Fighting Fraud with a Model of Models

Whitepaper

Utilising Artificial Intelligence
to Prevent Payment Card Fraud

KPMG

nets

Contents

Executive Summary	3
Nets Fraud Ensemble	4
Today's Fraud Landscape	5
Background	7
Leveraging Historic Data	8
Machine Learning in Fraud Prevention: A Primer	9
Using Fraudsters' Limitations Against Them	11
Threshold Optimisation	16
Model Explainability	17
Commercial Benefits Beyond Fraud Prevention	18
Conclusion	19

Executive Summary

Humans cannot compete with computers when it comes to data interrogation. That's why artificial intelligence (AI) and machine learning (ML) hold so much potential – because they present an opportunity to analyse and act on patterns too complex for the human brain to even identify.

Until now, the use of true machine learning to fight payment card fraud has been limited. Yet, we need it. With the total annual value of fraudulent transactions across Europe hitting €1.8 billion[1], the need to step up fraud prevention has never been greater.

Fraud prevention is an increasingly convoluted and nuanced business, as factors such as the mass adoption of e-commerce, increasing cross-border payments, and the growing popularity of new digital payment methods combine to add new layers of complexity.

This paper outlines a machine learning approach to fraud prevention: Nets Fraud Ensemble, which reduces fraudulent transactions by up to 40% for the benefits of banks, merchants and cardholders, as well as society in general.

At its core, Nets Fraud Ensemble uses historic data from a wide range of sources to generate a 'fraud score' that takes into account the inherent limitations faced by criminals attempting to make fraudulent transactions.

Fraud prevention is a service that Nets offers to issuers - but fraud is not just a problem for banks. Online merchants also suffer, as they often do not find out that the payment card used to place an order was stolen and the money refunded to the actual cardholder, until after they have shipped the goods, leaving them with a financial loss. It is also a burden to cardholders: even if they are fully refunded by their bank for all fraudulent transactions, they must still go through a dispute process in order to recover the money, not to mention being without a card for days or weeks until a replacement card arrives or the work associated with updating online subscription services. Finally, wider society is impacted, as the proceeds of crime is often associated with organised criminal activity, including human and drug trafficking / exploitation and terrorism funding[2].

[1] https://www.nets.eu/solutions/fraud-and-dispute-services/Documents/Nets-Fraud-Report-2019.pdf

[2] https://doi.org/10.1108/13685201111098879

Nets Fraud Ensemble

Nets, a European payment services provider, has collaborated with multinational professional services provider, KPMG, to develop Nets Fraud Ensemble, a next-generation fraud monitoring and prevention solution.

Nets Fraud Ensemble is an AI-powered anti-fraud engine that improves real time fraud prevention in an ever-changing landscape. By deploying true machine learning (i.e. a system that automatically identifies and reacts to existing and new fraud patterns), it represents a significant step forward from the rules-based models that are currently in use across the international banking industry.

The 'brain' of Nets Fraud Ensemble consists of multiple models working together to analyse each individual transaction within ten milliseconds – the time frame in which a transaction can be blocked. The solution learns automatically from patterns observed in the data and adjusts accordingly. This means that the longer historic data series available to it, the more fraudulent transactions are blocked, and the fewer false positives are raised.

The solution resulted in an immediate fraud reduction of 25% and an estimated 40% long-term potential.

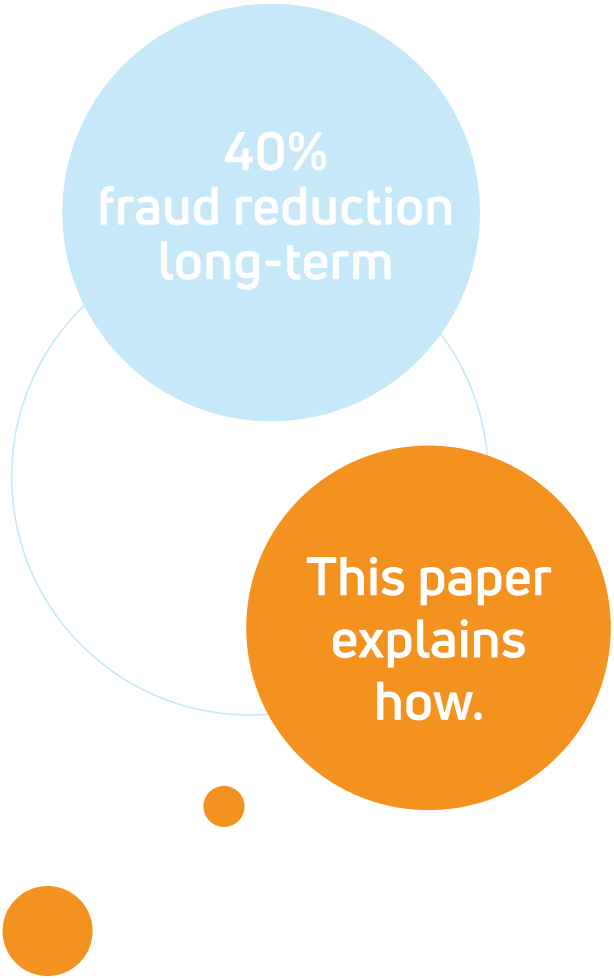

Today's Fraud Landscape

Third party payment card fraud is a growing problem that occurs in two ways. Either the fraudster is in possession of the card and can therefore make card present (CP) transactions, i.e. purchases in physical locations or cash withdrawals from ATMs, or they have only obtained the card details needed for making online purchases, meaning that they are limited to card not present (CNP) transactions.

For the purposes of this white paper, we will use the term 'merchant' to denote all possible entities from which payment card details could be fraudulently accessed and used, including ATMs.

As the popularity of ecommerce for consumer and corporate purchases continues to increase across the globe, CNP transaction volumes are growing steadily. In 2018, international ecommerce grew by 23.3% to reach nearly three billion transactions[3]. Where there is value, criminals will follow. CNP fraud now represents almost 80% of the total volume of fraudulent card transactions across Europe[4].

The modus operandi for traditional fraud prevention decision engines has been for humans to create rules in the 'If X and Y, then Z' format. The decision engine has two possible courses of actions if a transaction is flagged as potentially fraudulent – it either declines the transaction, or it allows it but raises an alert to a team of monitoring agents, who manually review the data and take appropriate action. This does work, but requires hundreds of rules to be effective – and it is highly labour intensive, and therefore costly, to create and maintain these rules, balance fraud prevention with the number of false positives, and maintain a stable alert stream for agents to review.

One last challenge with traditional rules is that a transaction that almost triggers several rules, but does not actually trigger any, will not be flagged by the system.

The card details and other Personal Identifiable Information (PII) required to make fraudulent CNP purchases can be acquired by criminals in a number of ways, which do not always involve phishing. Payment cards can be copied when they are physically handled, such as in restaurants and brick and mortar shops, or when details are provided over the phone, such as to hotels by consumers to make reservations.

Card details can also be illegally acquired after ecommerce websites suffer data breaches, as well as through skimming (either physical skimming via equipment mounted on ATMs or terminals, or digital skimming via malicious scripts embedded on ecommerce or third party providers' websites).

[3] https://www.statista.com/statistics/379046/worldwide-retail-e-commerce-sales/

[4] https://www.nets.eu/solutions/fraud-and-dispute-services/Documents/Nets-Fraud-Report-2019.pdf

Today's Fraud Landscape (cont.)

The sale of compromised payment card details has grown into an illicit online industry worth €1.8 billion[5]. Organised criminal enterprises are taking advantage of pre-packaged solutions, also known as Fraud as a Service (FaaS), which, together with vast amounts of payment card numbers, are available to purchase on the dark web. This is important for fraud monitoring and prevention as it means that there may not be any direct link between how the card details were compromised and the specific fraudulent transactions attempted afterwards. In fact, organised criminals selling payment card details on the dark web routinely mix card details from multiple breaches and split them into smaller, random batches before selling, to further impede investigatory efforts.

Further complicating fraud prevention efforts is the increased use of multiple different third-party providers for payment solutions by ecommerce merchants. If just one of those third party providers is compromised, then subsequently only a subset of payment cards used at that merchant will be compromised. Worse yet, there might not be any data available to the transaction processor that can differentiate between compromised and non-compromised payment cards.

This makes traditional fraud prevention tactics, such as preventively blocking cards suspected to be compromised, undesirable, as too many cardholders will be affected – many of whose card details will not have been compromised at all. This is an area of significant concern for merchants in particular, as 26% of cardholders have reduced their patronage of a merchant following a false decline, and 32% stopped shopping with the merchant entirely[6].

The methods introduced in this paper have been specifically designed to reduce false positives by creating a fraud score that balances multiple minor fraud signals – addressing both the challenges described above and taking advantage of the challenges faced by would-be fraudsters.

The solution was developed solely for the purpose of reducing loss and inconvenience due to payment card fraud. No commercial usage of the data or learnings resulting from the development process have been in scope at any time during the project. In short, Nets Fraud Ensemble has been created without commercial influence.

[5] https://www.nets.eu/solutions/fraud-and-dispute-services/Documents/Nets-Fraud-Report-2019.pdf

[6] https://www.javelinstrategy.com/coverage-area/overcoming-false-positives-saving-sale-and-customer-relationship

Background

Nets invested in an open-source software platform called Hadoop capable of handling the huge volumes of transactional data Nets processes – over 10 million transactions every day. The platform provides vast amounts of storage for any kind of data and enormous processing power.

Nets approached KPMG to support with realising the value of this platform and the two companies collaborated to develop a series of proof of concepts (PoCs).

One PoC was a machine learning alternative to manual rule creation with the objective of reducing payment card fraud by optimising the decision engine that reviews incoming transactions.

The outcome of the PoC revealed a potential for reducing fraud by 15-25% while generating the same number of alerts. It was clear that this potential should not be left untapped and the collaboration between Nets and KPMG continued into a project with two clear goals: realise the promise of the PoC by putting it into production and continue to improve the model behind the 15-25% reduction in fraud. This is what evolved into the Nets Fraud Ensemble model described in this white paper.

15-25%
reduction in fraud

Leveraging Historic Data

As the card processor for more than 240 banks, Nets has built up a unique historic data set consisting of billions of payment card transactions.

All transaction data from the past 2 years, including authorisation and clearing data, is available for analysis and modelling on the Hadoop platform. The data includes the entire ISO string for credit card transactions, properly anonymised and cleared of sensitive information such as the payment card number. Available information on individual transactions includes sender and recipient identifiers, a timestamp, transaction amount, currency, card type, input method, merchant category code (MCC) and much more.

In addition to the basic information on historic transactions, it is also essential to know if a transaction was fraudulent or not – the machine learning algorithm needs this in order to learn to differentiate between normal transactions and fraudulent ones. Over the years Nets has collected copious amounts of data by carefully tracking fraud cases from multiple sources (both directly reported to Nets, flagged by human agents and through card issuers). These data points were all integrated to obtain a vast log of credit card fraud cases. This information served as the starting point for the creation of Nets Fraud Ensemble.

The new solution has created a lot of interest from banks. The on-boarding plan for new banks enables a smooth transition taking into account both the data and application aspect.

Nets is the card processor for 240 banks

Machine Learning in Fraud Prevention: A Primer

It would be easy to prevent fraud if there were some straightforward pieces of evidence (or features, as they are called in the machine learning community) that separate fraudulent transactions from legitimate ones. Unfortunately, that is not the case.

Fraudsters use the same services that are used by genuine cardholders when committing fraud, for a simple reason: if an ecommerce merchant was used solely for fraudulent transactions, it would be identified and shut down very quickly. This leaves fraud prevention teams with the challenge of finding and accumulating multiple features.

Traditionally, this has been achieved, with reasonable success, by creating very specific rules that use multiple criteria to identify patterns. The process of creating and maintaining these rules is labour intensive and is complicated by the fact that less than one in every thousand transactions is fraudulent – making false positives very likely. Finally, it is human nature to independently write rules for each specific fraud scenario, meaning that rules are typically stand-alone and do not support each other.

A machine learning alternative to rule writing that – importantly – has an almost identical production implementation is a decision tree. In a decision tree, a series of yes/no questions are posed, creating a path from the top down through each level of the tree. Any specific observation travelling down the decision tree will meet only one question at each level of the tree, but the total number of questions within a layer doubles at each consecutive layer.

At the bottom of a decision tree, one can have either a label indicating fraud or not (a classification tree) or a score indicating the probability of fraud (a regression tree). It is typically more flexible to work with the score from a regression tree and this is the approach used by Nets Fraud Ensemble. It is a daunting task for a human to create, say, a depth ten decision tree, due to the way they scale. Although the first layer has just one question, the following three layers have two, four and eight questions respectively, and the final layer 1024 questions. Fortunately, efficient algorithms for identifying which questions to ask at each layer exist. Notably, these algorithms identify both which feature to base the question on and the appropriate values. For example, at one location in the tree the question might be 'Is the transaction amount greater than 5.53 EUR?' and at another it might be 'Is the number of online transactions made by the card under consideration greater than six in the past 24 hours?').

Machine Learning in Fraud Prevention: A Primer (cont.)

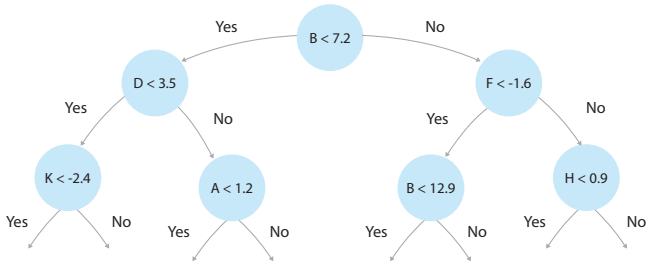

One such algorithm is XGBoost[7], which is popular in the machine learning community. The "boost" in XGBoost refers to a technique known as gradient boosting, where a large number of decision trees are created to form a more accurate model than any of the individual decision trees could, as the scores from each individual tree are added together to form the final score. A simple way to understand this is that each tree is being trained to correct for mistakes made by the previous trees. This method supports creating hundreds of decision trees on thousands of features, and works even on unbalanced data where only one in every 1000 transactions is fraudulent.

A large part of the effort that went into the Nets Fraud Ensemble model was the creation of strong features for XGBoost, some as the result of advanced models.

One important aspect of the XGBoost model is that all decision trees contribute to the final score. This contrasts with traditional rule writing, where one defines hundreds of individual binary rules, each with a specific focus. These hundreds of rules do not interact, partly because this would be very hard for humans to maintain. This means that a fraudulent transaction can go undetected if no single rule is triggered – even when multiple rules are approaching their trigger thresholds. With the XGBoost approach, a single holistic fraud score is created: the probability of a transaction being fraudulent when all data has been considered.

A single rule can then be created that defines at what threshold the fraud score requires the system to decline transactions or raise alerts – this threshold can, in addition to the score, take transaction amount into consideration, as described in the section 'Threshold Optimisation'.

The score approach has several advantages beyond preventing more fraudulent transactions than traditional methods, including reducing false positives, ease of integration and the facilitation of two-way cardholder communications. These are discussed in the section titled 'Commercial Benefits Beyond Fraud Prevention'.

[7] https://xgboost.ai/

Using Fraudsters' Limitations Against Them

This section describes the hundreds of additional features created from Nets' historic dataset, which the XGBoost model uses in addition to simpler features such as transaction amount. We have categorised the features as follows:

- Common Point of Purchase
- Merchant Characteristics
- Cardholder Characteristics
- Time Sensitivity

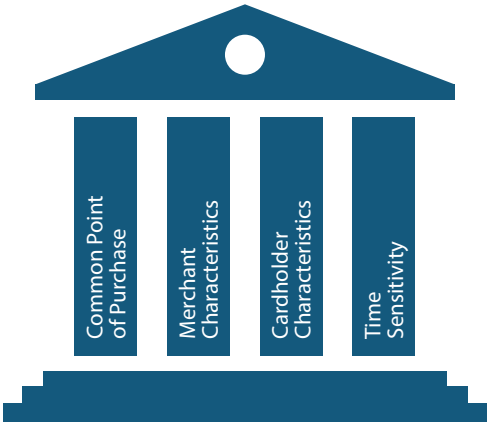

When creating the features, the various limitations faced by would-be fraudsters and the strengths offered by the production platform were considered. All features in each category are dynamic and adapt automatically to evolving fraud patterns.

Common Point of Purchase

Fraud is a two-step process. First, the payment card details are compromised, and second, the compromised payment card is abused. This is a limitation faced by criminals: they cannot commit fraud without access to compromised card information. This means that a model that predicts the likelihood of card details being compromised for each card in the portfolio can be used against attempted fraudsters.

Until now, fraud prevention teams have approached this by finding the Common Point of Purchase (CPP) i.e. the merchant that represents a common denominator among reported compromised cards. Once a suspected CPP is identified, a business decision can be made to preventively replace all cards that have visited this merchant within a given time window or leave it open and accept the risk (based on the assumption that perhaps only a subset of cards were compromised and that most compromised cards have already been misused and subsequently blocked).

Using Fraudsters' Limitations Against Them (cont.)

Traditionally, the CPP is identified by reviewing the cards used to make ecommerce transactions from merchants within a given time window and calculating the percentage of unique payment cards that have subsequently been used in fraudulent transactions. It is uncommon to analyse the entire history of transactions - sometimes as few as 500 reported cards are included in the analysis (and no cards without fraud reports are included).

This approach can result in artefacts, i.e. legitimate merchants being identified as CPPs. This can happen because of high transaction volumes (applicable to popular ecommerce and brick and mortar retailers) or because of confounding effects (many payment cards are compromised during travel, and many cards are used at airports during travel, but that does not mean that an airport merchant is the CPP).

Using the computing power of the Hadoop platform, it is possible to reduce artefacts. Nets Fraud Ensemble does this in four key ways.

1. It uses all payment card transaction data to identify the most likely CPPs, not just data from a subset of cards reported for fraudulent use. This encompasses over 20 million cards that have made more than five billion transactions.
2. It calculates all CPP probabilities simultaneously instead of individually, leading to a better signal-to-noise ratio.
3. It has a flexible structure that allows seamless CPP calculations on each combination of "merchant" and "time". Here, the "merchant" definition can be selected either as a particular payment terminal, merchant ID, merchant name (using fuzzy matching), etc. and the "time" definition can be any number of weeks (one and four being used by Nets Fraud Ensemble).
4. It aggregates the CPP probabilities into card-based probabilities we call CardProb. This means that a card used at multiple low-to-medium risk CPPs would still be identified as high risk. The CardProb is an important feature taken into account by the model, as it enables different reactions to low risk and medium-to-high risk cards. The XGBoost method automatically learns this and appropriately combines the CardProb score with other relevant features of the model. Everything is learned from historic data and no human input is needed in the model calibration.

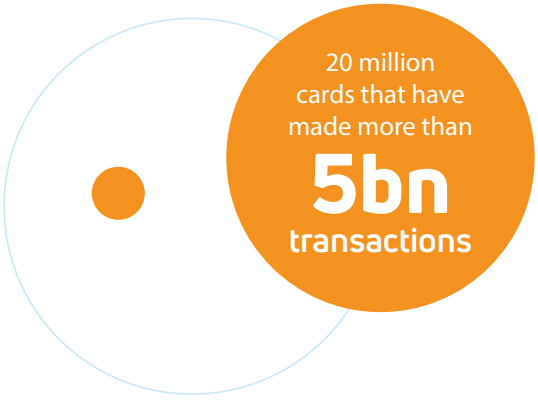

Using Fraudsters' Limitations Against Them (cont.)

Merchant Characteristics

A second significant limitation facing fraudsters is that only a small subset of merchants and bank managed ATMs constitute good places to actually commit fraud. It is a straightforward process to estimate the probability that a transaction is fraudulent based on the merchant (by calculating fraudulent transactions as a percentage of the total transactions accepted by the merchant), but this is not sufficiently accurate alone. Fortunately, once it has been added as a feature, XGBoost will identify how to use it in combination with other features. Similar features are created based on other characteristics of a transaction, such as payment method, MCC, currency and country. Essentially, any dimension where there is a distributional difference between fraudulent and legitimate transactions is added as a feature, enabling the model to combine and use these pieces of evidence.

In addition to the above "one-dimensional" features, Nets Fraud Ensemble also uses Bayesian techniques to aggregate multiple one-dimensional features into a multidimensional feature. This high-level feature is added to the feature pool and further features are created from it. This is a recurrent theme in our feature engineering – we create higher level, more complex features by building on simpler ones.

The background for creating the above mentioned features is that XGBoost is created to work on numerical features (such as amount) and not on categorical features (such as country – which can be coded as a number, but there is typically no natural relation between the country coded as 1 and the country coded as 2). There are many alternative ways of encoding categorical variables, such as mapping them into continuous values between 0 and 1 where the numerical value represents prevalence, i.e. values closer to 1 represent more frequently seen values. This gives the XGBoost model access to the entire complexity of the categorical variable as well as information on the prevalence.

There are other pieces of information about merchants that are simple to calculate yet still useful as features – these are referred to as simple merchant characteristics. Examples include the average number of transactions per day, average transaction amount, lower and upper quantiles of the amount distribution (the amount for each merchant which 25% of purchases are below and above respectively), various dispersion measures (for example, indicating if the merchant has spikes in payments or if the transaction amounts are stable over time) and the number of days since the merchant began accepting transactions. This last figure is important as sometimes fake merchants are created for the sole purpose of committing a significant amount of payment card fraud before they are identified and shut down. More advanced information can be gleaned by segmenting merchants, which is essentially a data driven alternative to MCCs.

Using Fraudsters' Limitations Against Them (cont.)

Cardholder Characteristics

A third limitation faced by criminals attempting to make fraudulent transactions is that they do not know what the normal behaviour of the legitimate cardholder looks like. So, although there are a number of possible merchants at which they could attempt making fraudulent transactions (subject to the limitations discussed earlier), they most likely do not know which will look least suspicious considering the legitimate cardholder's purchase history. This can be leveraged by building model-based features for assessing the transaction probability conditional on cardholder characteristics.

In simple terms, this is where it gets personal. Cardholder information is the most useful input into any fraud model, as a particular online transaction might be perfectly normal for one cardholder but extremely unlikely for another.

Basic metrics can be accessed from transaction history, including but not limited to card age, number of different currencies used, number of different merchants used, number of different MCCs used and normal spend (for example number of transactions per hour/day/month or average/maximum amount, both overall and for selected merchants, MCCs and currencies). This gives the model the ability to asses a new transaction compared to historic patterns – large transactions or multiple currencies might be perfectly normal for certain cards, such as corporate cards, but should raise suspicion for others.

A powerful feature is to assess whether a specific merchant and cardholder combination is likely. Nets Fraud Ensemble applies two separate methods for assessing this based on historic transactions. The first method verifies if a certain aspect of an incoming transaction (whether that be currency, country, merchant, MCC or a number of other selected dimensions) has been observed among past transactions for that specific card. This is in principle straightforward - the complexity arises from implementing this in a way that provides a result within the 10-millisecond constraint on a system handling many millions of transactions a day.

One limitation of this method is that it only indicates if the card has been used with the merchant in consideration before or not – it does not indicate if the card is likely to be used at the merchant. The second method addresses this limitation by creating a model that predicts whether the cardholder would be likely to make a purchase from a particular merchant that the cardholder has not necessarily previously bought goods or services from, based on historical transaction data for the card in question and all other cards. This is similar to recommendation engines on websites like Netflix and Amazon, but reversed. Instead of recommending transactions to cardholders, it calculates whether a particular transaction is likely or not given past card usage.

Using Fraudsters' Limitations Against Them (cont.)

Time Sensitivity

A fourth limitation is that high-value fraud transactions are easy to spot. As such, criminals typically attempt to perform multiple small transactions within a short amount of time, because their aim is to withdraw as much value as possible before the cardholder notices any unusual activity.

This can be taken advantage of by creating features that look at recent behaviour and compare it to normal behaviour. Much of the comparisons to normal behaviour have been covered in the preceding sections so this section focuses on how to measure deviation from the norm in real-time – all previously mentioned features are updated at least weekly, but many are not updated live. Let's look at an example. One useful feature of the popular SAS Fraud Management production platform is that it facilitates calculating information such as the number of transactions in the last five minutes, the total amount spent (converted to a joint currency) in online shops in the last hour, the average MCC-based fraud probability over the last four hours and the minimum model based transaction recommendation probability over the last eight hours for each card. These values are commonly referred to as 'time trains' and can enter on equal grounds with any of the other features used in the XGBoost model.

The time train features are typically the highest feature layer – as discussed in the section titled 'Using Fraudsters' Limitations Against Them', more advanced features are built on top of simpler features.

One specific example is the maximum and average time trains for various time windows of the joint multidimensional feature mentioned in the Merchant Characteristics section. Other filters can be applied, such as only considering CNP transactions.

Another example of applying time trains is in the case of a card attempting to make a transaction at a merchant with an increased risk score. The model knows from the transaction history that the cardholder mostly uses the payment card at low risk merchants, so that alone flags a certain level of risk. If, however, the average merchant risk score of the last ten transactions or two hours is also higher than average – i.e. if we have seen a significant change in spending habits over a short amount of time – that flags a much higher level of risk, and the XGBoost model can learn this.

Combining this depth of analysis with vast amounts of relevant historic cardholder information creates powerful fraud detection features. This also applies to other models, for instance the one provided by the SAS Fraud Management platform – Nets and KPMG's development of so-called "meta-features" on top of the model score supplied by the SAS Fraud Management platform makes Nets Fraud Ensemble significantly more powerful than any out-of-the-box solution.

Threshold Optimisation

Keeping the number of false positives low while preventing as much fraud as possible (measured in total monetary value, rather than the number of transactions) is the primary objective of a fraud prevention system.

To achieve this balance, it is considered best practice to include monetary limits when writing manual rules – meaning that even if a transaction is flagged as high risk, it will only be blocked if its value is greater than the specified limit. The primary issue with this approach is that fraudsters are not passive participants; if they can identify the limit, they will simply circumvent the rule by making fraudulent transactions for a slightly lesser amount. This is a key motivation for moving away from individual rules and limits and instead implementing a system that produces a holistic fraud score.

When combining the improved accuracy of a holistic fraud score and the need to minimise false positives, best practice would be to decline transactions above a certain fraud score and monetary value. Nets Fraud Ensemble enables fraud prevention efforts to go beyond this by finding an optimal decision boundary spanned by the model score and the transaction amount (in a single consolidated currency). Essentially, for a given model score, one is searching for the optimal amount cut-off while letting the cut-off vary depending on the model score. For example, if the model score states that a given transaction is 42% likely to be fraudulent, then we are looking for the transaction value above which we should decline transactions, and this value might be different depending on whether the model score is 38% or 46%. It is possible to formulate the search for these cut-off values as a 'constrained optimisation problem' – problems for which a quantity is to be minimised or maximised subject to constraints. Using this approach, it is possible to optimise the total fraud prevented while at the same time minimising the number of false positives.

This results in a much richer decision boundary that provides the optimal amount threshold for each specific risk score. In addition to increasing the amount of fraud prevented at a given level of false positives, it also makes life harder for would-be fraudsters, because the monetary threshold is unique to the risk score, which is not known to the would-be fraudster, making it unlikely to be reverse-engineered.

Furthermore, this approach allows for the creation of separate decision boundaries for prevention and alert generation (based on different false positive requirements). This optimises the value-add provided by the human agents reviewing the alerts as they can prioritise transactions and therefore maximise their capacity to prevent fraud. Finally, this decision boundary can be easily updated in order to adjust the number of alerts to the current capacity of the agent team. Importantly, this only requires adjusting a few numbers (the model does not have to be changed) – something that is not possible with a traditional rule-based approach.

Model Explainability

Model explainability is the procedure of exposing the internal mechanics of a decision made by an artificial intelligence solution in a way understandable to humans. There are clear benefits to being able to explain the predictions of a machine learning model. First, as part of the GDPR regulation, cardholders have a 'right of explanation' for why a transaction is declined. Second, it enables agents to quickly assess why the model has decided to raise an alert, allowing them to look for transaction patterns that enable them to either escalate or clear the alert.

Being able to examine the model's predictions also creates confidence among internal stakeholders, is useful for identifying improvements, and gives agents (who review alerts), analysts (who examine root causes and identify new trends) and data scientists (who add new components to the machine learning framework) a deeper understanding of the model – particularly where and, crucially, why it makes mistakes. This facilitates improvements and is fundamental when developing new features.

Nets Fraud Ensemble uses a framework for model explainability based on Shapley values, which have been proven optimal for model explainability[8]. Shapley values are a concept originally developed for the mathematical area called game theory and beyond the scope of this white paper.

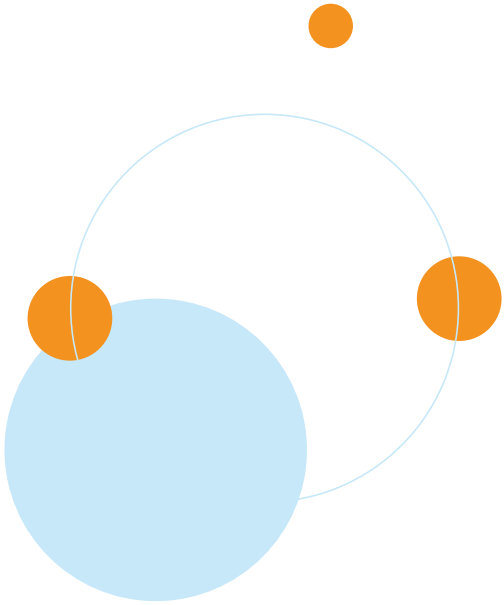

[8] https://arxiv.org/abs/1705.07874

Commercial Benefits Beyond Fraud Prevention

The holistic fraud score approach has several advantages beyond fraud prevention. It makes it much easier to control the false positive rate, which improves the cardholder's user experience – a crucial aspect of fraud prevention for banks.

Nets Fraud Ensemble has low maintenance requirements as the model can be trained (or recalibrated) regularly at the press of a button – a process not possible when using the traditional approach where the system consists of hundreds of human generated rules. The use of machine learning means that manually updating or removing old rules is no longer required. It also makes it easy to adjust the number of generated alerts to fit the capacity of monitoring agents.

Nets Fraud Ensemble automates a range of daily decisions. For instance, it is no longer necessary to deliberate over whether a card should be preventively replaced due to it being used at a suspected CPP, because the model automatically uses the CardProb in calculating the total score and can incorporate a high number of cross-interactions to other features of the model, thereby minimising cardholder inconvenience.

Furthermore, the score approach simplifies connecting the model's recommendations to other services provided by Nets as part of its full-stack business model, including soft and smart-block. Soft-block is a temporary block of the card in Nets' back-end, enabling rapid reopening, for example if the customer is abroad and has a need for temporarily using the card. The smart-block service enables the issuer to provide a flexible replacement of compromised cards to the cardholder, allowing them to continue using the compromised card in a restricted manner determined by Nets, based on Nets' assessment of risk and fraud patterns, whilst waiting for the new card to arrive. Automatic card blocking is normally done by identifying high performing rules and having these activate relevant card blocking procedures. It is labour intensive to maintain a list of high-performing rules and it means that medium-performing rules will never block cards automatically.

With the score approach, one can, in a single place, specify at what model confidence a specific card block should take place. Two way communication is, as the name indicates, a communication channel direct to the cardholder, enabling push of a button responses from the cardholder, to either a blocked transaction or a transaction exceeding a certain threshold and thereby classified as "suspicious".

Nets Fraud Ensemble also opens up bespoke thresholds, giving banks the freedom to customise what to prevent and review, enabling alignment of fraud prevention with their individual risk appetite and minimising false declines.

Conclusion

Creating a model of models has the advantage of collating both human and machine-generated information in a single framework that can then generate the most accurate fraud score possible.

It is impossible for humans to compete with machines when it comes to optimising thresholds – the human input adds the most value by teaching the system what features to look for. This symbiosis between man and machine enabled by Nets Fraud Ensemble resulted in an immediate fraud reduction of 25% and an estimated 40% long-term potential, but that is not the only advantage of this approach.

Merchants reduce their financial losses; cardholders benefit from a greatly improved user experience; and society benefits too. With every fraudulent transaction that is blocked, criminal and terrorist activity is hampered – making the world a safer place.

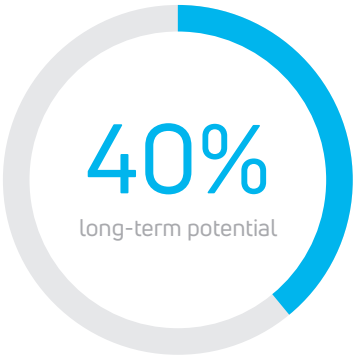

To learn more about Nets Fraud Ensemble, visit
Nets.eu/fraud-ensemble

About Nets

At Nets, we see easier products and solutions as the foundation for growth and progress – both in commerce and society. With headquarters in Copenhagen, Denmark, and 4,100 employees located in various European countries, we help financial institutions, businesses and merchants across the Europe make tomorrow a little easier for their customers while delivering unrivalled security and stability. This has made us a trusted partner to more than 700,000 merchant outlets, including 140,000 online merchant outlets, more than 260,000 enterprises and over 250 banks across Europe.

Powering payment solutions for an easier tomorrow.

About KPMG Denmark

KPMG is a global network of professional services firms providing advisory, audit and tax services. We operate in 153 countries and have more than 207,000 employees working in member firms around the world. We work closely with a broad range of clients, such as business corporations, governments and public sector agencies and not-for-profit organisations, working shoulder to shoulder making positive and sustainable changes in their organisations. KPMG in Denmark has more than 600 employees and is one of the fastest growing professional services firms in Denmark, with a revenue of close to DKK 700m in FY18.