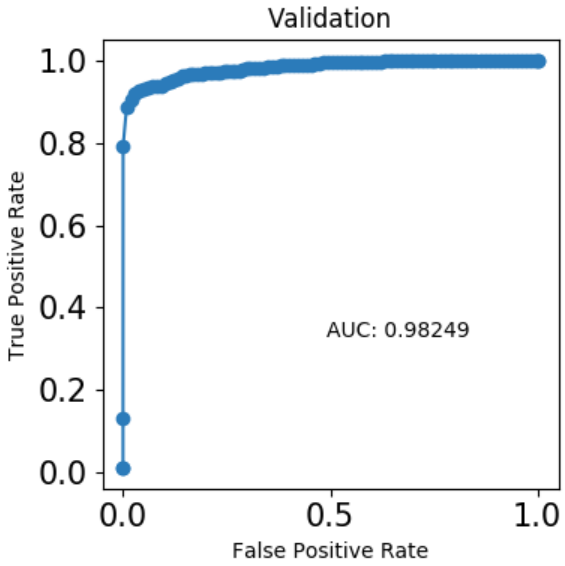

Figure 5. ROC curve for Experiment 1

Source: Authors' own experiments

The value of AUC = 0.98249 is good in itself, but when analyzing the unbalanced data set the final

decision on the choice of the classifier can be made only after the AUC values are compared to all classifiers.
The main differences of Experiment 2 from Experiment 1 were:
– accuracy was the main metric for the quality of classification in the synthesis of models;
– to reduce the calculation time, the input sample of data was reduced to 100,000 lines.
The confusion matrix of the experiment's results is given in Table 2.

Table 2. Confusion matrix for Experiment 2
Source: Authors' own experiments

		Transaction genuine class	
		P	*N*
Predicted transaction class	*P*	191	46
	N	32	99,730

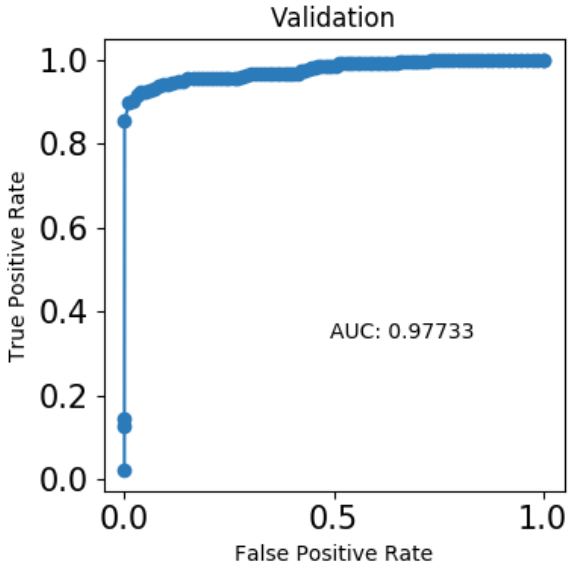

Figure 6. ROC curve for Experiment 2
Source: Authors' own experiments

The analysis of Figure 6 shows that the AUC metric for Experiment 2 (0.97733) was slightly lower than for Experiment 1 (0.98249), but this difference is negligible.
The comparative analysis of the effectiveness of the models' ensembles – synthesized during Experiment 2 and Experiment 1 by calculating the *Precision* and *Recall* characteristics (Kolodiziev et al., 2020) – produced the following results:

$$Precision_1 = \frac{388}{388 + 67} = 0.85,$$

$$Recall_1 = \frac{388}{388+104} = 0.79,$$

$$Precision_2 = \frac{191}{191+46} = 0.806,$$

$$Recall_2 = \frac{191}{191+32} = 0.857.$$

Thus, in the second experiment, we managed to achieve some better (by 8.5%) recall metric results, showing how many fraudulent transactions were able to be identified from the total number of transactions.

The precision metric, which shows how many transactions that received the fraud label were indeed fraudulent, was 5.1% worse for the second experiment than that for the first.

From the point of view of the bank or payment organization, the result for the recall metric is somewhat more important than the result for the precision metric. Thus, in practice, the ensemble of models obtained as a result of Experiment 2 is likely to be selected.

Another progressive feature of modern software is the ability to generate a Microsoft Word .doc file as a report with all of its findings and the lifecycle of the analysis, which can be printed out and read at any time. This shows everything that the system did, methods, how long it took, how effective it was, shifts, and importance. This saves at least a week of human analyst work.

The generated model can be turned into a java or python application that will generate a set of APIs and import a raw data set with transaction variables, and the system will show whether the transaction is fraudulent or not and how sure it is in that decision. This may be used to decide whether to allow the processing of a transaction or to stop it immediately.

These tools can be used by top-level management, CTOs, and even marketing departments to generate valuable insights into business operations and answer the following question: When is the best time to push a notification about a new product – is it when consumers make the most transactions, or the opposite?

AutoML can help find answers and improve business decision-making through data analysis.

This solution is not a magic bullet against all fraudulent transactions or the only correct method for rolling out such a model. This experiment has provided enough information regarding what is inside AutoML, what can be worked on, and what else can be explored. However, it is reassuring that a different research team that spent six months working on the same data set reached the same conclusions that this research determined in around a week.

New attack vectors

Payment card fraud is limited by card expiry dates, limits, and security notifications. The method explained above can help find and stop fraudulent transactions made by perpetrators, but what if customers unwittingly transfer money to criminals by themselves?

Currently, according to the European Central Bank, the majority of fraudulent transactions with payment cards are carried out without the presence of a card (Figure 7).

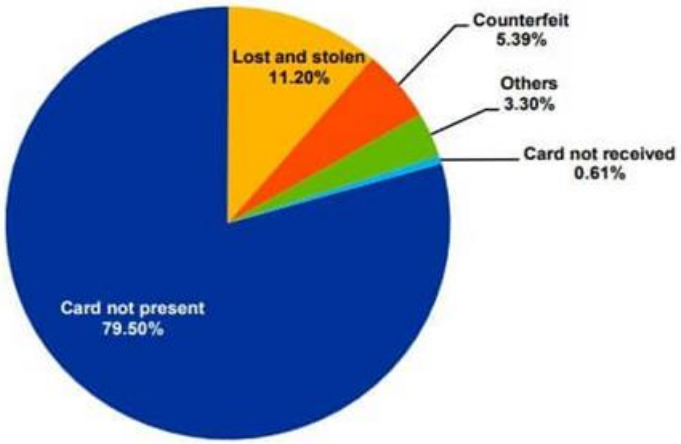

Figure 7. Value of fraud types as a share of total card fraud using cards issued within SEPA
Source: European Central Bank

The COVID-19 pandemic forced many financial and credit institutions to reconsider their approach to doing business. Previously, many organizations had their own physical points of service and customer identification widely distributed in all major cities – in banks, offices of credit organizations or simply in crowded places such as shopping centers.
Under the conditions of strict quarantine and subsequent restrictions, companies met their customers online and began to offer remote methods of service and identification. This innovation reduced the level of checks on customer profiles and documents, which gave rise to a new wave of fraudulent attacks on financial and credit institutions.
Organizations that issued instant credit cards or loans in the form of payment cards with a credit limit were attacked. Conventionally, attacks can be divided into several categories.

First-party fraud
In this case, the scammer uses fake income documents to secure a higher credit limit or a cash loan. Documents confirming the place of residence can also be forged in order to confuse and throw off the exchange system of creditors who exchange data on unscrupulous borrowers, as well as complicate the search for a borrower by collection organizations (Figure 8).
Typically, such a scam can be carried out once, taking loans in a short time from many available organizations, and then disappear by turning off the phone.
After the name and surname of such a client enters the bases of unscrupulous borrowers exchanged by credit organizations, repeated loans will no longer be issued.

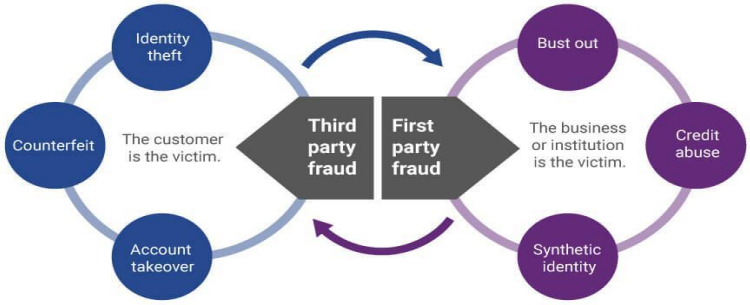

Figure 8. First-party fraud and third-party fraud
Source: Authors' own development

Third-party fraud

This method implies that the loan is obtained on behalf of a person whose documents were stolen in digital or physical form (Figure 8).

An attack from a third party is most often obtained in the form of a transfer to the card number of a criminal, because a credit institution cannot obtain the name of the real owner by the card number, and banks do not disclose such information. Or, if the attacker has access to the victim's online banking, the received loan in the form of wire transfer funds is withdrawn to a front organization or an individual who is engaged in cashing out.

Contacting the police and searching for an individual or the owner of a one-day company who received the victim's funds can take several years and not produce results. At the same time, the victim's credit reputation also suffers, because most likely the victim will find out that a loan was opened for them when the debt has already been sold to a collection agency and the client's profile is included in the database of unscrupulous borrowers – from which it is almost impossible to remove oneself without going through the courts and other bureaucratic institutions.

Such an attack can be multiplied by an attacker as many times as possible to steal the victim's documents or access their account.

Synthetic personalities

This type of attack involves taking loans and credit for a person who never existed. This is the most technologically advanced method in the attacker's arsenal.

With the development of artificial intelligence and machine learning systems, it became possible to fully synthesize the profile of a living person, including their photo, video and voice.

This attack vector uses fake IDs that attackers acquire on the black market and utility bills that confirm the fact of residence, and also synthesizes account statements from well-known banks confirming the solvency of the borrower.

The relevance of the problem is illustrated in Table 3, which shows a significant increase in the number of cases of this type of fraud.

Table 3. The synthetic identities problem

Source: Collected by authors from different sources

Type of identity theft	Number of reports in 2020	Change from 2019
Government documents or benefit fraud	406,375	1,663%
Credit card fraud	393,207	44.7%
Other identity theft	353,152	63.7%
Loan or lease fraud	204,967	95.8%
Employment or tax-related fraud	113,529	149.2%
Phone or utilities fraud	99,539	19.2%
Bank fraud	89,476	52.4%

The attack is carried out in two stages.

First, an account is opened for a synthetic person in a financial institution. In the second stage, loans are received from another company – sometimes from the same one that opened the account, but such cases are rare. Companies rarely issue loans immediately to clients with new profiles without a credit history, but they are always ready to receive a transfer from another financial institution in the form of a loan.

Thus, the attacker deceives two financial institutions at once – the one that opened the account and

the one that issued the loan.
After receiving a loan, the withdrawal scheme is similar to the third-party fraud scheme – the funds are transferred to a company or a person who is an accomplice of the attacker.
The problem with this attack is that the actual borrower does not exist at all, and the number of attempts by the attacker is limited only by how many synthetic profiles are missed by the systems or personnel of both companies that are being attacked.
This method of attack is the most dangerous at the moment, because as techniques are constantly being improved, it becomes more and more difficult to identify synthetic personalities.
Let us consider several examples. In 2019, an executive of a UK-based energy firm thought he was speaking on the phone with his boss, the CEO of the firm's German parent company, who asked him to send €220,000 to a Hungarian supplier (Stupp, 2019). The caller said that the request was urgent, directing the executive to pay within an hour, which he did. Instead of his boss, the executive spoke to a voice recording generated by artificial intelligence-based software that successfully impersonated the CEO.
The live facial biometric data that many digital-only banks rely on to authenticate their customers is not fraud-proof either. Cybercriminals have found a way to recreate 3D models of faces using recorded videos that can be used to log in by generating head tilts and turns on demand.
Social engineering plays a significant role in modern fraud cases. A man behind an Instagram account with 2.5 million followers flaunting his opulent lifestyle told people that they could earn as much as him by sending him money (Dawkins, 2021). He was arrested after stealing over $400 million from individuals and businesses worldwide (Karimi, 2020).

Table 4. The minimum cost of a new identity
Source: compiled by the authors based on Anderson (2020)

	USA	**Canada**	**Australia**	**UK**	**Europe**
New identity: passport, ID card and birth certificate	$1,152	$1,175	$1,355	$1,255	$1,125
Education: high school diploma and bachelor`s degree	0.1699 BTC	0.1699 BTC	0.1699 BTC	0.1699 BTC	0.1699 BTC
Finance: bank account, credit card and 5,000 of counterfeit currency	$115+ 0.0984 BTC	$115+ 0.0984 BTC	$115+ 0.2593 BTC	$115+ 0.1162 BTC	$115+ 0.0894 BTC
Total	$1,267+ 0.2683 BTC	$1,290+ 0.2683 BTC	$1,470+ 0.4292 BTC	$1,370+ 0.2861 BTC	$1,240+ 0.2593 BTC

There are many examples of money flippers on social media who promise to turn $100 into $1,000, $500 into $5,000, and so on (Akpobi, 2020). Suffice to say that people do not receive their investments back. If the recipient is not blacklisted, has a business, and receives money regularly, training a system to detect such type of fraud is challenging, if not impossible, for now.
Payment card and identity fraud are closely tied to criminal activities that aim to launder money and conceal identities. Modern compliance and AML investigations check social media accounts for suspicious posts and activities. To get around these checks, criminals buy inexpensive accounts created and maintained for a few years to develop a plausible online identity.
People who want to take on a different identity can buy a passport and a new identity with social media accounts, diplomas, and other documents for relatively little cost (Table 3). On the on hand, it is easier to obtain a new identity than ever before. On the other, regulators and service providers are tightening security and making it more difficult to evade their checks.
In the attempt to balance convenience and security, security is losing. Customers do not like long

passwords and additional verification methods. Frankly, they do not care if their information leaks because "they have nothing to hide and don't have that much money anyway." AutoML has the potential to slow down the advancement of financial fraud, and the only question that remains is: for how long?

Conclusions

The AutoML approach could take on all of the routine and repetitive tasks (such as data set clean, spot anomalies, and balance) that come with in-depth data analysis, and can extract insights from raw data. This technology can drastically reduce data analysis time.

Our experiments involved data on the actual credit card payments of customers in Western European banks. The data set contains 284,807 observations, including 492 cases of fraud, representing 0.172% of the total sample. Thus, the data sample is highly unbalanced, which is typical of the considered problem.

The result of our experiments, carried out by using automatic machine learning algorithms that make it possible, in a short time, to sort a large number of variants of models and the composition of input data, is the quality of classification – in terms of the AUC metric, from 0.97733 to 0.98249. This is a strong result and exceeds the effectiveness of the classifiers developed by traditional methods. The ensemble of models has allowed us to detect up to 85.7% of fraudulent transactions in the sample. At the same time, the accuracy of detecting fraudulent transactions is also fairly high (79%–85%).

New ways of payment fraud are constantly emerging. The use of AutoML methods allows the adaptation of the security systems of banks and payment organizations to be accelerated, and losses from fraud to be reduced.

References

1. Akpobi, W. J. (2020, September 15). *How Instagram Scammers Make $10k/Month Using the Dumbest Strategies*. Better Marketing. https://bettermarketing.pub/how-instagram-scammers-make-10k-month-using-the-dumbest-strategies-6c481593ac42
2. Anderson, S. (2020). *Dark Web: The Average Cost of Buying a New Identity in 2021*. Safety Detectives. https://www.safetydetectives.com/blog/dark-web-the-average-cost-of-buying-a-new-identity/
3. Brando, R. (2018, May 25). *Facebook and Google hit with $8.8 billion in lawsuits on day one of GDPR*. The Verge. https://www.theverge.com/2018/5/25/17393766/facebook-google-gdpr-lawsuit-max-schrems-europe
4. Dal Pozzolo, A., Caelen, O., Waterschoot, S., & Bontempi, G. (2013). Racing for unbalanced methods selection. In *Intelligent Data Engineering and Automated Learning – IDEAL 2013. IDEAL 2013. Lecture Notes in Computer Science* (Vol. 8206, pp. 24–31). Berlin, Heidelberg: Springer. https://doi.org/10.1007/978-3-642-41278-3_4
5. Dawkins, D. (2021, February 19). *Nigerian Influencer Ramon 'Hushpuppi' Abbas Laundered Funds for North Korean Hackers*. Forbes. https://www.forbes.com/sites/daviddawkins/2021/02/19/nigerian-influencer-ramon-hushpuppi-abbas-laundered-funds-for-north-korean-hackers-says-us-department-of-justice/
6. Fu, K., Cheng, D., Tu, Y., & Zhang, L. (2016). Credit Card Fraud Detection Using Convolutional Neural Networks. In Hirose, A., Ozawa, S., Doya, K., Ikeda, K., Lee, M., Liu, D. (eds.), *Neural Information Processing. ICONIP 2016. Lecture Notes in Computer Science* (Vol.

9949, pp. 483–490). Springer, Cham. https://doi.org/10.1007/978-3-319-46675-0_53
7. *H2O Driverless AI* (2021). https://docs.h2o.ai/h2o/latest-stable/h2o-docs/automl.html
8. Heller, M. (2019, August 21). *Automated machine learning or AutoML explained*. InfoWorld. https://www.infoworld.com/article/3430788/automated-machine-learning-or-automl-explained.html
9. Karimi, F. (2020, July 12). *He flaunted private jets and luxury cars on Instagram. Feds used his posts to link him to alleged cybercrimes*. CNN. https://edition.cnn.com/2020/07/12/us/ray-hushpuppi-alleged-money-laundering-trnd/index.html
10. Kolodiziev, O., Mints, A., Sidelov, P., Pleskun, I., & Lozynska, O. (2020). Automatic machine learning algorithms for fraud detection in digital payment systems. *Eastern-European Journal of Enterprise Technologies*, *5*(9 (107)), 14–26.
11. Machine Learning Group. (2013). *Credit Card Fraud Detection* (Version 3) [Data set]. https://www.kaggle.com/mlg-ulb/creditcardfraud
12. Minsky, M., & Papert, S. (1969). *Perceptrons*. MIT Press.
13. The Nilson Report. (2020, December 1). Card Fraud Losses Reach $28.65 Billion. *The Nilson Report*, Issue 1187, 4–6. https://nilsonreport.com/mention/1313/1link/
14. Stupp, C. (2019, August 30). Fraudsters Used AI to Mimic CEO's Voice in Unusual Cybercrime Case. *The Wall Street Journal*. https://www.wsj.com/articles/fraudsters-use-ai-to-mimic-ceos-voice-in-unusual-cybercrime-case-11567157402
15. Tedder, K., & Buzzard, J. (2020). *2020 Identity Fraud Study: Genesis of the Identity Fraud Crisis*. Javelin Strategy and Research. https://www.javelinstrategy.com/coverage-area/2020-identity-fraud-study-genesis-identity-fraud-crisis
16. Wikiwand. (2018). *Loss functions for classification*. https://www.wikiwand.com/en/Loss_functions_for_classification