Scalable Data Analytics,
Scalable Algorithms, Software Frameworks
and Visualization ICT-2013 4.2.a

Project	**FP6-619435/SPEEDD**
Deliverable	**D7.1 (update)**
Distribution	**Public**

http://speedd-project.eu

User Requirements and Scenario Definitions (update)

Chris Baber, Sandra Starke, and Natan Morar
(University of Birmingham)

Ivo Correia
(Feedzai)

Status: Final

July 2015

Project

Project Ref. no	FP7-619435
Project acronym	SPEEDD
Project full title	Scalable ProactivE Event-Driven Decision Making
Project site	http://speedd-project.eu/
Project start	February 2014
Project duration	3 years
EC Project Officer	Alina Lupu

Deliverable

Deliverable type	Report
Distribution level	Public
Deliverable Number	D7.1
Deliverable Title	User Requirements and Scenario Definitions (updated)
Contractual date of delivery	M16 (June 2015)
Actual date of delivery	July 2015
Relevant Task(s)	WP7/Tasks 7.1
Partner Responsible	Feedzai
Other contributors	UoB
Number of pages	30
Author(s)	C. Baber, S. Starke, N. Morar and I. Correia
Internal Reviewers	A. Artikis
Status & version	Final
Keywords	Fraud analysis, Human Factors, user requirements

Contents

Tables

Figures

1. Executive Summary

This report is an Appendix to D7.1 “User Requirements and Scenario Definitions” for the Credit Card Fraud Use Case. In the original document there was no consideration given to the requirements of users of the SPEEDD prototypes. There were two reasons for this omission. The first lies in the nature of automated systems and the underlying assumptions relating to the need to remove humans from the control loop. In credit card fraud detection, a key goal is in reducing decision time to fractions of a second. Obviously any human intervention in this decision process would be costly in terms of time and therefore it makes sense to remove human intervention from this process. The scenarios described in D7.1 were primarily informed by the ambition to automate the decision process and this meant that the need for user requirements was obscured. The second reason lies in the challenges that the SPEEDD consortium have faced in gaining access to organisations that were willing to discuss approaches to fraud detection and analysis. This is an ongoing challenge for the project. However, as D7.2 as shown, the consortium has begun to make in-roads to such organisations.

From discussions with Subject Matter Experts, it is apparent that there are several different roles involved in the response to credit card fraud. These range from the Call Centre operative who responds to an alert raised by the automated systems relating to a specific credit card transaction, and follows a well-defined script to speak with the cardholder to check the transaction or to explain the reasons for a card being declined, to the Supervisor who oversees Call Centre operations and looks for patterns in transactions being referred or declined, to the Case Analyst who develops hypotheses for the types of pattern being identified and reports these patterns, to the Fraud Analyst who develops and reviews the rules that the organisation applies in handling fraud. A key differentiator between these roles is the number of transactions that they handle, e.g., a Call Centre operative might deal with some 200 transactions a day while a Case Analyst might deal with 10 transactions. This assumes that transactions are individual events involving individual accounts and cardholders. It is apparent that the role of the Supervisor and the Fraud Analyst is often to explore larger collections of transaction data to look for common patterns across these.

While this differentiation of roles has been instructive, it is apparent that different organisations arrange the roles and duties in their own manner and that there is little commonality across organisations or across countries in defining such roles. Further, roles such as Call Centre operative could be out-sourced and performed by personnel not directly employed by a credit card issuer. This raises interesting questions regarding access to material and data held by the issuer, particularly where this material could help in determining whether an account is suspicious but which might be deemed secure (and thus not shared) by the owner of such data. Banks have been developing consortium-level pooling of relevant data to support investigation but this is carefully managed to protect confidentiality. Even when data is available and accessible, it is usual for this to be presented on a system which is different from the one that the analyst uses for everyday work. Thus, the analyst will have a number of screens in the workspace and each screen will be used to access different sources of information –

with different user interfaces and different formats for the data. Understanding user requirements, therefore, is partly a matter of knowing how the different analysts work and partly a matter of knowing where the data they require is sourced and how it is shared. This means that the user interface (UI) requirements for SPEEDD will differ according to the roles of the people who will be using the system and the purposes for which they will use it. This report presents an outline of the key roles and information requirements that have been identified to date. As with any requirements document, the intention is for this to be updated and treated as a 'living' document over the next 12 months of the project; as further discussions with analysts are undertaken and as initial prototype designs for the UI are developed, the set of requirements will be consolidated.

In conclusion, it is worth comparing the approach taken in WP7 Credit Card Fraud with that taken in WP8 Road Traffic Management. In the latter, there is a well-defined domain of activity (the Grenoble South ring-road) and a well-defined and accessible set of operators, working in the control room in Grenoble. This means that defining requirements on the basis of the operator goals and activities can follow conventional Human Factors practice, i.e., field observations and interviews, Hierarchical Task Analysis and Cognitive Work Analysis, leading to specifications for UI to support specific types of task. For the former, it is more difficult to apply such approaches because of the uncertainty surrounding the nature of the activity being supported. Knowing that there are several types of analyst involved in credit card fraud and appreciating differences between these analysts has helped to develop the approach to UI design in SPEEDD for WP7.

2. Introduction

History of the Document

Version	Date	Author	Change Description
0.1	19/06/2015	Chris Baber	First version of the document
0.2	25/06/2015	Chris Baber	Edit with Feedzai contribution
0.2	29/06/2015	Alex Artikis	Review comments
1.0	03/07/2015	Chris Baber	Final version

Purpose and Scope of Document

The purpose of this document is to extend D7.1 to include an analysis of user requirements for the design, development and evaluation of the SPEEDD prototypes for the Credit Card Fraud Use Case. The reason for the extension was to allow for partners to meet with Subject Matter Experts and develop an appreciation of the requirements of prospective users of the SPEEDD prototypes.

Relationship with Other Documents

As noted in the previous section, this document is related to the D7.1 User Requirements and Scenario Definitions, D7.2 Evaluation, and D5.1 Design of User Interface for SPEEDD Prototype.

3. Defining Credit Card Fraud Use Case Requirements

3.1 Introduction

As noted in D7.2, a key objective of SPEEDD is to provide a user interface which will be accepted by the users and which will help them during work, specifically in terms of their ability to make. For the Road Traffic Use Case, the definition of 'users' and the nature of the work that they perform was easy to determine and to analysis. For the Credit Card Fraud Use Case, the definition of 'users' is complicated by two factors. First, a global aim of SPEEDD (and the fraud detection industry in general) is the elimination of human analysis from the fraud detection cycle. This explains why there is no mention of User Requirements in the first version of D7.1. Elimination of the human operator, in this context, takes two forms: (a.) relegation of the role of the human analyst to a Call Centre operative, following a script to check that a transaction on a card can be accepted, and (b.) increasing demands on automated systems to make accurate decisions within milliseconds. In both cases, the notion that a human operator should be in the path from Point of Sale to bank decision makes little sense. Consequently, the user interface should either support Call Centre operatives (often dealing with flagged transactions on a specific account) or with Supervisors (often monitoring and overseeing groups of Call Centre staff, and detecting patterns in fraud activity), or Analysts (developing and refining the algorithms used by the automated systems). While the SPEEDD project has gained insight into the roles of the Call Centre staff and Supervisors, it has still not been possible to have conversations with fraud analysts 'on the record', i.e., in a manner which would allow the conversations to be reported. Thus, the requirements to date reflect the access gained and the reportable conversations.

3.2 Outline Requirements

D5.1 presented an initial set of requirements: "explaining the results of the models in a human-friendly way", "reducing false alarms to reduce alert fatigue", "ability to move from explanation visuals (what is happening now) to exploration visuals (why something happened", and "dealing with time-changing results and dealing with many dimensions and variables."

D7.2 elaborated these requirements through discussions with FeedZai, FICO and UK Cards Association. We divide these requirements into those which relate to the content of the User Interface (UI), those which relate to the tasks of the analysts and those which relate to the overall goals of the systems in which the analysts work. Noting that there are several roles for analysts in the system, care needs to be taken to ensure an appropriate match between UI design and role; presenting information which is not appropriate for a role can lead to confusion or distraction. In broad terms, we have identified those analysts who concentrate on the transaction as an individual activity which has been flagged as suspicious, those analysts who concentrate on collections of suspicious transactions, and those analysts who create and review the rules used by the automated analysis systems.

Discussion with FICO suggested that the ‘gold standard’ for determining whether an act was fraudulent or not was the information obtained from conversation with the card holder. This suggests a degree of skill involved in the interviewing of card holders when a transaction has been flagged for review or has been declined. Paradoxically, much of this ‘skill’ is enshrined in scripts which are used to guide the conversation between an analyst in a Call Centre and the cardholder. This suggests that the focus is less on determining the nature of fraudulent activity and more on ensuring that the cardholder is a legitimate user of the card. Such a focus explains why banks and card companies are moving from Call Centre conversations to either sending Short Message Service (SMS) requests to the cardholders, or automated telephone calls to cardholders, to seek confirmation that the card is in their possession and by used by them. From this perspective, the UI for the Call Centre analyst *could* contain information about the current transaction and the cardholder’s account, but it is equally possible that such information will be accessed and reviewed during the automated process (see chapter 4). A UI to help analysts to review suspicious account activity should be able to a scripted conversation with the account holder and should provide sufficient information to enable the analyst to review activity on that account. This would require:

- Client and card history
- Location - physical (cardholder, shipping address, billing address, merchant)
- Location – digital (IP address of cardholder, merchant)
- Time of day of transaction
- Time of year of transaction (day/ month, season, festivals)
- Client risk score
- Amount in current transaction
- History of transactions on the account
- Account summary
- Customer summary (payment schedule, delinquency)
- Comments on current and previous reviews

For a supervisory role, an analyst might wish to review cardholders or accounts which have similar risk scores or similar case tag (i.e., classification as fraud or specific type of fraud). This information might be presented on a timeline or on a map to allow patterns of be explored.

In terms of the tasks that analysts might perform, the discussions in D7.2 suggest the following:

- Tagging fraudulent transactions
- Understanding the output / results of automated systems
- Exploring multidimensional, multi-source data
- Communicating findings, decisions and explanations (to other analysts, to other agencies, to cardholders)

We think that it will be particularly important to understand the relationship between the ‘situation space’ which defines the context in which frauds occur and the patterns of fraud which can be discovered through exploring the visualization, and the ‘decision space’ which defines the response that the analyst can make. This response is not simply a matter of Authorise / Refer / Decline (which, anyway, are primarily made by the automated systems which support fraud detection), but are a matter of interpreting the information provided by the automated systems, the information provided by the cardholders (if contacted), and the decision to pursue a line of enquiry (given the cost of collecting and reviewing multi-source data and the amount of money that could be recovered).

Finally, the role of the analyst will be determined by the goals that the overall system in seeking to achieve. In D5.2, the Abstraction Hierarchy (from Cognitive Work Analysis) was used to highlight the range of competing subgoals (termed ‘values and priorities’) which need to be borne in mind during analysis. Given this set of subgoals, there will be trade-offs and competition between sets of these, and the role of the analyst (particularly at the supervisory level) will be to manage these trade-offs within the policies of the organization in which they work and in terms of the assumptions driving the automated systems supporting analyst activity. Figure 1 shows the revised Abstraction Hierarchy (modified as a result of the discussions reported in D7.2).

There are several differences between this figure and the one developed in D5.1. The most notable relates to the number of items under the lowest level (Physical Objects) description. As the project team gains more insight into the nature of the work that fraud analysts undertake and as we visit and speak with more analysts so we are gaining a clearer notion of what sources of information they use to perform their work. A second point arising from our analysis is the notion that there are several types of analyst in the system. This is illustrated by figure 2, in which we have indicate a ‘Call Centre analyst’ (who speaks with the cardholder), a ‘Case analyst’ (who considers flagged transactions at the account level), a ‘Supervisor’ (who oversees the work on the case analysts and is looking for patterns in batches of transactions), and the ‘Fraud analyst’ (who is looking to define and review the rules applied by the organization to recognize and handle fraudulent activity).

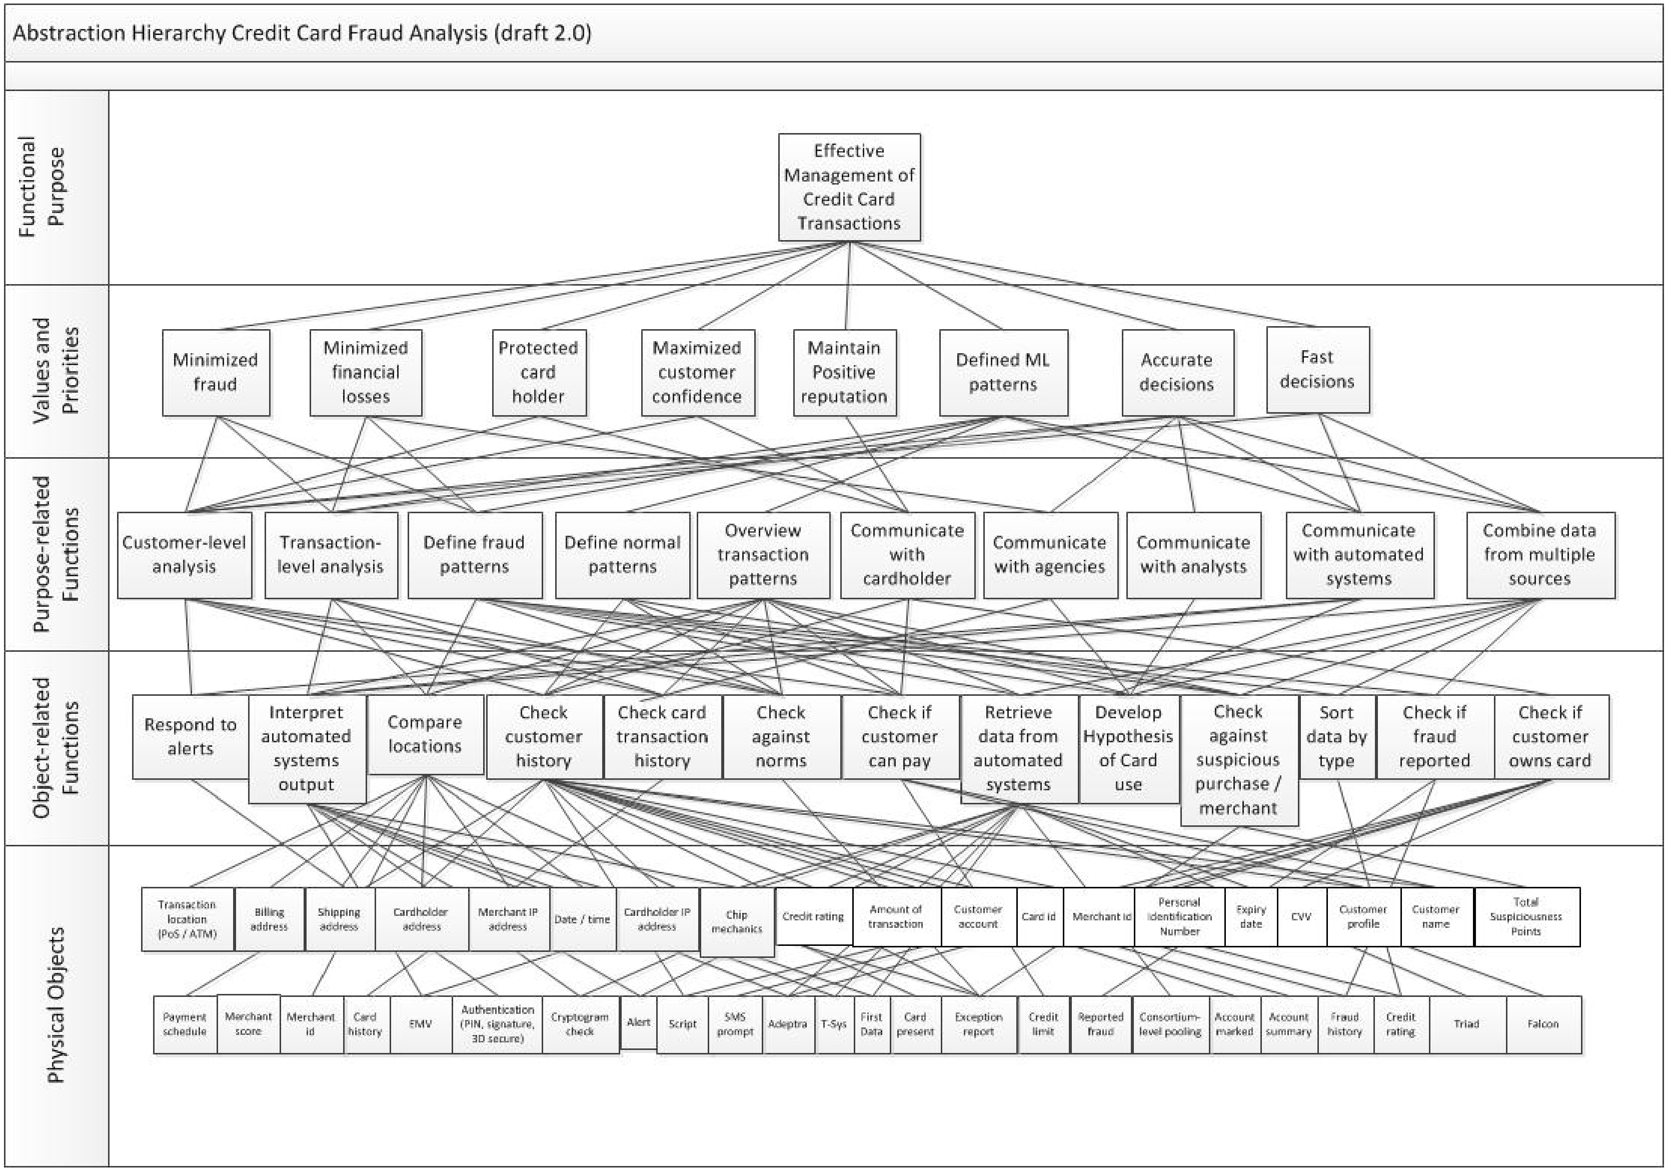

Figure 1: Abstraction Hierarchy

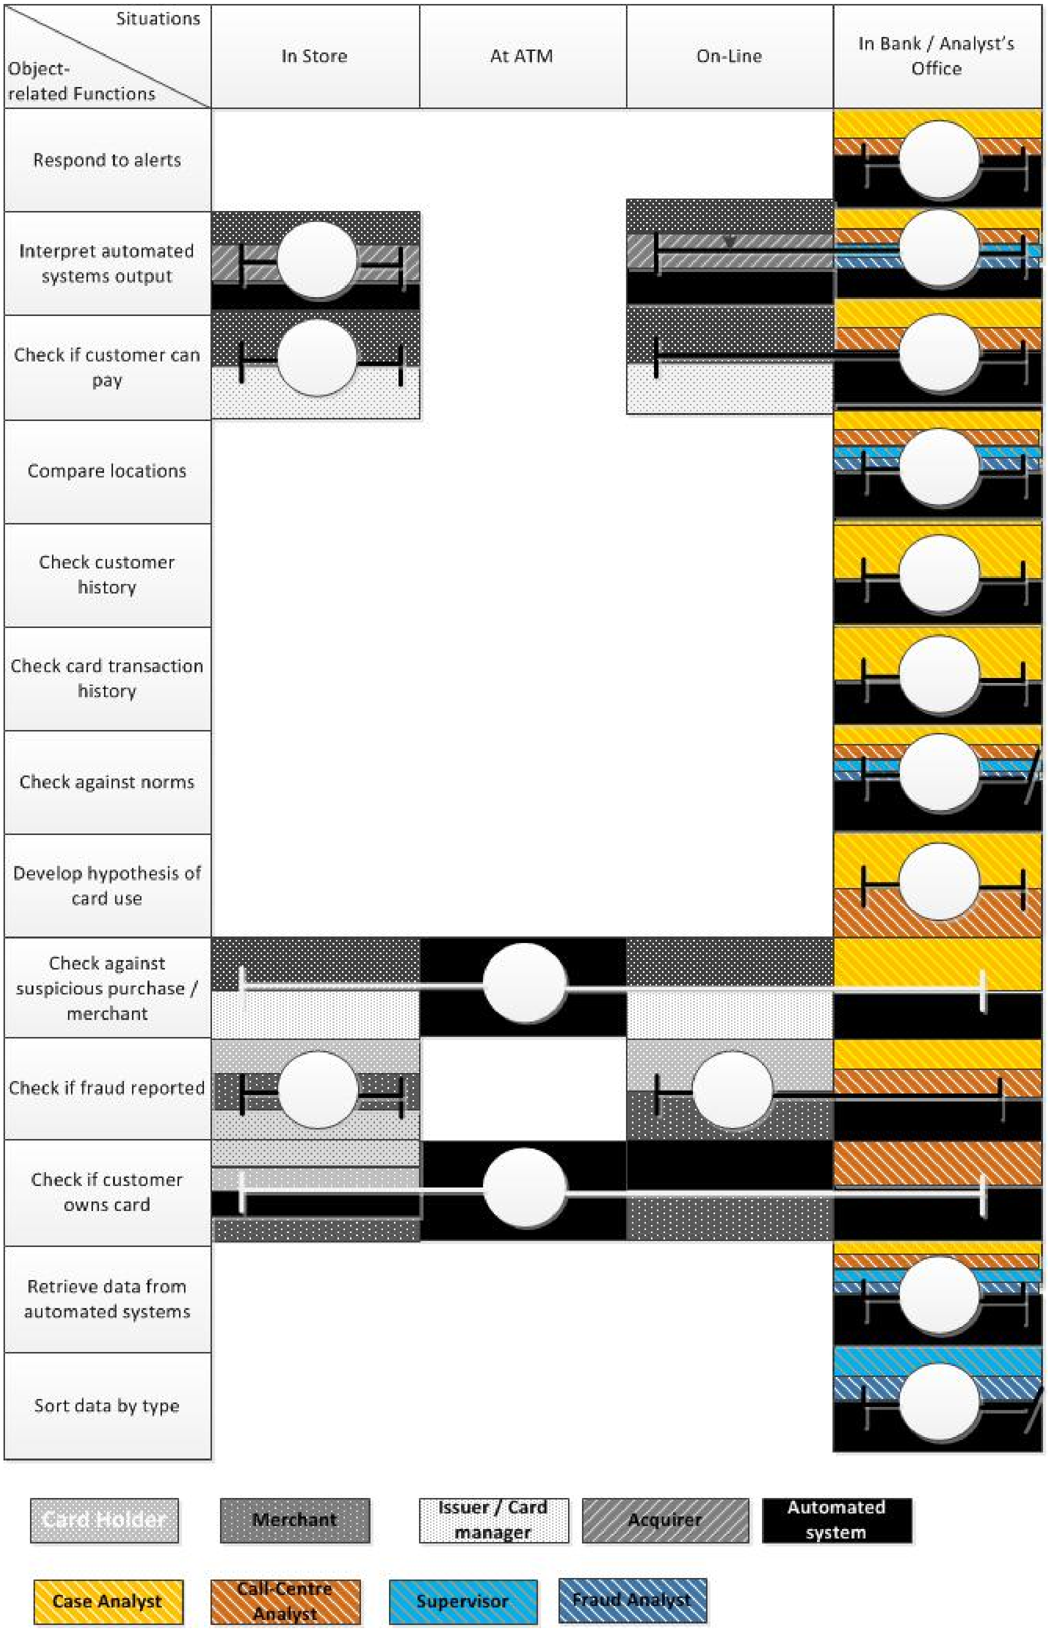

Figure 2: Social Organisation and Cooperation Analysis (SOCA) showing different analyst roles mapped on to Object-Related Functions

4. Processes of Fraud Investigation

4.1 Introduction

In this section, the processes involved in fraud analysis and investigation are summarised. Drawing together material from D5.1 and D7.2, the project team is developing a detailed understanding of the tasks of the analyst. As noted in section 3, there are several different types of analysts in credit card fraud investigation and this will mean that UI design might need to be adapted for the needs of each analyst (most probably as a suite of UI designs). In this section, our concern is with the type of information that analysts might need to use and the manner in which this information might be used.

4.2 Process model of fraud analysis

In D5.2, a simple process model of fraud analysis was presented. This assumed that the analysis was performed in response to the output from an automated system. In order to situate this activity in the wider context of credit card use, figure 2 shows how different decisions are made from Point of Sale to transaction authorisation.

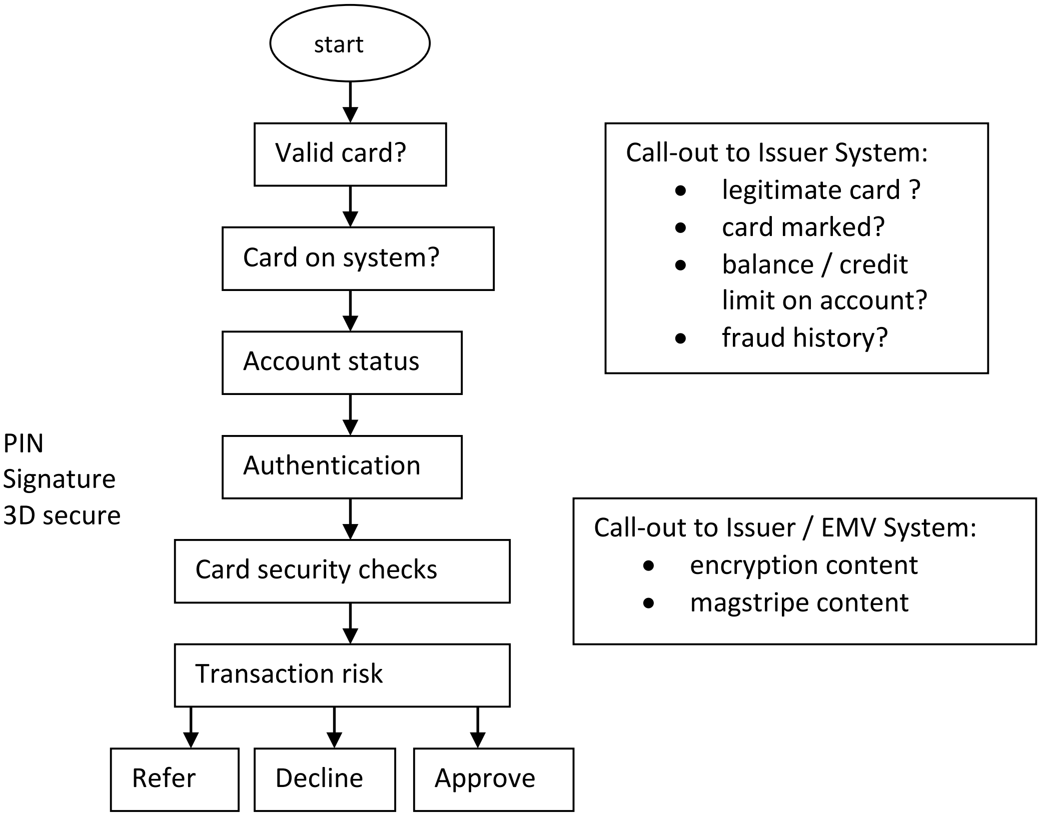

Figure 3: Card checks (part 1)

The majority of the decisions indicated in figure 3 (and the management of call-outs to other systems) are automated. The majority of these systems are independent and do not talk to each other. Consequently, any analyst is likely to see only a partial view of the transaction. This means that a skill of the analyst lies in piecing together fragments to which they have access in order to determine whether or not any part of the transaction could be deemed suspicious.

4.3 Call Centre Analyst's - Checking transaction risk

When analysts in a Call Centre respond to a 'refer' or 'decline' alert, they need to review transaction information, e.g., using systems such as T-Sys or First Data (figure 4). They might also have access to the account information which allows them to review the activity of the cardholder and the use of the account.

In terms of responding to alerts, most systems operate a Priority Mode, with the most significant needing immediate response. For Call Centre analysts, the role is customer-facing and managed through conversation with the customer. This would require key information to be available on the UI to the analyst as they manage the conversation. As noted previously, the conversation typically follows a script. However, the analyst will also seek to ask questions which are not directly related to the transaction (such as the weather in the cardholder's location or the type of purchase made). The aim is to both put the honest cardholder at ease and to catch the dishonest card user off guard.

At the end of the working day, transactions are batch-processed and the results of the batch processing can be passed on to Case Analysts who might focus on the behaviour of individual cardholders.

4.4 Fraud Analysts - Analysing Fraud Patterns

While one response to alerts is to contact the cardholder, another response is to explore trends and patterns in fraudulent activity. In real-time activity, the Supervisor could monitor the transactions that are being handled by the Call Centre analysts to look for possible trends and issues. From a batch of processed decisions, a Fraud Analyst could explore patterns and trends in the data. D7.2 showed that analysts tended to approach this activity in different ways but would seek to sort a set of transactions (by date or by risk score) and then search for common features. These features could be the cardholder name, the merchant ID, or the country of billing appearing in several of the suspicious transactions. Alternatively, discrepancies between cardholder address, country of billing and country of shipping could be deemed of interest. In this instance, the UI should allow the analyst to define and explore patterns in the data.

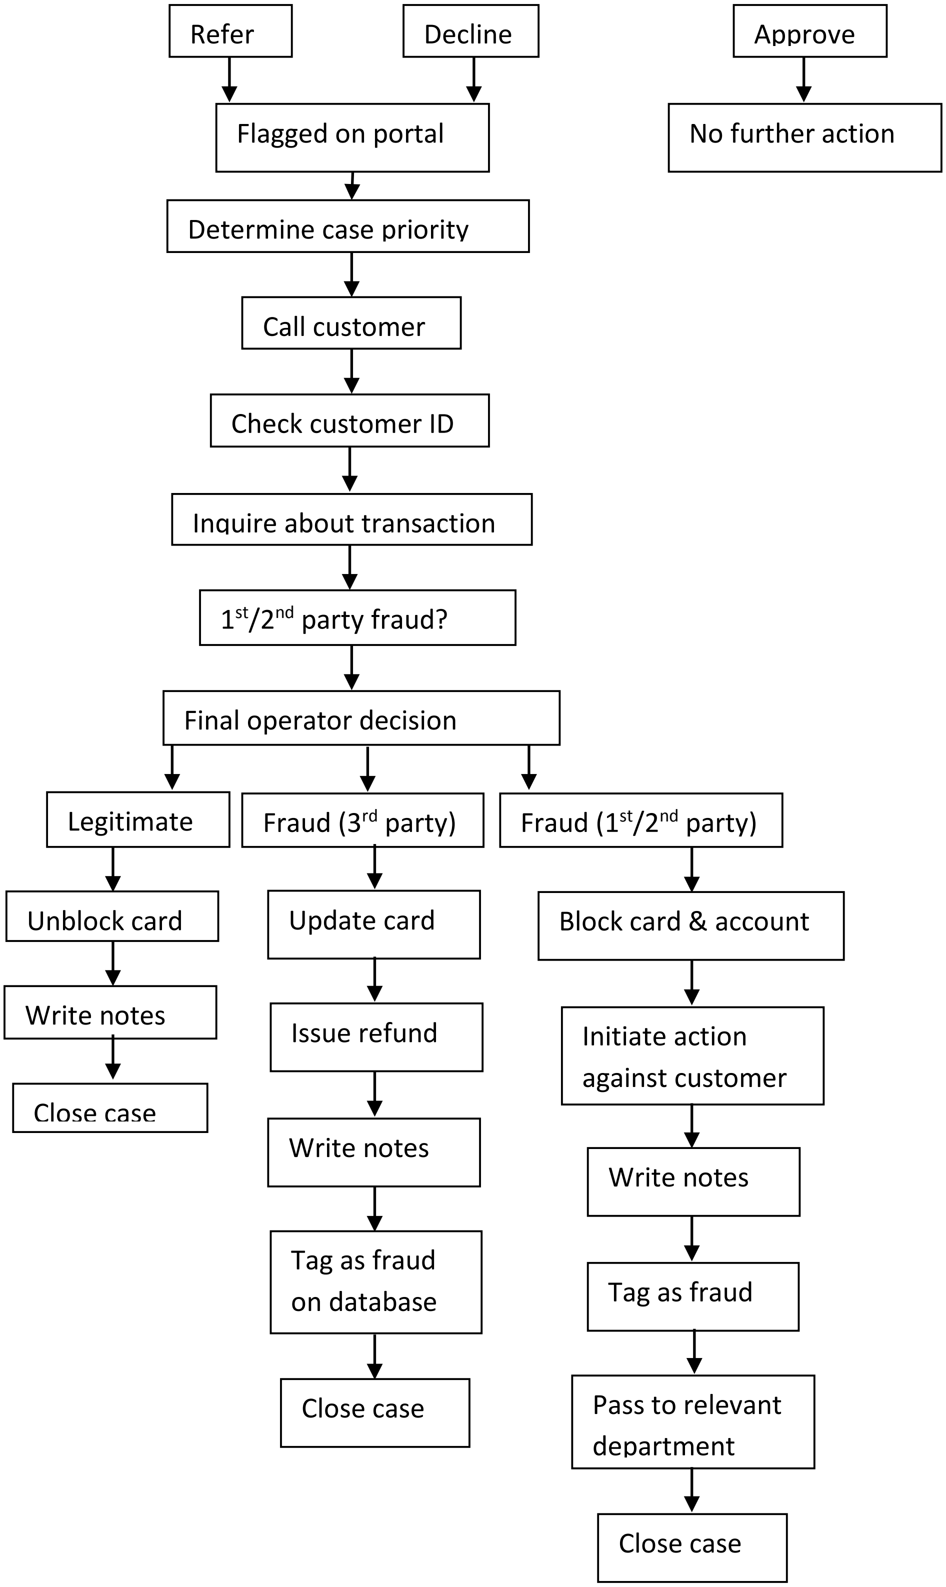

Figure 4: Card checks (part 2)

4.5 Merchant Fraud Analysis Workflow

In this section, the workflow for merchant fraud analysis is considered. The workflow here is discussed can generally be applied to most of the eCommerce platforms. The description is based on a combination of reports from Sift Science and experiences at Feedzai. Sift Science is a fraud detection company focused on the digital business, with a workflow very similar to that employed by Feedzai. The two companies describe themselves as follows:

"Sift Science fights fraud with large-scale machine learning. Machine learning lets a computer program recognize patterns of fraudulent behavior based on past examples. "[1]

"Feedzai believes every business can unlock the power of big data and machine learning. We deliver enterprise software to make management of risk and fraud better."[2]

eCommerce usually differs on normal transaction processing by the fact that real-time responses are not usually required. If it involves the shipment of any goods, the customers will not be expecting to receive the orders in the next minute. This gives more time for the fraud analysts on the merchant side to carefully analyze all the suspicious transactions and avoid sending merchandise to fraudsters. The software available to merchants who perform merchant-side fraud detection includes for example Feedzai Risk & Fraud, Sift Science Anti-Fraud Ring and Kount Central. These software solutions perform fraud analysis based on individually developed machine learning algorithms, handle merchant transactions and provide a user interface which allows employees to examine each incoming order for fraud.

Associated with the software are databases that hold the merchant's record of incoming transactions, associated attributes and metadata required by the company supplying the software. A company-based fraud analyst hence has access to new incoming transactions with associated customer details as well as the history of past transactions and customer information. At present, merchant-side fraud detection systems are being developed and optimized in order to facilitate analysis of the transactions and helping concluding the whole shipment process, with an emphasis on lowering the risk of a chargeback against the merchant or declining a genuine order (see figure 5).

[1] https://www.linkedin.com/company/sift-science [last accessed 25/06/2015]
[2] https://www.feedzai.com/ [last accessed 25/06/2015]

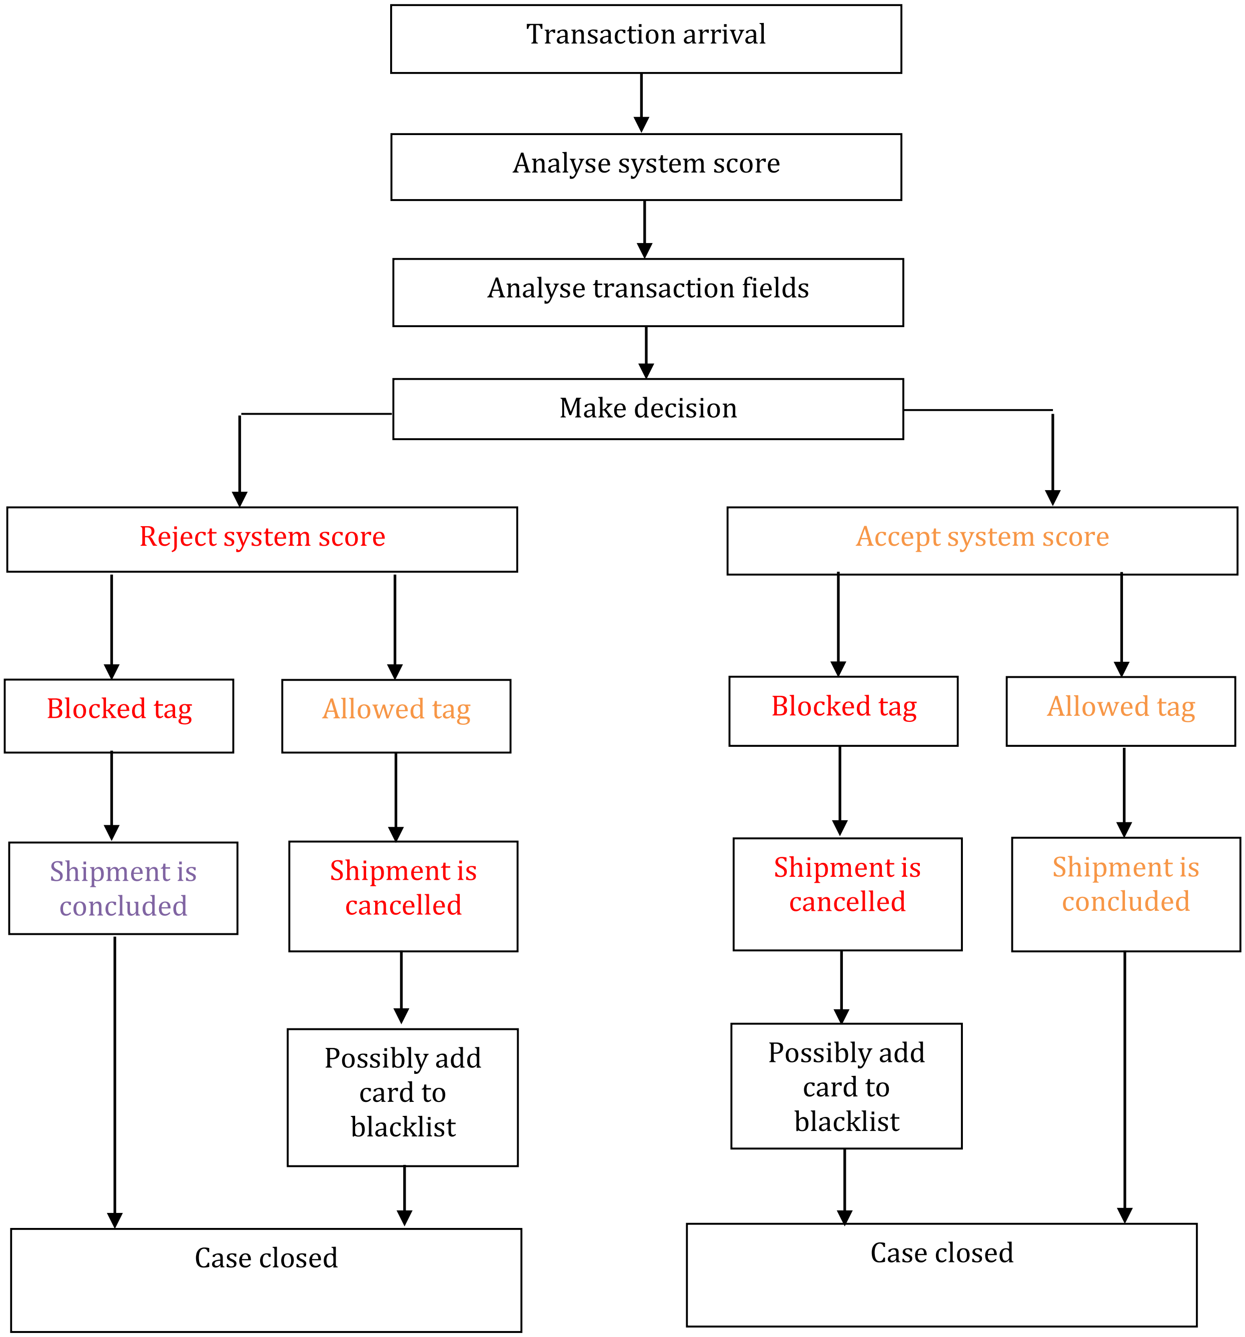

Figure 5: Flowchart for decision process at merchant level

Current challenges include improve the accuracy of the machine learning models usually used in these kind of systems, as well as the reduction of false positives and alerts. Full automation of the entire process is also one of the biggest challenges for the industry. At present, there is no interchange between merchant-side fraud detection and bank-side fraud detection outcomes, except for those cases where banks detect fraud and issue a chargeback to the merchant.

4.5.1 Workflow overview

Figure 6 illustrates the decision path that is taken from the moment a new order is made until it is shipped. Upon the arrival of a new request, all the information is channeled into the system, which after running its internal models, will attach a score to the transaction.

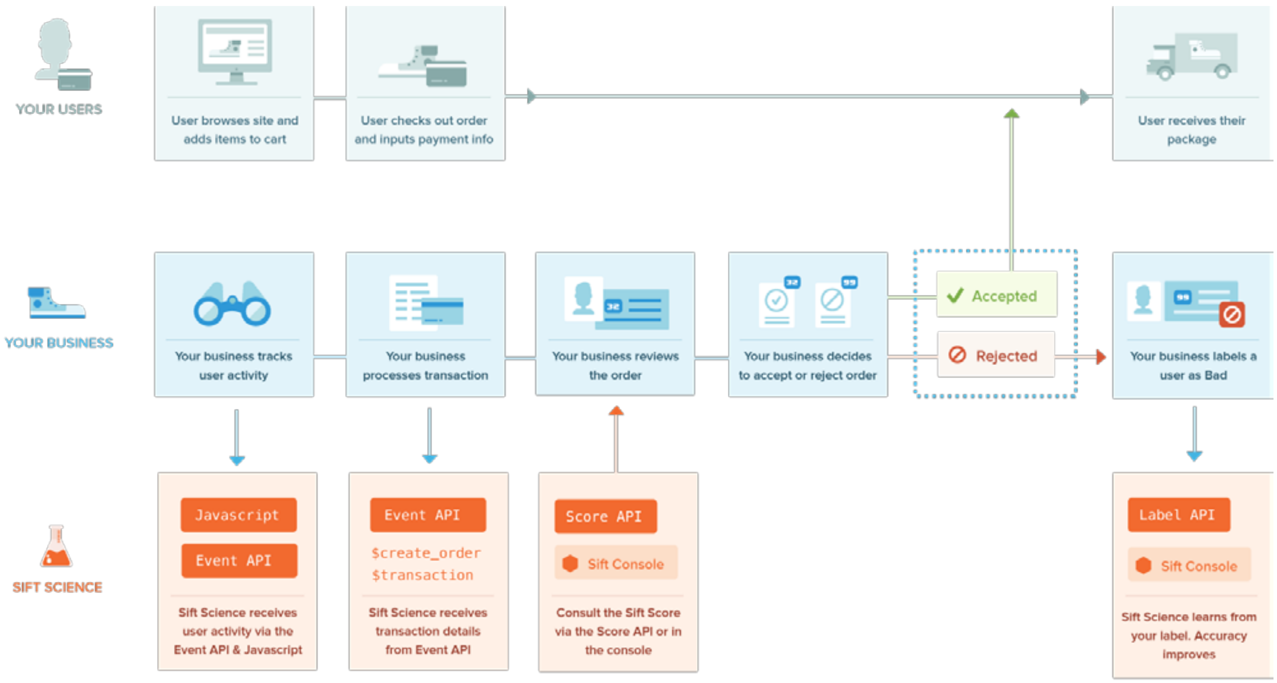
YOUR USERS
User browses site and adds items to cart
User checks out order and inputs payment info
User receives their package
YOUR BUSINESS
Your business tracks user activity
Your business processes transaction
Your business reviews the order
Your business decides to accept or reject order
Accepted
Rejected
Your business labels a user as Bad
SIFT SCIENCE
Javascript
Event API
Sift Science receives user activity via the Event API & Javascript
Event API
$create_order
$transaction
Sift Science receives transaction details from Event API
Score API
Sift Console
Consult the Sift Score via the Score API or in the console
Label API
Sift Console
Sift Science learns from your label. Accuracy improves

Figure 6: Fraud detection on merchant side[3]

The score and the transaction details are then carefully analyzed by the merchant fraud team, which will have the final word about whether the transaction should be shipped or not. Although the eCommerce solution score gives the merchant some guidance about whether the transaction is fraudulent or not, there is total freedom concerning the final decision. Even if the models score the transaction as highly suspicious, it can still be shipped if the merchant thinks accordingly. The same might happen in the opposite case, where the merchant blocks a supposedly genuine transaction.

[3] https://siftscience.com/resources/tutorials/integration-guide [last accessed 25/06/2015]

Guidance to the interpretation of scores is typically given by the company executives. Further, scores for the same transaction may vary between companies that use the same fraud detection software, as individual risk factors, weighting and other parameters can be custom set.

The definition of thresholds depends from merchant to merchant. Usually, the fraud detection company provides standard values (such as in a score between 0 and 100, all the transactions above 50 are tagged as potentially fraudulent), but it can also happen the case where the software company makes a previous study on the historical data provided by the merchant and can personalize these thresholds to match the best results on the datasets. Tags correspond to the decision that is made by the models, depending on the defined threshold. It means there is usually only two tags, one for values above the threshold ("blocked" transactions) and another for the values under that same threshold ("allowed" transactions). Although, merchants may want to add more tags to the system, possibly if there is more than a single threshold.

Note that the size of the merchant fraud team can vary from merchant to merchant, depending on its size. For the smaller merchants, typically it is the owner of the business that makes the decision alone, while for bigger companies, the process might even be automated. In case of a fully automated system, the orders are simply denied or accepted based on the score, as machine learning has already internally analyzed the relevant fields on the data before creating the score.

Once the transaction is labeled, the shipment order is sent or denied. The systems are usually capable of receiving *post-mortem* feedback. This means that if an order was shipped and it turned out to be fraudulent, the merchant can always give some feedback to the system for its own future improvement, as they are usually capable of online learning. For feedback, the system usually allow access to the handle transactions on the user interface, giving an option to confirm if the decision made was accordingly or not to that specific case. If nothing is altered, the system assumes that the correct decision was taken.

4.5.2 Analysis

Now that an overview to the processing workflow was made, let us pay more attention to each step performed.

1. Upon the arrival of an event, no human processing has been made yet. The data goes through the models within the fraud transaction system and a score is outputted. Depending on the thresholds defined for the application, a "blocked" or "allowed" tags are added to this transaction in particular.
2. All the transactions are then in a hold state, waiting to be accepted or declined. Depending on the number of incoming transactions and the number of analysis team on the merchant side, several scenarios might take place:
    a. All the transactions are analyzed.
    b. Only the "blocked" transactions go through the fraud analysts.
    c. Only the "blocked" transactions with lowest scores go through the fraud analysts. For the rest of the "blocked" transactions, they are declined.
    d. Only the "blocked" transactions with lowest scores and the "allowed" transactions with highest score go through the fraud analysts. For the rest of the "blocked"

transactions, they are declined. For the rest of the “allowed” transactions, they are accepted.

3. Independently of the previous scenario, all the transactions that go through the fraud analysts are usually evaluated looking to the fields in a certain order. Although it can naturally differ from analyst to analyst, the order of verification follows these steps:
    a. First of all, the score given by the fraud detection models is the most relevant aspect taken in consideration by an analyst. It usually defines if more fields should be verified or not.
    b. The analysis in the digital world is naturally different than for card present transactions. Fraud analysis in these cases usually turns into the analysis of the IP addresses of the computer used to make the transaction, the value of the merchandise and the shipment and billing address. After looking to the score, fraud analysts give preference over the analysis of the IP addresses associated with this customer. Associated with the IP addresses, most of the systems should be able to provide some data enrichment information, such as the use of Internet proxies and Tor networks[4]. These may be a strong evidence that the use is trying to hide his or her true identity and this can be a sign of potential fraud. IP addresses of certain countries are also more prone to fraud and this is also taken into account.
    c. While the IP addresses gives an idea of the geo location from where the transaction is taking place, the shipment and billing addresses are also quite often taken into account during decision process. The addresses can not only be associated with more risky regions, but non matching addresses may also indicate fraud. Imagine that a card is stolen. If the shipment and the billing addresses are different, in this case the billing address would be the one from the stolen card, while the shipment address would belong to the person committing fraud.
    d. The value of the purchase is also very relevant in the final decision. The amount may not only indicate if the transaction is fraudulent or not, but also for higher amounts, the merchants may be more reluctant is accepting risky transactions. Usually, the chargebacks of a fraudulent transaction will end up on the merchant bank account[5] and therefore, in order to protect their interests, merchant will most likely reject high volume transactions that are not clearly genuine. The last metrics to be analyzed are the ones that are also most commonly shared with the card present scenarios, namely the card information such as card number and cardholder name, or the details provided by the user, such as the e-mail or phone.
    These type of information can usually be crossed with static rules and information. For example, merchants may look if the user name matches with the one present in the card; if the card is associated to any black list; if the e-mail address was suspiciously formed or if the domain is marked as potentially spamming mail host and finally, if the phone number matches the country associated with this user in particular.
4. Once all these components are taken in consideration, the fraud analyst must take a decision. If the transaction is accepted, the shipment order is concluded and the customer should expect the package within a few days. If the transaction is blocked, the shipment must be cancelled and this step is usually associated with a customer notification about the reasons of the cancelling[6].

[4] https://www.torproject.org/ [last accessed 25/06/2015]

[5] https://www.wepay.com/api/payments-101/payments-fraud-and-loss [last accessed 25/06/2015]

[6] https://www.feedzai.com/developers/rest-api/ [last accessed 25/06/2015]

5. Defining UI content requirements

5.1 Introduction

The automatic checks performed for each credit card transaction generate a large number of attributes that become relevant for those transactions that are flagged and followed up. As described in Sections 3 and 4, checks are performed on separate sub-systems. Each of these systems has rules and generates data associated with the check. In order for a human to retrace details of a flagged transaction, in theory all features of all systems could be made available. However, practically, the displayed information will have to fit the purpose of the user. Based on discussion with Subject Matter Experts (detailed in D7.2), the displayed information will differ between call centre agents, supervisors, case analysts, fraud analysts and will also differ from the data used by merchant-side staff checking incoming transactions for fraud. In this section, we have mapped operator activities to individual features generated by the automated systems; these features will need to be integrated into a future UI that is fit for the market and that would push the boundaries of current systems.

5.2 Mapping of operator activities to required UI content

As outlined in sections 3 and 4, the fraud check and management workflow consists of three stages: firstly, transactions are automatically scanned for risk and abnormalities by computerized systems. Secondly, any flagged transactions are followed up by a human call center operative (or sometimes another automated system) to establish the legitimacy and background of the suspicious transaction. Thirdly, all transactions confirmed as fraud are examined by fraud analysts to for example establish emerging patterns. This activity goes hand in hand with updating of the machine learning algorithms / fraud models that flagged transactions in the first place.

In the tables 1, 2 and 3, the stages of the process and associated activities are mapped to features that could be shown in a UI to support operator activity if a human has to back-trace a transaction and establish whether it might be fraudulent. The mapping was performed based on discussion with domain experts at UK Cards Association and FICO as well as study of the relevant literature.

Table 1: UI content to support operator activities when following up on a flagged transaction based on features associated with automatic transaction and card checks.

Stage in fraud detection	Activity details	Information usage during fraud processing	UI information content
Automatic card check	Computer. Check of card fundamentals such as card number	Call Centre Agent: Past behaviour associated with card number; abnormalities in start/expiry date (e.g. card used before start date or card used close to expiry date) that explains a false positive or potentially true fraud; customer scoring in context of wider customer risk assessment etc. Fraud Analyst: Patterns associated with specific account number(s) / types / issuing banks; patterns associated with card usage close to a certain date; systematic exploitation of risk weighting settings associated with multiple cards etc.	**Basic card details** - Card number - Start and expiry date - Associated risk weighting by banks
Automatic account check on sub-system 1 [card issuer]	Computer. Call-out to check account details / status	Call Centre Agent: Any abnormalities associated with account status which the customer can be questioned about or which can be used to confirm that the account status is normal. Explanation to customer why payment history may confound purchases and advice on resolving debt issues etc. Fraud Analyst: Patterns of certain demographic exploiting credit card applications etc.	**Account status** - External status: card marked lost/stolen; previous fraud on account; account holder bankrupt / delinquent / insolvent etc. - Internal status: customer payment history; this is often also used in order to set the risk thresholds for checking the transaction

Stage in fraud detection	**Activity details**	**Information usage during fraud processing**	**UI information content**
Automatic authentication check	Computer. Check of customer authentication by card owner	Call Centre Agent: Check whether and how the customer authenticated the transaction and potentially question him/her about it; this is an important feature when exploring the possibility of 1st / 2nd party fraud; check with customer whether passwords are kept safe etc. Fraud Analyst: Patterns associated with hacked security numbers, leading for example to crime groups able to run through multiple PINs until one succeeds; emergence of trends in circumventing / exploiting security systems; check whether certain types of passwords frequently get hacked (current ongoing research into memorable but safe passwords)	**Authentication details** - Commonly PIN entered/not entered, result of authentication and history of falsely entered PINs - A signature may have been logged as 'present', especially if the card does not have PIN - Potentially further authentication outcomes via e.g. VISA verify or MasterCard Secure (3DSecure)
Automatic card security check on sub-system 2 [chip / EMV system]	Computer. Call-out to check card security	Call Centre Agent: information regarding invalid chip security information to guide discussion with customer Fraud Analyst: specialist analysis of emerging trends in hacking chip mechanics, systematic counterfeiting, security gaps, guidance for updated technology etc.	**Chip checks (EMV)** - Security code details - Cryptograms details - Large number of features / codes embedded in chip that are not necessarily used but available

Stage in fraud detection	Activity details	Information usage during fraud processing	UI information content
Automatic transaction risk scoring on sub-system 3 [e.g. FICO]	Computer. Call-out to calculate a risk score for the transaction	Call Centre Agent: Sense-making of the flagged transaction: why was the transaction flagged? Is there an immediate reason obvious why it is likely a false positive? What is the customer's spending behaviour and how does it relate to the flagged transaction? Are there past transactions that were similar but did not pass the threshold? Is this a high-risk customer? The information has to allow the call centre agent to explain a blocked card to the customer so that customer loyalty is not jeopardised; guidance of interaction with customer via script Fraud Analyst: new trends in systematic fraud that has not yet been picked up by machine learning algorithms, or very specific fraud patterns; patterns associated with specific merchants, ATMs, global regions etc.; examination of fraud that may have been misinterpreted by call centre operative; trends for fraud occurring just sub-threshold; emerging fraudulent behaviour not seen before etc.	**Transaction attributes** - Access to all features associated with transaction - Access to past buying behaviour - Both basic features and aggregated features; for features used in the literature, please see D7.2 - Exact number of features varies between issuers, up to around 70 to 80 fields [Krivko 2010; Whitrow et al. 2009]. After calculation of aggregated metrics, the number of features can be even larger [Whitrow, Hand, Juszczak, Weston and Adams 2009] and may go into the 100s - Display of relationship between features - Display of features flagged by computational fraud detection - Display of thresholds used by computational fraud detection - Explanation of reason for computation fraud detection to flag the transaction - Details on the 'normal' aspects of the transactions in similar style as described for fraud pointers - May have to allow for 'consortium level pooling', respectively intelligence across banks

Table 2: UI content to support operator activities when following up on a flagged transaction based on features associated with customer engagement.

Stage in fraud detection	**Activity details**	**Information usage during fraud processing**	**UI information content**
a) Transaction flagged, passed on directly to frontline staff	Human. - Call centre agent calls card holder to verify card holder ID and legitimacy of card transaction	Call Centre Agent: Take notes of conversation, possibly via interaction with drop-down menus or check boxes; check previous conversations with client logged during conversations in the past; make sure to follow the script so that the examination process is repeatable Fraud Analyst: Patterns of call centre agent exploitation; potential amendments to scripts to improve process	**Log of call centre agent reasoning** - Script which call centre staff have to follow - Log of conversation - Log of thought process / reasoning - Log of actions taken
b.1) Transaction flagged, passed on to automated system	Computer. - Automatic SMS / phone call	Call Centre Agent: check original outcome of automated check Fraud Analyst: check time/date of automatic check and outcome; check whether a specific system associated with a specific issuer is more prone to risk	**Log of automatic contact** - Log all attributes associated with automatic check
b.2) Transaction declared fraudulent by owner after contacted by automated system	Human. - Contact customer to establish details	Call Centre Agent: see actions in a) Fraud Analyst: Patterns of confirmed fraudulent behaviour similar to a)	**Log of customer service rep contact** - Script - Log of conversation - Log of thought process / reasoning - Log of actions taken

Stage in fraud detection	**Activity details**	**Information usage during fraud processing**	**UI information content**
c) Transaction referred	Human. - No further transactions possible or account 'watched' until customer was spoken to	Call Centre Agent: using all details past and present, establish whether action performed with card was unusual Fraud Analyst: emerging trends of unseen behaviour resulting in difficult to classify cases	**Log of customer service rep contact** - Script - Log of conversation - Log of thought process / reasoning - Log of actions taken
d) Transaction assessed as legitimate	Human or computer. - Log of check following assessment as legitimate	Call Centre Agent: as in a), with the outcome that the transaction was legitimate Fraud Analyst: check whether there are fraud patterns for which some specimens go unnoticed by call centre agent or computerised verification; check whether there are clusters of transactions tagged as legitimate associated with specific customers, merchants, ATMs, global regions etc	**Log of incident** - Script - Log of conversation - Log of thought process / reasoning - Log of actions taken

Table 3: UI content to support fraud analyst activities using transactions that were confirmed as fraudulent.

Stage in fraud detection	Activity details	Information usage during fraud processing	UI information content
Updating of machine learning algorithms / fraud models	Human.	Call Centre Agent: N/A Fraud Analyst: Communicate patterns to technical staff for implementation in machine learning algorithms; likely to communicate patterns to bank directors, colleagues and other stakeholders; information about changes in implemented machine learning algorithm and associated pattern of fraudulent behaviour	**Updates** - Date of changes and nature of implemented changes within fraud model - Notifications regarding emerging / global patterns or 'hot' patterns, likely region-specific
Continuous updating of fraud model via tagged transactions	Computer.	Call Centre Agent: N/A Fraud Analyst: check for emerging fraud patterns picked up by machine learning; frequency of occurrence of new fraud patterns; location-specific emergence of 'updated' fraud patterns / fraudster learning	**Updates** - Log of updates to fraud model

6. Performance Metrics and Baseline Determination

6.1 Introduction

As specified in section 3, discussions with Subject Matter Experts have led to at least four key roles being identified, all bearing the title of fraud analyst. In order to easily distinguish between these three roles they have been given the following names: call centre analyst (investigates individual transactions), supervisor (oversees groups of call centre operators), case analyst (investigates links between individual transactions) and fraud analyst (searches for new fraud patterns). Their responsibilities are very different and so will be the support they require in order to make correct decisions. Call centre analysts usually deal with a large number of transactions (i.e., in the region of 200 per day) flagged by an automated system. For each of flagged transaction, the analyst will need to make a decision whether to block an account from trading or not. This decision can have a great impact on the financial institution they are working for as it can either increase financial losses or lead to a negative outlook on the company. Therefore, not only data that is presented to the operators is crucial but also the time necessary to reach a decision. Supervisors, apart from being responsible for managing teams of call centre operators (low-level analysts), they might also have to deal with more problematic cases, referrals, and refining scripts used by the call centre operators. The nature of their job is higher-level and less time-critical. The Case analyst will perform similar activity to the Supervisor, in terms of analysis, but work from batches of transactions. The highest-level analyst, referred to simply as fraud analyst, is responsible for finding new fraud patterns that can feed into the automated scoring systems.

6.2 Quantifying Performance

Because of the high security nature of fraud investigation jobs, not only getting real data but also conducting studies in the work environment (of the task ecology – such as that reported in D8.3 for the traffic management use case) proves to be problematic. However, after discussions with key players in the area of fraud investigation and prevention (see D7.2) we have gained a better understanding of what the tasks of the different roles identified might be. Furthermore, based on the review of fraud types presented in 5.1, datasets containing fraud patterns can be generated. Therefore, tasks using realistic synthetic data can be simulated and tested with both subject matter experts and non-experts (due to very low availability of and restricted access to experts).

As previously mentioned, different roles have been identified to bear the name of fraud analyst. They can be characterised in terms of data and time requirements as shown in table 4.

Table 4: Performance metrics for analysts

Role/Requirement	Case Volume	Data Volume / Case	Data Resolution	Time Criticality	Time Resolution
Call Centre Analyst	High	Low	Account Level	Very High	Very High
Supervisor	High	Moderate	Customer Level	Moderate	Low/High
Case Analyst	Moderate	High	Account / Institution Level	Low / Moderate	Low / High
Fraud Analyst	Low	High	Customer/Institution Level	Low	Low/High

Due to the fact that the requirements of each role are different, the user interfaces for each of the different roles will be different as they will need to support different goals.

6.2.1 Call Centre Analysts

In the case of call centre analysts, the UI will need to provide clear and concise information about the context of a transaction so as to help reach a decision as soon as possible. This could include an explanation of the automated scoring system (reason and/or reasoning), account history, date, time and location of the transaction investigated and of the most recent transactions. Call centre operators usually deal with around 200 transactions per day which means they have less than 3 minutes to make a decision per transaction, time during which they might have to contact the account holder. A number of metrics can be used in order to quantify the performance of operators while using different UIs:

- Information search time
- Number of correct/incorrect decisions
- Decision time
- Decision confidence
- Understandability of the automated score

In terms of determining the operator performance baseline with current state of the art UIs, the synthetic data generated can be fed into the current FeedZai UI and an experimental task can be simulated around this role. The results of this experiment will be characterised in terms of the aforementioned metrics. A subsequent experiment will be run with the same data but, this time, with the UI developed as part of the SPEEDD project and the results characterised in terms of the same metrics. The results of the two experiments can be then compared in order to determine changes in operator performance.

6.2.2 Supervisors and Case Analysts

While a performance baseline is relatively straightforward to define in the case of low-level fraud analysts (call centre analysts), for the higher-level roles (supervisor and case analyst) it is not as clear how such a baseline can be determined. This is due to the fact that the latter roles do not have a consistent description across institutions and there are no standard state-of-the-art systems or user interfaces to support these roles. While there have been some attempts to design UIs for these higher-level roles, to the best of our knowledge, none of them have been adopted and implemented commercially.

Supervisors have the responsibility of looking over groups of call centre operators, but on top of this they might have to deal with the more difficult cases (higher potential losses, high-value customers, inconclusive cases or referrals) and with refining the scripts used by the lower-level analysts. Considering the high variability in their job description some of their requirements might be: being able to see snapshots of what the call centre operators are dealing with, accessing statistics regarding the decisions made by the operators and being able to receive referrals from them. They would probably investigate a much smaller number of cases and spend more time on them. Also, compared to the call centre operators they would require access to a higher volume of data (often from different sources – even external from the institution), at a customisable resolution. Under these considerations, a number of metrics can be defined as follows:

- Time to establish communication with call centre operators
- Ease of accessibility to multiple information sources
- Number of correct/incorrect decisions
- Decision confidence

6.2.3 Fraud Analysts

The role of Fraud Analysts shares some similarities with Supervisor roles: it deals with multiple data sources at a selectable resolution. However, instead of flagging individual transactions and blocking accounts, their main goal find new fraud patterns in datasets of already closed cases. Their job is more explorative in nature, implying a less strict time constraint. Due to the fact that this role deals with the most sensitive data and generates knowledge that is formalised and used as inputs to automated scoring systems, it is also the one that is hardest to get access to. Nevertheless, relying on what we have found out about this role we can identify the following metrics as performance indicators:

- Ease of accessibility to multiple information sources
- Number of correct/incorrect patterns identified
- Time to pattern identification
- Understandability of the relationship between the automated score and operator decision