



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

Project **FP6-619435/SPEEDD**
Deliverable **D7.2**
Distribution **Public**



<http://speedd-project.eu>

Initial Evaluation Report

Ivo Correira
(FeedZai)

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

Status: FINAL

May 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	D7.2
Deliverable Title	Initial Evaluation Report
Contractual date of delivery	M13 (March 2015)
Actual date of delivery	May 2015
Relevant Task(s)	WP7/Tasks 7.2
Partner Responsible	FeedZai
Other contributors	UoB
Number of pages	57
Author(s)	I. Correira, C. Baber, S. Starke and N. Morar
Internal Reviewers	A. Skarlatidis, A. Artikis
Status & version	Final
Keywords	Evaluation, User Interface Design, Human Factors, Eye Tracking

Contents

1. Executive Summary	6
Document Structure	6
1. Introduction	7
1.1 History of the Document	7
1.2 Purpose and Scope of Document	7
1.3 Relationship with Other Documents	7
1.4 Sources of Information	7
2. Requirements Revisited	8
2.1 Introduction	8
2.2 FeedZai.....	9
2.2.1 Client and card history	9
2.2.2 Tag fraudulent transactions.....	9
2.2.3 Locality information.....	10
2.2.4 Common patterns and trends for analysis	10
2.2.5 Communicate decisions to clients and other entities	11
2.2.6 Communication between analysts	11
2.2.7 Summary of FeedZai consultation	12
2.3 UK Cards Association (David Baker).....	12
2.3.1 Background	12
2.3.2 Background to transactions and automatic checks.....	13
2.3.3 The role of humans in the credit card fraud management process	14
2.3.4 Conclusions from discussion with UK Cards Association.....	15
2.4 FICO (Brian Kinch and Scott Zoldi)	16
2.4.1 Background	16
2.4.2 Transaction scoring.....	17
2.4.3 The role of humans in the credit card fraud management process	17
2.4.4 History of credit card fraud prevention	18
2.4.5 Information used and user Interface content	20
2.4.6 Conclusions from discussion with FICO	20
3. Comparing User Interface Designs for Fraud Analysis.....	21
3.1 Entering Credit Card Details for Online Purchase.....	21

3.2	Transactions on an individual account	22
3.3	Dashboards	25
3.4	Analysing Patterns in Collections of Transactions	27
3.5	Other visualisations	30
3.6	Graphical design in SPEEDD	31
4.	Formative Evaluation and Transition from first to second prototype	33
4.1	Introduction	33
4.2	Usability of Evaluation of the Two Versions	37
4.3	Additional Comments	38
5.	Understanding the Fraud Analysis Process	39
5.1	Introduction	39
5.2	Using Fraud Analysis Software: a case study conducted at FeedZai	40
5.2.1	Fraud analysis workflow: interviews following work through the 10 sample transactions 41	
5.2.2	Fraud analysis workflow: detailed workflow explained by dedicated fraud analyst.....	41
5.2.3	Gaze behavior while using the FeedZai UI.....	43
5.3	Laboratory Task to Investigate the Fraud Analysis Process.....	44
5.3.1	Background for the work	44
5.3.2	Experimental design	44
5.3.3	Participants	47
5.3.4	Preliminary results	47
5.3.5	Summary and outlook.....	51
5.4	Defining Baseline Performance Metrics	51
6.	References	52
7.	Appendix 1 – SUS questionnaires in English and Portuguese	53

Figures

Figure 1: Usage of cardholder historical behaviour during decision process.....	9
Figure 2: Example of a map showing with most fraud is happening.....	10
Figure 3: Different patterns for fraudulent and genuine transactions.....	11
Figure 4: UK Cards Analysis of Card Fraud Loss	13
Figure 5: History of fraud prevention measures summarised by FICO.....	19
Figure 6: Credit Card purchase on www.Amazon.com.....	21
Figure 7: Setting limits on card payments	22
Figure 8: Displaying transactions on an individual account	22
Figure 9: User Interface design for MIVA product.....	23
Figure 10: User Interface for Magento Commerce UI.....	24
Figure 11: User Interface provided by CyberSource for merchant-based fraud management.....	25
Figure 12: Dashboard for card transaction activity	26
Figure 13: Dashboard of NICE Actimize	26
Figure 14: Displaying patterns of card fraud in Europe.....	27
Figure 15: Displaying card skimming in the US.....	27
Figure 16: Displaying financial transaction data using Tableau.....	28
Figure 17: Combining sources of data	29
Figure 18: Tree map.....	29
Figure 19: Monitor by exception in high density visuals	30
Figure 20: Complex link analysis.....	31
Figure 21: WireVis: Visualization of Categorical, Time-Varying Data From Financial Transactions	32
Figure 22: Display of donation patterns	32
Figure 23: Evaluation of UI version #1 and #2 at FeedZai in Lisbon, Portugal.....	33
Figure 24: The first version of the user interface.	33
Figure 25: An overview of the SPEEDD user interface.....	34
Figure 26: Screen for historical data.....	35
Figure 27: Screen for transaction details.....	36
Figure 28: Screen for the details of a country.	36
Figure 29: SUMI evaluation of initial prototypes.....	37
Figure 30: Possible strategy for fraud analysis	39
Figure 31: Screenshot of FeedZai UI for fraud management.	40
Figure 32: Subset of stages within a workflow during examination of a suspicious transaction.....	42
Figure 33: Distribution of fixations across the monitor while working with the FeedZai UI.....	43
Figure 34: Simplified UI displaying nine information sources related to a credit card transaction.	45
Figure 35: Information showing fraud pattern #1, several small transactions on the day of the current transaction.....	45
Figure 36: Information showing fraud pattern #2, a single large transaction.....	46
Figure 37: Feedback given on each transaction evaluation.	47
Figure 38: Responses to individual transactions by four participants.....	48
Figure 39: Visual attention allocation to the different information sources across all trials.....	49

1. Executive Summary

This deliverable reports the evaluation methodology and presents the results of the initial user evaluation performed after the release of the SPEEDD prototype for the Credit Card Fraud Management Use Case. The deliverable includes examples of competitor user interfaces and explains how these can be taken as the state-of-the-art for design in this domain. The initial prototype sought to reflect these designs. In addition to evaluating the design of the User Interface, the report also presents initial work on defining performance metrics for fraud investigation. The goal is to use these metrics to evaluate the impact of the User Interface (and SPEEDD architecture) on decision making by fraud analysts.

Document Structure

The document is divided in two main parts. In the first, the User requirements presented in D5.1 are revisited and extended with additional requirements elicited from an informal review at FeedZai. This is complimented by the results of discussions with two organisations involved in the management of credit card fraud. This is followed by a review of the state-of-the-art in User Interface design for financial management and analysis software. Following this, evaluation of the initial prototype for the SPEEDD Credit Card Fraud Management use case is presented. The evaluation involves an informal review by FeedZai personnel followed by the application of the Software Usability Scale (SUS) (Brook, 1988). This scale was used in D8.3 and provides a simple but effective means of formative evaluation of the User Interfaces. The next section presents initial studies into decision making in fraud analysis. This begins with a step-by-step analysis of the ways in which FeedZai employees interact with their current system to evaluate (fictitious) fraud cases, and is followed by a study in which fraud data are presented in a controlled manner to explore how people search information sources. The final section presents a summary of baseline metrics which might be beneficial for subsequent evaluation activity.

1. Introduction

1.1 History of the Document

Version	Date	Author	Change Description
0.1	30/03/2015	Ivo Correia	First version of the document
0.2	26/03/2015	Chris Baber	Second version of the document
1	29/05/2015	Chris Baber	Final version (addressing review comments)

1.2 Purpose and Scope of Document

The purpose of this document is to report an initial evaluation of the SPEEDD prototype in the Credit Card Fraud Management use case. As part of this activity, requirements from the potential users of the prototype are explored together with evaluation of the current prototype. These requirements represent developments on the initial set of requirements for the Credit Card Fraud Management use case (D7.1) and the specification of the User Interface of the prototype (D5.1). In terms of evaluation, the aim is to show how the prototype should evolve and how it is going to be used by the fraud operators. The target audience of this document will be all parties involved in the implementation of the fraud use case.

1.3 Relationship with Other Documents

As noted in the previous section, this document is related to the deliverable 7.1 User Requirements and the deliverable D5.1 Design of User Interface for SPEEDD Prototype. It is also related to D8.3 Initial Evaluation Report of the Road Traffic Management use case (particularly in terms of evaluation methods).

1.4 Sources of Information

For the initial plans, FeedZai counted on the fraud detection organization Paywatch to provide access to their analysts, in order to gather feedback to be applied in the SPEEDD prototype. Being the only processor in Portugal and renowned world-wide SIBS-FPS (2011), their insight would be most valuable for the project. Unfortunately, the final agreement was not achieved and so, it is not expected that any help will come from Paywatch for the SPEEDD project. In order to keep the project going, work-arounds had to be found.

Information was gathered from the FeedZai personnel, who could provide useful insight due to their contact with several credit card fraud analysts, providing this way and indirectly, what is expected to find in a user interface built towards fraud management. FeedZai personnel has worked in several fraud detection projects, including work with Paywatch and the development of its own Fraud API, an e-commerce service available online. Meetings were also arranged for project partners to visit two organisations in the UK: the UK Cards Association and FICO.

2. Requirements Revisited

2.1 Introduction

One of the main objectives 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. Therefore, it is important to know what users value, what are the main features that they find fundamental in the current applications and what kind of features they are expecting to find in future releases of software.

While D7.1 outlined the main requirements for the SPEEDD System for Fraud Use, it did not consider user requirements. D5.1 outlined a set of requirements which were used for the initial design of the User Interface (UI) for Work Package 7. It was noted, in D5.1, that a key determinant in financial fraud involved assessing whether a pattern of behaviour was normal for a given cardholder or not, with the implication that an aspect of fraud detection was the definition of 'normal'. However, given the speed of response that credit card industry expects from analysis, it is unlikely that placing the human analyst 'in the loop' will provide sufficient benefits to make the delay acceptable. As quoted in D5.1, *"Automatic systems are essential since it is not always possible or easy for a human analyst to detect fraudulent patterns in transaction datasets, often characterized by a large number of samples, many dimensions and online updates."* (Dal Pozzolo et al., 2014).

In D5.1, an initial set of requirements were derived from discussion within the consortium and from the application of Cognitive Work Analysis. Requirements, to support analyst decision making included "explaining the results of the models in a human-friendly way", "reducing false alarms to reduce alert fatigue", and "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." These requirements were refined through literature review and initial discussions with financial institutions in the United Kingdom, Germany and Romania but there was no direct contact with credit card fraud management. The three institutions with which initial discussions were made are related to financial compliance, cheque-clearing and money laundry investigation. In this report, the initial requirement set is supplemented through discussion with organisations involved in credit card fraud.

2.2 FeedZai

This list addresses all the requirements that were collected from the analysts, creating a wish list for credit card fraud management use case. Not all the requirements are essentially for an analyst to perform his or her work, so this list was created by putting the items with most priority in the top, so decisions of what to keep and what to exclude can be made accordingly.

2.2.1 Client and card history

One of the most important aspects when trying to detect credit card fraud is to know if the observed behaviour is uncommon or normal from that cardholder (figure 1). For the same transaction, it can be normal for a certain cardholder, or strange for another. After all, people come from different backgrounds and naturally, it is expected that they use their cards differently as well. Consequently, the user interface should be able to present to the analyst the recent history of the card, showing information such as the number of transactions in the last month or what were the average amounts of transactions performed by that card.

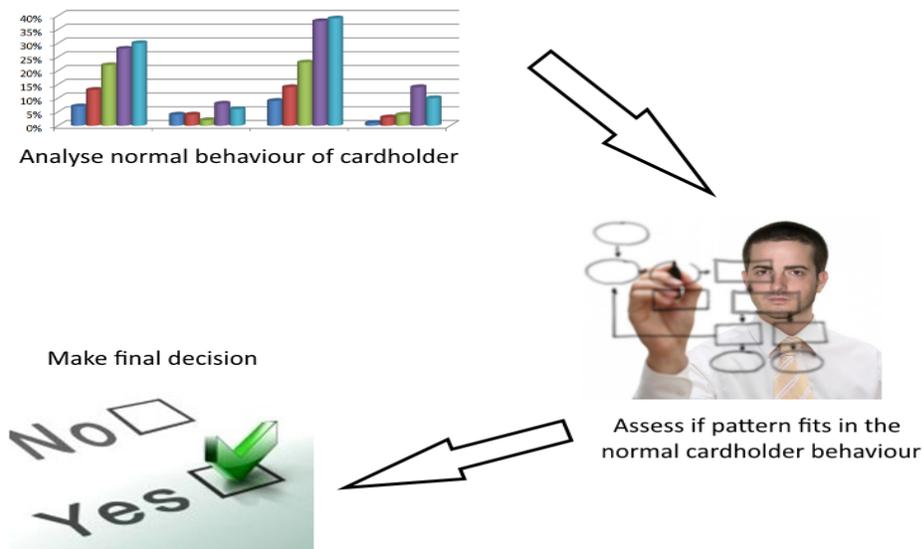


Figure 1: Usage of cardholder historical behaviour during decision process.

2.2.2 Tag fraudulent transactions

Although it may seem an obvious requirement, it cannot be omitted that analysts should be able to tag the evaluated transaction as either fraudulent or genuine. What happens after the analyst marks the transaction, is not really concerned to the user interface, but to the decision making module and to the fraud company itself. Still, the analysts should be able to tag the transaction they are observing and also, permit an undo operation, as mistakes may occur and it should be allowed for the analysts to correct their actions.

2.2.3 Locality information

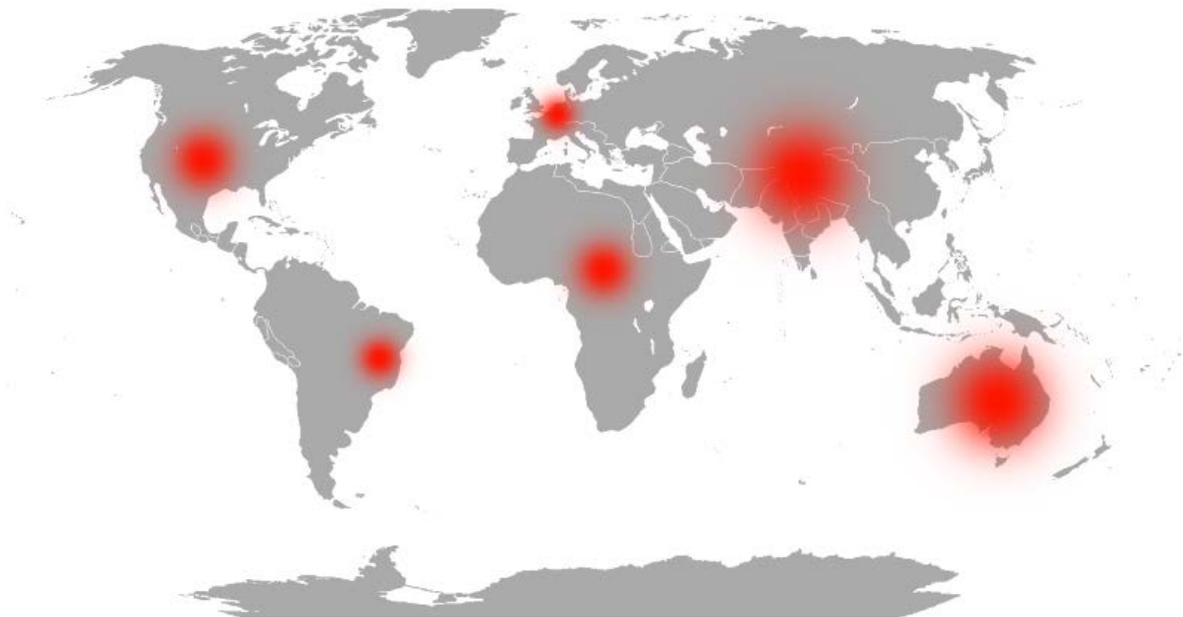


Figure 2: Example of a map showing with most fraud is happening.

Although it is not common in older management systems, the focus on locality information has been gaining relevance in the past few years. It is not only important that locality information is shown, but also how it is shown (figure 2). Locality information can for example, allow the analyst to see if the cardholder usually triggers transactions out of his home country, or if it common from him or her to be in a given part of the world.

As referred, it is important to note that the way this information is shown to the analyst is a factor of extreme importance. A list of the countries from the last ten transactions can be replaced by a map, with the points pinned on it. Visually, it will be easier for the analyst to process and correlate all the information. If using a map, it can also be possible to colour or contrast areas that are most visited by the cardholder. In a more general perspective, the areas where more fraud is occurring can also be shown in maps, and this information can be directly crossed with the location of the current transaction.

2.2.4 Common patterns and trends for analysis

Sometimes fraud analysts may not be aware of the latest fraud patterns. If the previous point was more related to specific cardholders and their behaviours, this point related to a more general overview of the market. Analysts find useful having information regarding the context. In this case, if the analyst knows that the pattern he or she is seeing has been marked as common recently, it can be a positive indicator that the transaction is effectively fraudulent. On the other hand, general

information such as the amount of fraud recorded for a specific country in the last year, can also indicate to the analysts if he or she should pay more attention to the shown transaction (figure 3).

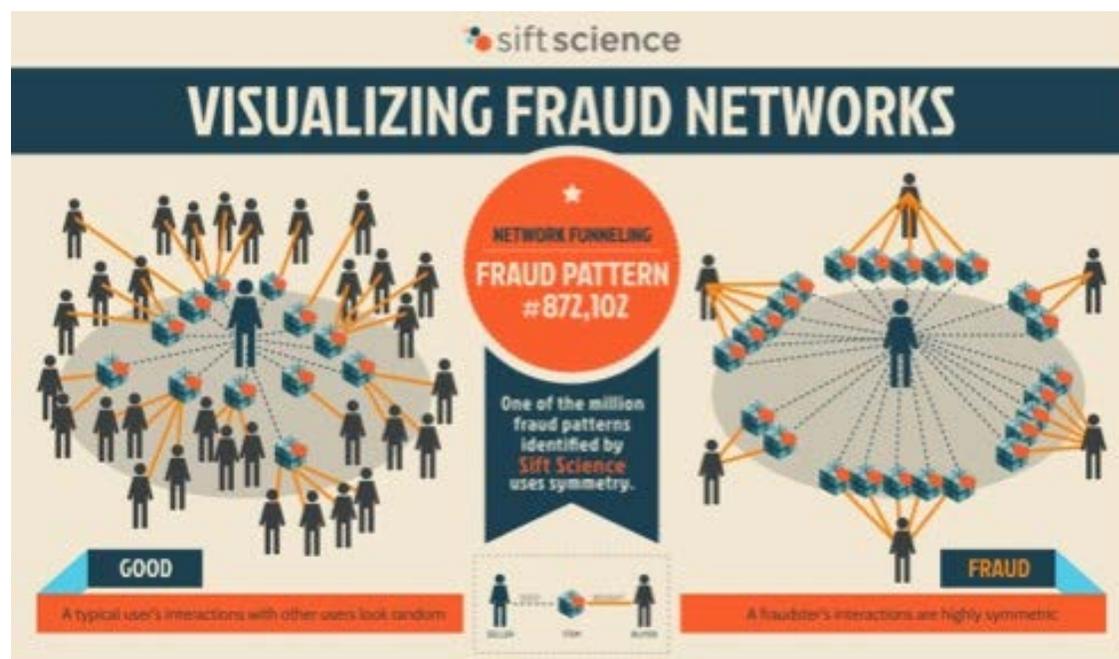


Figure 3: Different patterns for fraudulent and genuine transactions¹

2.2.5 Communicate decisions to clients and other entities

Finally, one aspect to consider, once the transaction has been tagged as genuine or fraudulent, is what are the further actions to follow. Although in the case of SPEEDD, this point may be more associated with the decision process rather than the user interface, it is not uncommon for analysts themselves being able to make reporting decisions over transactions. These kind of actions can refer to generate sample reports, which can be consumed both internally or externally. If actions can be taken by the analyst, then the user interface should also provide a way to associate to a given transaction, what subsequent action was completed. This point can have extra importance when auditing procedures are required within the fraud detection companies.

2.2.6 Communication between analysts

The process of credit card fraud involves a lot of reasoning and discussion, as different analysts might have different opinions about the same transaction. Therefore, in order to promote a more dynamic environment, analysts should be able to share their experiences. Also, it can be helpful if analysts are able to share transactions between each other. Sometimes, an analysts may find that he or she is not

¹ Accessed 26052015. <https://gigaom.com/2013/03/19/sift-science-says-it-can-sniff-out-cyber-fraud-before-it-gets-expensive/>

capable of correctly identifying a transaction that was assign to him or her. In these cases, it can be a positive point if the analysts are able to send back the transaction to a waiting queue or directly to a colleague. Although this is not a hard requirement, as analysts, in the last instance, are expected to be able to speak face-to-face, it can be a positive point to take in consideration when implementing the SPEEDD prototype.

2.2.7 Summary of FeedZai consultation

In this Section, we went through the list of requirements elaborated to address the requirements defined by the users consulted. We have seen six issues that should be considered, dividing them into the most important (client and card history, tag fraudulent transactions, locality information and common patterns and trends for analysis) and the secondary requirements (communicate decisions to clients other entities and communication between analysts).

2.3 UK Cards Association (David Baker)

2.3.1 Background

The UK Cards Association is the umbrella organisation overseeing security and development of any matters related to plastic card and other payment methods. The main aims are to make card security as good as possible in a trade-off between customer experience and security, especially online. Further, the aim is to standardise transactions so that the customer experience is similar independent of the location and issuer. David Baker is the Head of the Card Payment Innovations Unit.

*'The UK Cards Association is the trade body for the card payments industry in the UK, representing financial institutions which act as card issuers and acquirers. Members of the Association account for the vast majority of debit and credit cards issued in the UK - issuing in excess of 55 million credit cards and 95 million debit cards - and cover the whole of the payment card acquiring market. The Association promotes co-operation between industry participants in order to progress non-competitive matters of mutual interest; informs and engages with stakeholders to shape legal and regulatory developments; develops industry best practice; safeguards the integrity of the card payments industry by tackling card fraud; develops industry standards; and co-ordinates other industry-wide initiatives such as those aiming to deliver innovation [...].'*²

The UK Cards Association's work on fraud includes Financial Fraud Action UK, the Financial Fraud Bureau, a Dedicated Card and Payment Crime Unit, a Fraud Intelligence Sharing System and an Industry Hot Card File.

² http://www.theukcardsassociation.org.uk/aims_objectives/index.asp

2.3.2 Background to transactions and automatic checks

While the fraud rate has remained fairly stable, the total loss due to fraud has increased linearly with the rise in purchases using credit cards. Transactions have now shifted to be most frequent in the digital space (figure 4), and online transactions are now more frequent than face-to-face transactions during peak shopping periods and especially over the Christmas period. Accordingly it is important to consider whether purchases are made with the card present or card not present; more than half of fraud is committed with the card not being present (remote purchase). There are also large differences between countries: in the UK, more than 90% of transactions are online authorised, whereas in France this is only around 30% (a much higher proportion of transactions is authorised by chip there). In contrast, the US has not yet adopted chip+PIN yet. The emerging new technologies for payments are phone payments, such as Apple Pay. The outlook is that payment options on the phone will match those currently available for plastic cards.

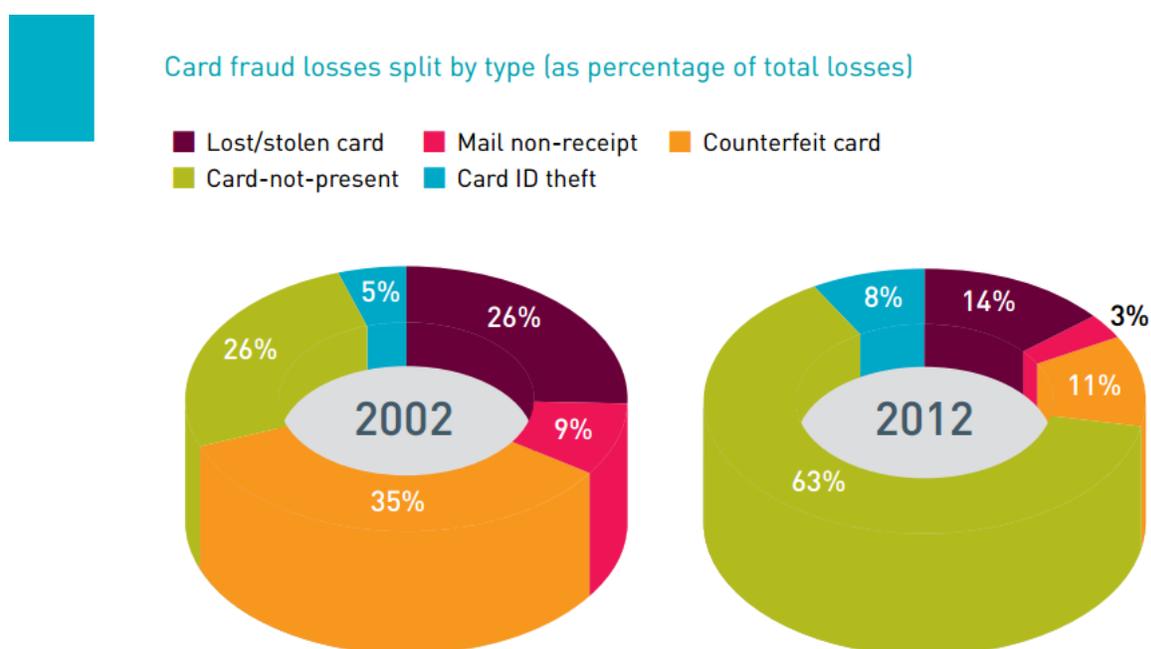


Figure 4: UK Cards Analysis of Card Fraud Loss³

The basic transaction mechanics have changed little since the system was developed. They can be broken down into fundamental mechanics, EMV (Europay, MasterCard and Visa) respectively chip card standards, risk analysis and fraud detection. The process is a linear series of checks, which includes call-outs to different sub-systems.

³ Accessed 27052015.

http://www.theukcardsassociation.org.uk/wm_documents/3533%20Fraud%20The%20Facts%20FINAL.pdf

The fundamental mechanics check includes the following:

- Check whether the card number is valid and on the system
- Check the start and expiry date; this indicates an initial risk, and the risk weighting depends on the specific issuer
- Check the account status [call-out to account system of the issuer]. The external status includes whether the card is marked as lost or stolen; whether there has been fraud identified on the account; or whether the account holder is bankrupt / delinquent or insolvent. The internal status includes the payment history of the customer; a customer that did not pay back outstanding amounts for 3 months would trigger a different check than a customer who did not pay back for 2 days
- Authentication of the customer (in earlier days this was the signature, now it is the PIN or 3D Secure)
- Check security of the card for the bank to confirm that the card is genuine and issued by the bank [another call-out in the checking process is made for EMV checks]. These checks nowadays concern the security code and cryptograms generated by the chip on the card. The EMV data checks give information about the card usage and potential counterfeiting. All security information from the chip can be accessed when following up fraudulent transactions, but this is currently not embedded and requires time.

If the transaction request passes these checks, a call-out is made to assess the risk of the transaction; this could be for example FICO's Triad system. Triad uses FICO scoring (Flacon) in a risk engine over volume and value of transactions. Risk factors depend on the model and the risk score is based on a tailored consumer model that accounts for past consumer behaviour (see Section 3.4.2). Factors can include for example:

- Excessive authorisations
- Purchase amount relative to account balance
- Number of cash withdrawals at ATM

Following these checks, the transaction is either authorised, referred or declined. A transaction can be flagged prior to transaction completion or afterwards. Referred or declined transactions are then passed on to operators in a call centre, via processors such as 'First Data' or 'T-sys'. The rules used by operators in the call centre are usually bank-specific and parameter driven.

2.3.3 The role of humans in the credit card fraud management process

The roles for humans are to either work as call centre operatives or as higher level fraud analysts. Further, technical staff is employed typically an IT roles to go through potentially fraudulent transactions or potential scams at the most technical level.

Call centre agents respond to flagged events. Their role is to follow a script to verify the customer and check with the customer whether the transaction was genuine. Customer verification can include asking for information only known to the genuine card holder, such as past transactions and the mother's maiden name. Information available to the call centre agent is for example account information and transaction information. VISA and MasterCard are working towards removing human-centred referral calls from the process completely. If agents are removed from the decision process, their role would be client management over the longer term with a view to maintain customer loyalty and situation management. 'In flight' referrals will likely be redundant in future because banks do not want to break the transaction/purchase.

Fraud analysts look for more general and emerging patterns in fraudulent behaviour. This could include working on 'exception reports' which are produced at the end of batch processing, in which transactions are followed up that, after an over-night cycle, appear unusual. Technical staff / IT become involved in the process when expertise in security is needed or for example when the potential fraud is associated with irregularities in the card's pin content. Only the very difficult cases are referred to IT.

2.3.4 Conclusions from discussion with UK Cards Association

The first conclusion drawn from the meeting with UK Cards Association is that there are several different forms of 'fraud analysis'. Not only are the different forms undertaken by different people in an organization, but different organisations are involved in each form. For example, financial transactions can be scored in terms of likelihood of fraud. Many organisations employ a standard scoring algorithm, such as that supplied by FICO. The algorithm highlights whether a transaction requires investigation. In D5.1 we distinguished between analysis which focused on transactions on individual accounts and collections of transaction from multiple accounts. The use of a telephone call (or text message) to an account holder represents a response to a transaction on an individual account. Knowing what constitutes 'normal' for each account and tagging this as fraudulent would, therefore, be relevant in this instance (although the tagging is far more likely to be done by the algorithm, with the analyst seeking to interpret the score). On the other hand, overview of a collection of transactions, say by location would be less useful for the individual account analysis (unless the analyst was suspicious of multiple transactions in temporal proximity or unless the analyst was seeking to define 'normal' activity). In such a case, the overview would be useful for a different type of analysis.

The main issue identified with the present working of cases is the missing integration of different systems into a single UI and little or no interlinking between components. For example, EMV details will be accessible in a completely different software suite than the outcomes from the neural network analysis and automatic fraud checking. This requires a lot of work to pull the different information sources together; for fraud on a small scale, this effort would be more expensive than the value of the fraudulent transaction. Hence, certain investigations are dropped as they would create more cost than they could recover.

Another issue is the familiarity of fraud analysts with the individual features associated for example with chip mechanics and the meaning of certain code structures. This holds especially true for new systems that were introduced over the years, especially EMV and fraud risk scoring. Reports created for investigated transactions commonly contain a large amount of codes and technical summaries. If agents are unfamiliar with the meaning of individual items, they are of limited use purely due to an understanding barrier. In future, it would be valuable to present evidence in a standard format that everyone can understand. This would also be very much appreciated by the financial ombudsman and legal representatives dealing with disputed transactions and fraud cases.

Currently missing in fraud analysis at different levels are visual aids to assist the human. This is paired with very limited scope for active thought process, especially at the call-centre level, since most processes are scripted. This results in the exploitation of inexperienced call centre agents by organised crime groups, which may even know more about the systems than the people 'inside' and who can hence bypass security checks. In future, it would be great to develop a system that allows fraud operatives to make better decisions by presenting more information more intuitively; "spreadsheets and code leave people cold".

2.4 FICO (Brian Kinch and Scott Zoldi)

2.4.1 Background

FICO provides market-leading commercial products for financial fraud detection and fraud management. In the company's own words,

*'[...] FICO provides analytics software and tools used across multiple industries to manage risk, fight fraud, build more profitable customer relationships, optimize operations and meet strict government regulations. [...]FICO clients include more than half of the top 100 banks in the world, more than 600 personal and commercial line insurers in North America and Europe including the top 10 US personal lines insurers, 400+ retailers and general merchandisers, including one-third of the top 100 U.S. retailers, 95 of the 100 largest financial institutions in the U.S., and all the 100 largest U.S. credit card issuers and more.'*⁴

The modular products which FICO supplies for fraud management cover 2/3 of all credit cards globally. The modules from which customers can select are separated into Rules, Case management, Data management, Analytics and Customer Engagement.

⁴ http://www.fico.com/en/about-us#at_glance

2.4.2 Transaction scoring

Transactions are scored for their risk based on a large number of (confidential) attributes. The weighting of attributes and cut-off level to classify a transaction as suspicious can be customised by each issuer. This hence varies between banks, each seeking a balance between customer satisfaction and loss. Customer satisfaction becomes compromised if attempted purchases are too frequently aborted by the bank due to the perceived risk. This requires the customer to use a different card and / or go through the process of verifying that the purchase was genuine (if the issuing bank has processes in place that follow up a blocked transaction fast enough, the card may still be used). This is perceived as hassle by the customer.

Fraud models are typically region specific for reasons such as differences in card mechanics (the US has not yet adopted CHIP and PIN), fraud patterns and data quality / sparsity. These regions typically correspond to countries. The bank using the software is able to specify additional rules if it feels that specific fraud patterns which analysts perceived to be present are not represented by the fraud model. Aside the static component, the model is adaptive to changes in fraud patterns and learns based on feedback from flagged transactions that were verified after contacting the card owner using an automated method or human operator.

Scoring is performed through artificial neural networks (ANNs) based on features associated with the transaction. The response to a transaction response is to either approve, decline or refer a transaction. Referrals are very rare and usually mean that no further transactions will be allowed on the card until an operator has spoken to the customer. The weighting of the features is used to inform operators to explain to the customer why a transaction was flagged.

2.4.3 The role of humans in the credit card fraud management process

While transaction risk scoring is completed at the computerised level (neural networks) in near real time, humans are involved to work through cases, verify transactions and provide support. Humans are commonly either employed as call-centre agents to respond immediately to blocked/flagged transactions, or as 'Analysts' to look at fraud patterns in broader terms, for example if regions-specific patterns emerge. These two roles are very different:

Call-center agents follow a clearly defined script in order to establish whether a) the person at the other end of the phone is the genuine card holder and b) whether the card holder made the purchase or whether it was made fraudulently. For this role, agents have access to transaction and customer details, both past and present. This role is today sometimes replaced by computerised solutions, which call or text a customer automatically to confirm the validity of a purchase. Calling the customer is the only option to find out the true state of a transaction, at least in third party fraud (card used by an unknown person). In contrast, first party fraud (malicious intent by card holder) or second party fraud (card user known to card owner) is more difficult to establish, as answers to the agent cannot be trusted. For first and second party fraud, human intelligence is needed in order to determine whether answers to questions are genuine or fabricated.

Alerts are dealt with according to the risk level; alerts with the highest level of risk are worked on first (this is called 'priority mode') and cue the remaining alerts. The bank will set the score threshold so that the number of alerts to be dealt with per day can be matched by the number of employed staff. Banks may vary between 250 and 1500 FTE staff, and they will also show variation in interaction style depending on the customer demographic. The total number of cases to be processed by an operator following referral due to a flagged transaction is around 200 per day on systems such as Adeptra. The most successful banks have a clear policy on how to work cases and employ specialists; banks that struggle more tend to employ less qualified staff (as an example, college kids that do not think about each specific case are less likely to pick up on inhomogeneous information). Analyst quality has a great influence and banks are starting to recognise this (the case load is much lower for fraud on current accounts, where it may be around 10 cases per day).

Higher level Fraud Analysts have access to past transactions, which includes the risk score and the outcome of the decision which the agent made after calling the customer following a blocked transaction. At this level, humans are able to examine fraud patterns and possibly interact with currently flagged cards. If for example several transactions are flagged or confirmed fraudulent at a specific ATM, further requests for withdrawals might be blocked. This is called 'consortium-level pooling' of information. Also, merchant behaviour can be monitored: if three cards in a cue are confirmed fraud, all other cards in that cue may be blocked for reasons of caution.

2.4.4 History of credit card fraud prevention

The history of systems developed to combat credit card fraud are outlined in figure 5.

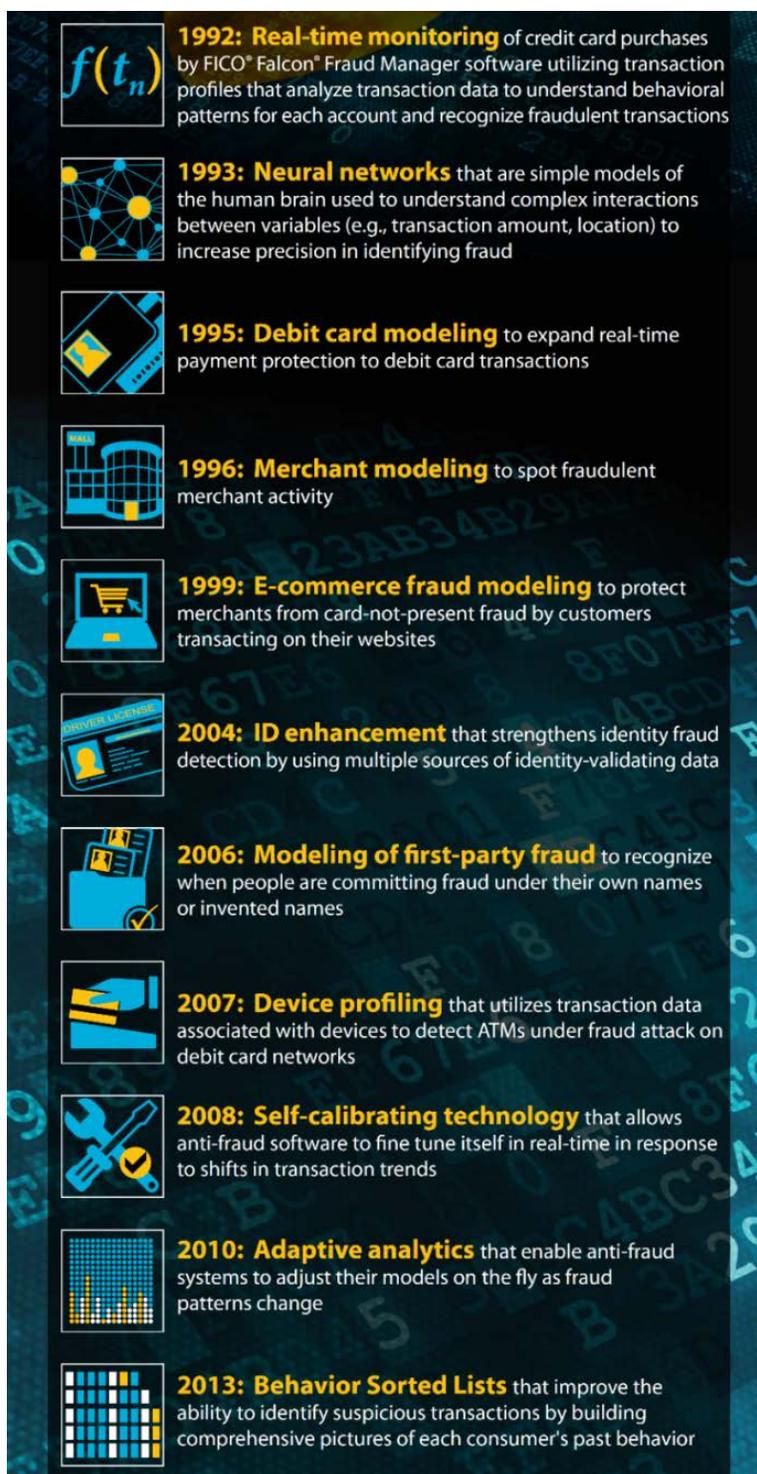


Figure 5: History of fraud prevention measures summarised by FICO⁵

⁵ Accessed on 20052015, <http://www.fico.com/en/node/8140?file=5567>

2.4.5 Information used and user Interface content

The case management UI is typically centred on one case (commonly the bank account), which may contain several alerts. Influential weights (from the neural network output) are displayed to allow the operator to explain the call. Hence, it is important for UI development that the reason for the flagged transaction is given in the UI as it acts as a cue to interact with the customer when calling. The aim is to process a flag as fast as possible; this means that an operator has seconds rather than minutes to think a case through and establish whether the person answering the phone is the genuine card holder. Individual account holders are profiled and alerts depend on this profile; activities of one customer may not trigger a flagged transaction, whereas exactly the same purchase made by a different customer may trigger a flag. The profile builds up over time and becomes more reliable the longer a customer has been with a bank. However, the profile does not know the 'global' spending pattern of the customer across all of his/her credit cards that belong to different banks; it only works based on the bank-specific customer information. The profile is created at three levels, namely card number, account information and customer information. Accordingly, the UI contains information associated with the transaction and customer. This includes:

- Transaction amount, date, type etc.
- Merchant category, ID, country etc.
- Customer summary
- Account summary
- Transactions and history

The investigation log contains information such as

- Outcome (closed / fraud)
- Case tag
- Account activities
- Operator comments

At present, the graphical design is minimalistic, information being displayed as text and numbers, mostly in tables. Different information sources can be accessed via tabs / pop ups / scroll bars and sorted based on feature values. At the higher analyst level, the UI should display top-level information, such as a map of the world to highlight fraud activity; this suggestion fits with the current SPEEDD UI layout. Displayed data should be aggregated and showing commonality.

2.4.6 Conclusions from discussion with FICO

The outlook with regards to software development for the call center operative concentrates largely on data presentation; this could include highlighting of important information. UIs for higher level fraud analysts are being in development (there was one project resulting in a prototype). The content in such systems would have to allow determining and tracking of fraud patterns on a global scale.

3. Comparing User Interface Designs for Fraud Analysis

This section presents a selection of User Interface designs related to the detection and analysis of credit card use. We have selected examples that relate to the management of credit cards by banks (e.g., MIVA), the use of credit cards by individual account holders (e.g., Billguard), and vendors of visual analytics (e.g., Tableau). The aim is to provide an exhaustive review of the designs so much as to highlight similarities and differences in the design in order to position the SPEEDD UI designs in a competitor space.

3.1 Entering Credit Card Details for Online Purchase

Online credit card payment systems, such as the one in figure 6, can require the cardholder to first register a credit card (using name, card number and expiration date). While systems such as PayPal allow users to register the card with a single source, the intention is much the same: to reduce the barrier between making a purchase and completing payment. As such, the link between an account (as represented by card details) and the transaction is mediated by other software. The point is that increasing prevalence of fraud involving 'card not present' (CnP) requires an understanding of how transactions are supported by online purchase systems. It might be useful for the fraud analyst, for instance, to know what type of transaction was being performed, e.g., did the cardholder enter all of their details, did they select a card from a predefined list, was there additional security questions or information requests (such as providing the Card verification code)?

Credit or Debit Cards
Amazon accepts all major credit and debit cards:

▼ Add a Card

Enter your credit card information:

Name on card	Card number	Expiration date	
<input type="text"/>	<input type="text"/>	01 ▼ 2015 ▼	<input type="button" value="Add your card"/>

VISA EUROCARD MasterCard Maestro
American Express DELTA

Figure 6: Credit Card purchase on www.Amazon.com

For the organisations receiving payment by credit card, there is a need to determine how best to manage receipt of such payments. Figure 7 shows an example of an authorization screen which can be used to set up limits for card payments.

Similarly, MIVA provide e-commerce solutions and an example of the interface design they offer for the evaluation of credit users is shown in figure 9. Here, the cardholder is defined in terms of type of user ('bad user') and a score ('86') on the top of the screen. Further information on the transaction is provided in terms of a map and a list of locations (for billing, shipping and the IP address of the computer used to make the transaction).

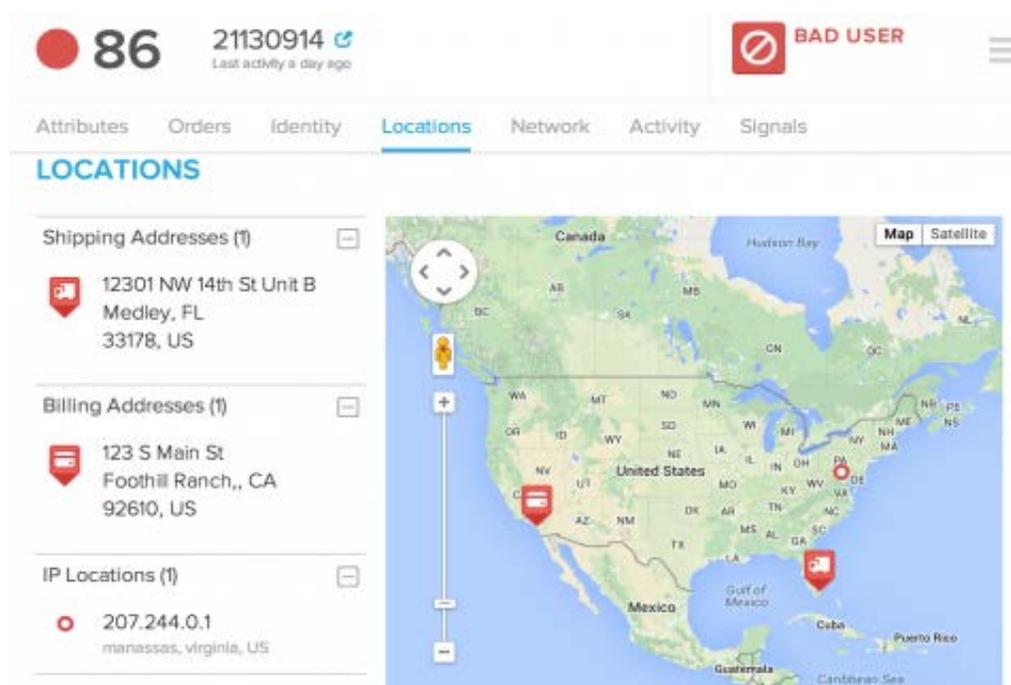


Figure 9: User Interface design for MIVA product⁸

Having this information to hand could be useful for some types of fraud analysis. However, it implies a forensic investigation into account use which might be time consuming and outside the scope of the remit of the analyst. This highlights one of the problems in designing a single User Interface for 'credit card fraud management': there are, as noted previously, several forms of analysis, each with different information demands and different types of decision that need to be made. A similar UI, also using SiftScience tools, was developed by Magento Commerce (figure 10).

⁸ Accessed 26052015. <http://apps.miva.com/product/SIFT-100.html>

The screenshot displays the user interface for Magento Commerce UI, showing details for a user named Crabbe386@deathheaters.org. The interface is divided into several sections:

- User Profile:** Shows the user's name, email, and a large red box with the number 89. Below this, it lists the user ID (09agm4k3h_zeev402208), last login (2014 Apr 20, 7:12 am), and score (2014 Apr 20, 12:34 am).
- Label Section:** Includes a "Label" dropdown, checkboxes for "Bad User" (Chargeback, Spam, Funneling, Fake, Referral, Duplicate Account) and "Not a Bad User", and a "SUBMIT LABEL" button.
- Notes:** A text area for adding notes for this user.
- Statistics:** A list of statistics including "1 ORDER", "1 IP ADDRESS", and "2 ADDRESSES", each with a "SHOW" link.
- Recent Transactions:** Shows a transaction for £8.00 (\$12.93) GBP, successful, dated 7:12 am - Apr 20, 2014. It includes transaction type (886), order ID (36275a0-046a-446f-a093-911a0551a294), and transaction ID (18173d1-293-436a-8530-1e15766a567b). Billing and shipping addresses are also listed.
- Suspicious Signals:** A list of signals with their likelihoods:
 - USER AGENT: Mozilla/5.0 (Windows NT 6.1; Win64; x64; rv:25.0) G... (8x more likely among bad users)
 - USERS PER DEVICE: 2 (7x more likely among bad users)
 - USERS PER COOKIE: 2 (6x more likely among bad users)
 - USERS PER BILLING ADDRESS: 2 (6x more likely among bad users)
 - SHIPPING/BILLING ZIP MATCH: false (5x more likely among bad users)
- User Network:** A diagram showing connections between the user and other users based on various fingerprints (Device, Cookie, IP, Shipping address, Billing address). The network includes users like acromantab@deathheaters.org, sstamend@752@yahoo.com, and sstamend75@gmail.com, with labels indicating if they are "Labeled bad" or "Not labeled".

Figure 10: User Interface for Magento Commerce UI⁹

Another company focussing on e-commerce, CyberSource, describes itself as “the World's First eCommerce Payment Management Company [...] CyberSource provides a complete portfolio of services that simplify and automate payment operations.” The software provided allows merchants to cancel orders that are suspected fraudulent in order to prevent chargebacks. The Merchant Portal designed by CyberSource is shown in figure 11.

⁹ Accessed 26052015. <http://www.magentocommerce.com/magento-connect/sift-science-automated-real-time-fraud-detection.html>

CyberSource[®]
the power of payment

CyberSource Business Center » LOGOUT

Merchant Portal - Merchant A Welcome, Merchant A | Time Zone: EDT

Alerts | My Settings | Contact Us

Alert Queue | Reports |

Alert Timeframe: Transaction Type: Outcome:

Initiated By:

Alert Queue ?

Search Results: 1 - 25 of 27 First Prev. 1 2 Next Last

	Alert Date/Time	Age	Auth. Date/Time	Amount	Card Number	Merchant Descriptor	Outcome	More
	2013-03-25 18:18:27	0d13h	2013-03-25 04:51:03	645.30 (2 Alerts)	506006*****6006	Cyber Global Payme...	Not provided	
	2013-03-25 18:18:27	0d22h	2013-03-24 20:41:07	225.95 USD	402002*****2002	Cyber Payment Secu...	Not provided	

Figure 11: User Interface provided by CyberSource for merchant-based fraud management¹⁰

3.3 Dashboards

Combining multiple information sources for transactions on a single account can give an overview of the activity on that account or by that cardholder. Figure 12 presents an example of such a dashboard, developed by Splunk, to highlight transaction patterns on an account. This combines a map with data on the account activity over time and a display of relative 'risk' associated with the account.

¹⁰ Accessed 27052015.

http://apps.cybersource.com/library/documentation/dev_guides/Fraud_Alert/html/wwhelp/wwhimpl/common/html/wwhelp.htm#href=getting_started.html&single=true

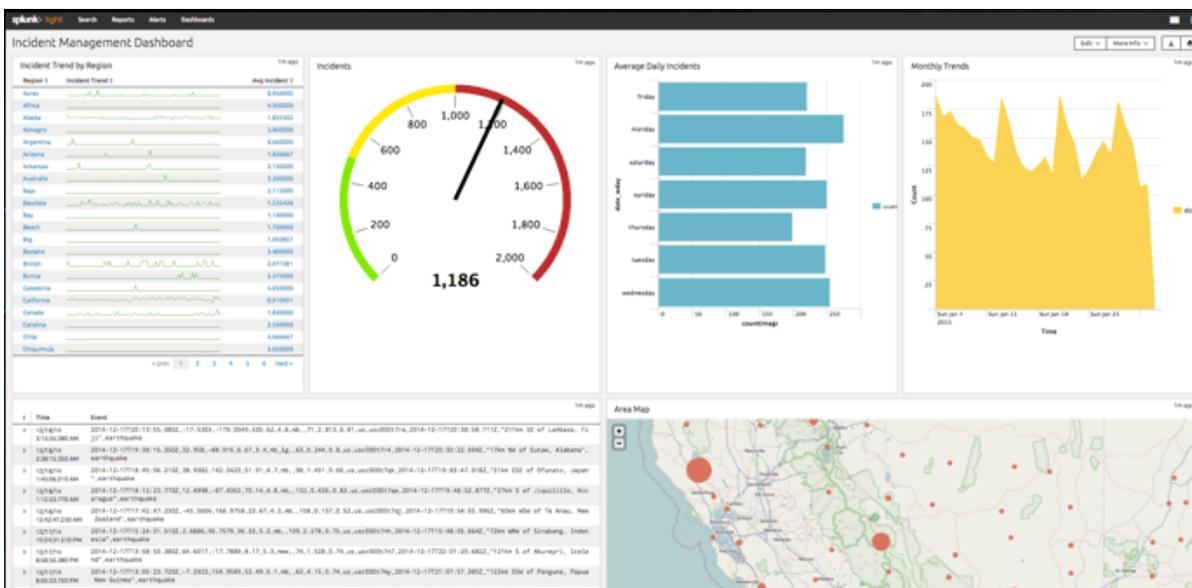


Figure 12: Dashboard for card transaction activity¹¹

NICE Actimize (figure 13) includes another example for a fraud management dashboard, in which various fraud patterns can be examined by means of multiple graphical designs. The company advertises its product to provide “high-volume, real-time, context-driven transaction decisioning, alerting, and resolution for signature, PIN, ATM, and Card Not Present transactions”.



Figure 13: Dashboard of NICE Actimize¹²

¹¹ Accessed 26052015. http://www.splunk.com/en_us/products/splunk-light.html

¹² Accessed 27052015. http://nexttop.ru/wp-content/uploads/2013/03/chart_investigate1.jpg

3.4 Analysing Patterns in Collections of Transactions

While the display of individual activity might help investigate specific cases, fraud analysis also involves appreciating trends and patterns across multiple transactions. In this section, a set of visual display designs are presented which target this high-level perspective on fraudulent activity. For example, FICO provide interactive tools to explore the incidence of credit card fraud across Europe (figure 14) or card skimming in the US (figure 15).



Figure 14: Displaying patterns of card fraud in Europe¹³

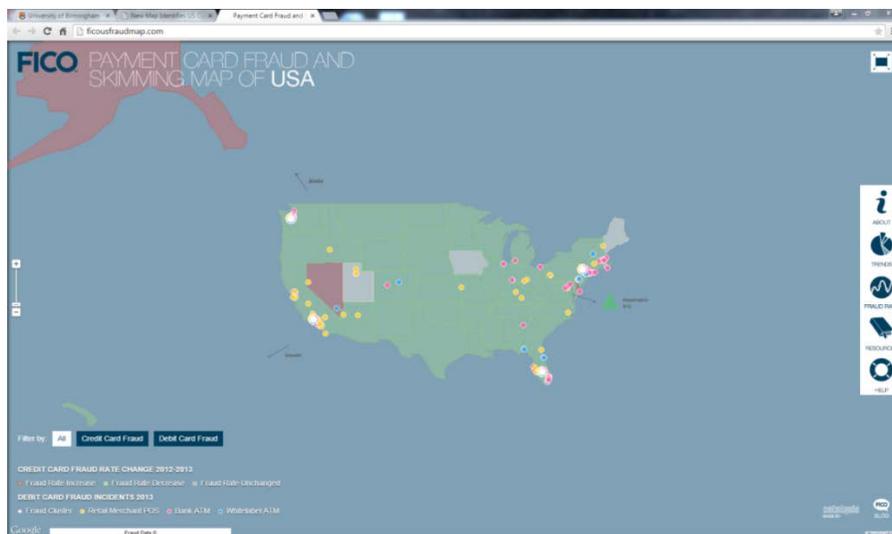


Figure 15: Displaying card skimming in the US¹⁴

¹³ Accessed 26/05/2015. <http://www.fico.com/landing/fraudeurope2013/>

¹⁴ Accessed 26/05/2015. <http://ficusfraudmap.com/>

A popular suite of tools for displaying data is offered by Tableau. These tools can take standard data formats, such as Excel spreadsheets or comma-separated value data sets, and create high quality visual displays. Figure 16 shows a display of a collection of data related to financial transactions. In this case, in addition to presenting data in terms of spatial location, the display also displays the relation between location and amount and outstanding days. In this instance, the focus is not on individual accounts by collections of transactions in different US states.

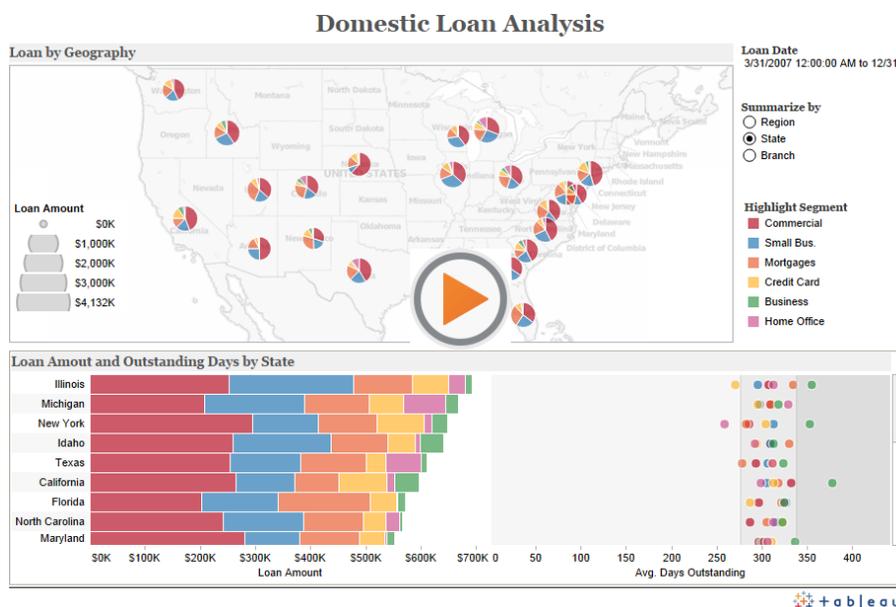


Figure 16: Displaying financial transaction data using Tableau¹⁵

DataWatch.com offers services and tools to visualise financial data. A common approach in the monitoring of credit card transactions is the use of tables showing transaction data. For example, figure 10 shows tables, with some use of colour-coding to highlight salient material. Combining several sources of data can result in displays like the one shown in figure 17.

¹⁵ Accessed 26052015.

https://public.tableau.com/static/images/Do/DomesticLoanAnalysis_0/DomesticLoanAnalysis/1_rss.png

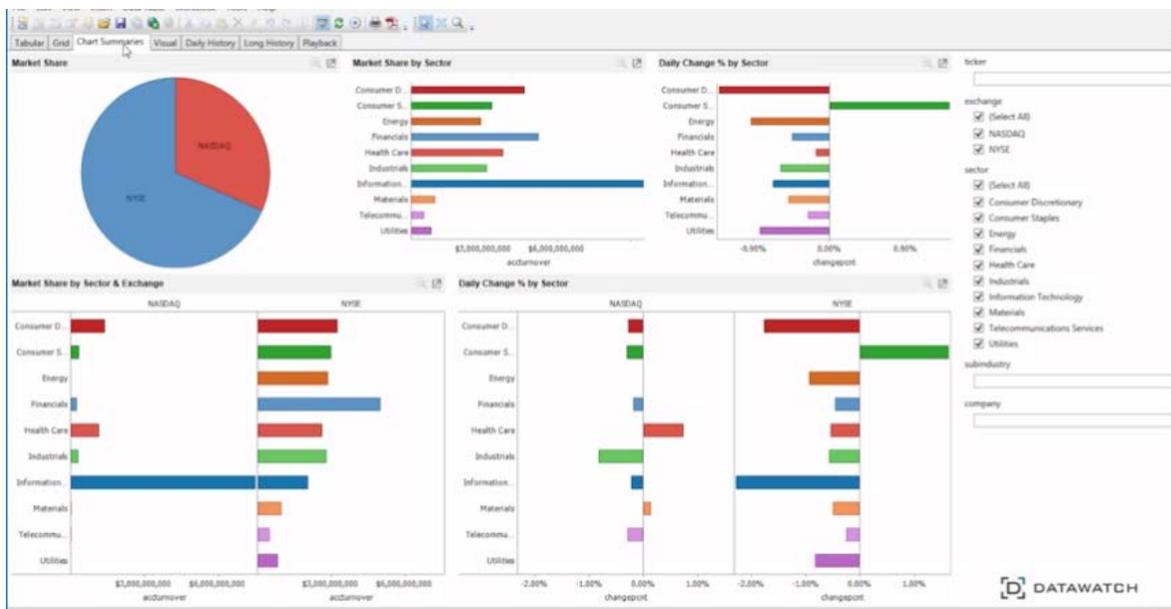


Figure 17: Combining sources of data¹⁶

Rather than displaying the data or relations between values, displays can show trends and patterns in the data. The 'tree map' (figure 18) combines multiple data sources into a single view. The analyst can then monitor this display to spot elements that are changing or which are outside acceptable criteria.

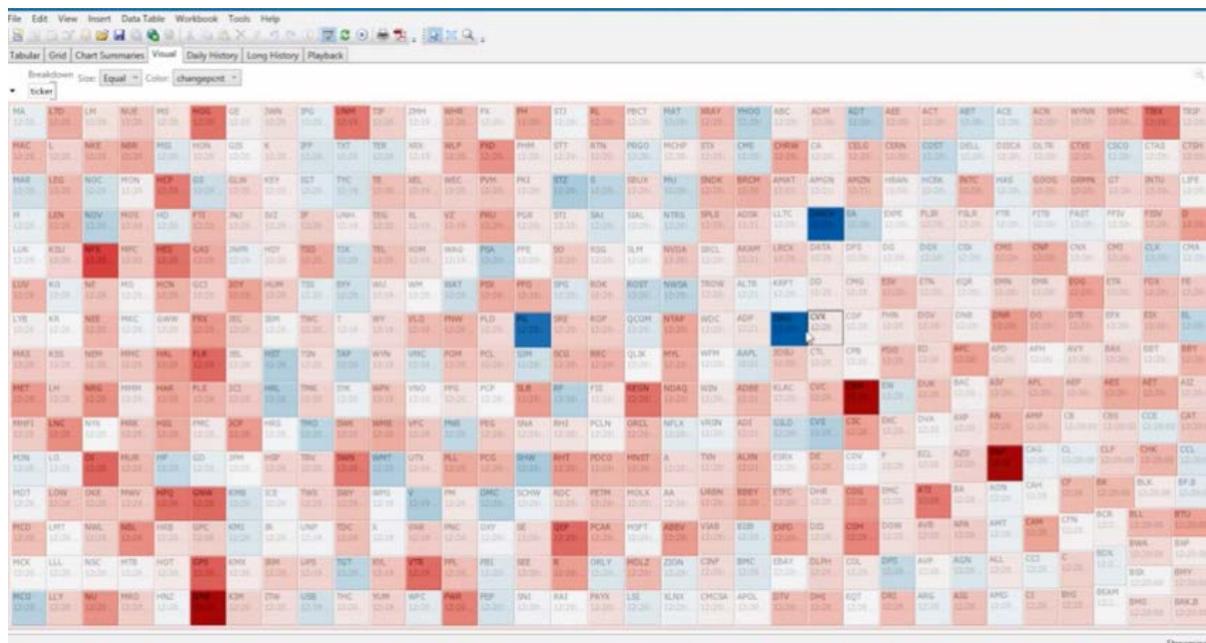


Figure 18: Tree map¹⁷

¹⁶ Accessed 26052015. <http://www.datawatch.com/solutions/industries/capital-markets/compliance-fraud/>

¹⁷ Accessed 26052015. <http://www.datawatch.com/solutions/industries/capital-markets/compliance-fraud/>

Combining treemap, trend displays and data sources produce detailed visualisations like that shown in figure 19.



Figure 19: Monitor by exception in high density visuals¹⁸

3.5 Other visualisations

So far, the report has presented examples of displays which use common forms of visualisation. Recent trends have been towards developing unusual ways of displaying data. In particular, displays like that shown in figure 20 are intended to present relations between multiple sources of data in a way that allows the viewer to recognise changing patterns in these relations.

¹⁸ Accessed 26052015. <http://www.datawatch.com/solutions/industries/capital-markets/compliance-fraud/>

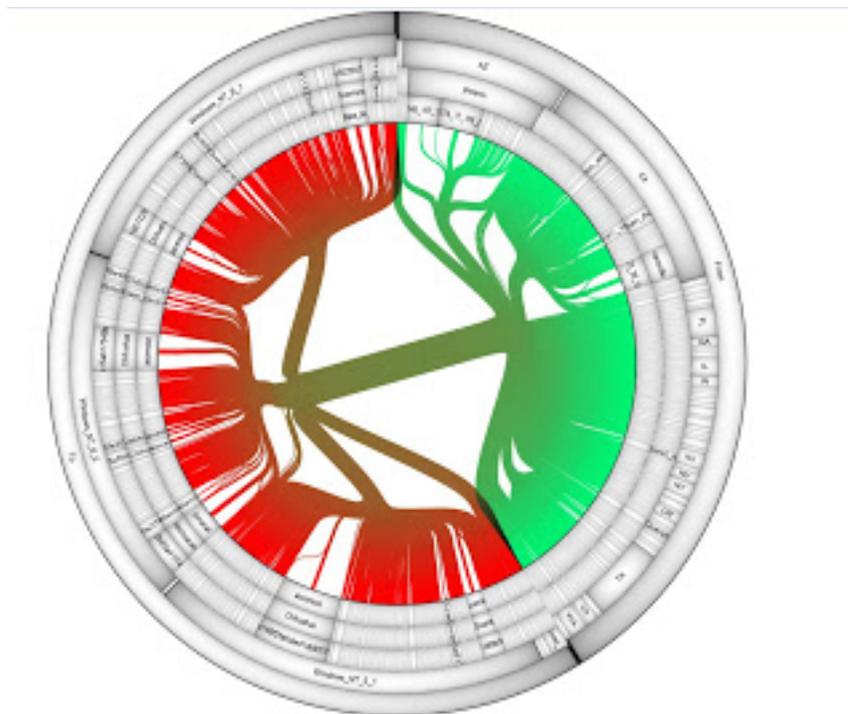


Figure 20: Complex link analysis¹⁹

The approach taken in the WireVis graphical display is to combine multiple views of the same data (relating to the transfer of money by 'wire'). Figure 21 illustrates the ways in which these data are displayed. Similarly, figure 22 shows money transfer in terms of donors and recipients of financial (charitable) aid in different parts of the world.

3.6 Graphical design in SPEEDD

In the SPEEDD project, the graphical design for the initial prototype sought to retain the graphical language that would be familiar to potential users while also allowing integration with the SPEEDD architecture and the specific scenario used in the first year of the project. Consequently, the initial prototype employs a dashboard layout, with map table or geographical map to display data relating to the SPEEDD scenario. In later developments, we will explore alternative forms of visual display. However, the aim is to produce a User Interface which can support analysis of fraud and which the end-users find useful and usable. This aim might be at odds with requirements to produce unusual or novel graphical displays.

¹⁹ Accessed 26052015. http://linkanalysisnow.com/2011_09_01_archive.html

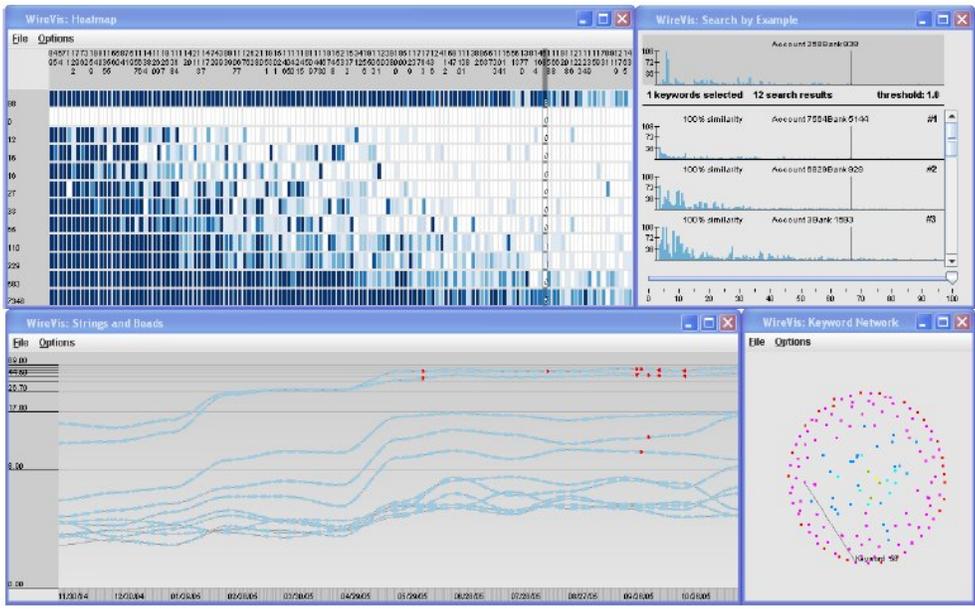


Figure 21: WireVis: Visualization of Categorical, Time-Varying Data From Financial Transactions²⁰



Figure 22: Display of donation patterns²¹

²⁰ Chang et al., 2010

²¹ <http://ilya.boyandin.me/works/2012/09/01/aiddata/>

4. Formative Evaluation and Transition from first to second prototype

4.1 Introduction

As part of the User Interface (UI) development process, a formative evaluation of the first prototype was conducted by FeedZai. This evaluation involved pairs of analysts accessing the UI through the SPEEDD architecture and reporting impressions of the ‘look and feel’ and use of the UI (figure 23).



Figure 23: Evaluation of UI version #1 and #2 at FeedZai in Lisbon, Portugal.

The first credit card fraud management user interface presented to FeedZai is shown in Figure 24.

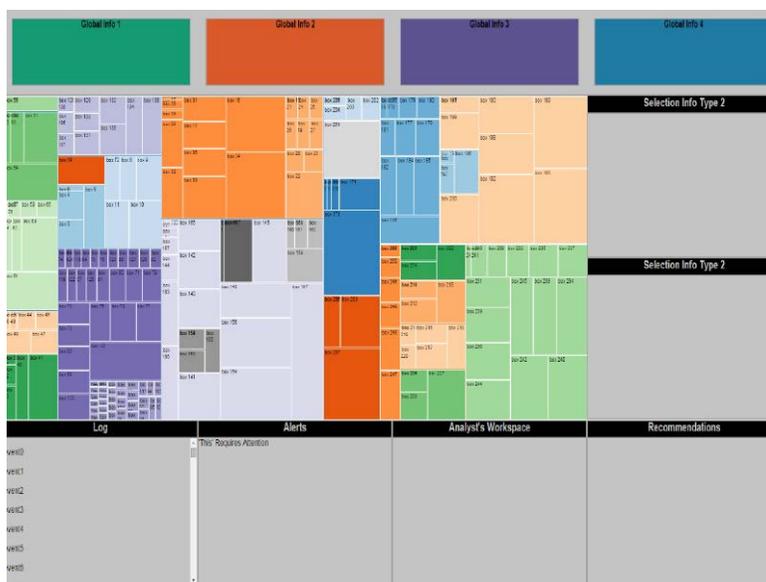


Figure 24: The first version of the user interface.

The locality information is given in the central pan, where the coloured squares indicate the amount of fraud per each country. Some of the critics pointed towards this first version of the user interface was that the colours of the tree map and the ones used for the four top boxes (which report general processing metrics, such as the number of transactions already analysed) could be interpreted as having some kind of correlation due to their similarity, which was not the initial intention.

On the other hand, it was also pointed that the locality information was not being easily transmitted through the tree map. As it aggregates the countries by the amount of fraud, the position of a country can appear anywhere in the map, which makes the analysis process much slower as the analyst may have to look to each square individually to get the information regarding the country of the incoming transaction. This information was pointed out to the partners from University of Birmingham, who produced a new version of the interface.

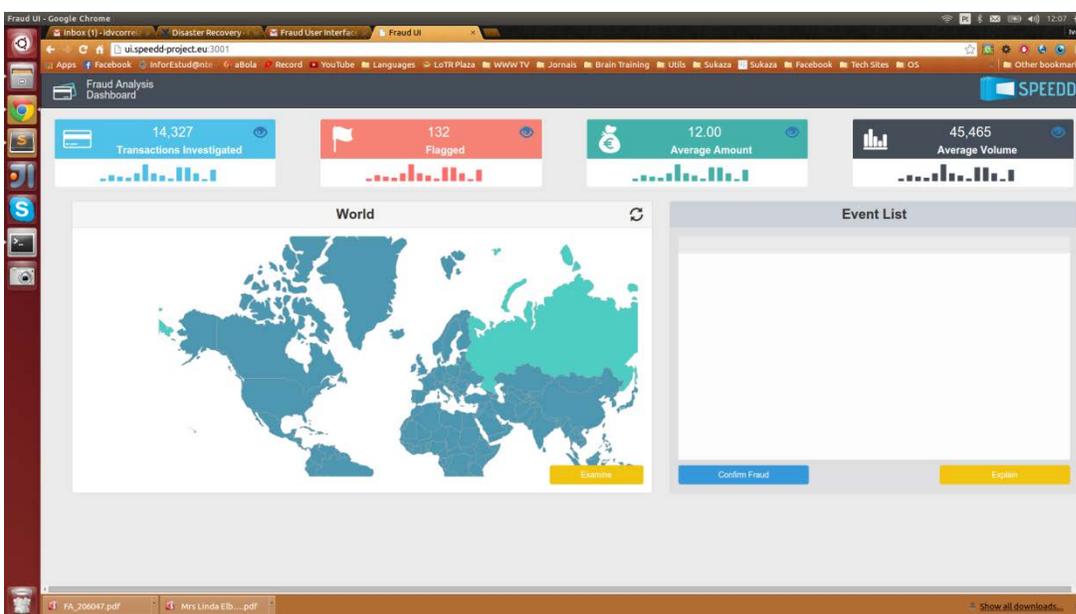


Figure 25: An overview of the SPEEDD user interface.

Looking to Figure 24, and comparing to what was shown in Figure 25, one can immediately see that many improvements were incorporated in the new version of the user interface.

First of all, we have now a world map instead of a tree map, which makes the task of looking for a specific country much easier. On the other hand, in order not to lose the fraud distribution information that the tree map was giving, there is the option to deform the original map accordingly to the selected parameter. This option is accessible by clicking in one of the eyes inside one of the four boxes located in the top of the interface, regarding the number of transactions investigated, flag, average amount and volume.

On the other hand, it is now also possible to get more details concerning the specific transaction in analysis, as well as getting an overview of what is going on around the world. These options are available through the historical data screen (shown in Figure 26), as well as in the transaction and country details (Figures 27 and 28). You can access the historical data by clicking in the bars inside the four boxes in the top, returning data for each context. This fulfils the user requirement for presenting data. However, this still does not attend fully one of the specifications, as currently, the user interface is only showing general historical data, and not yet specific to a card or to a cardholder. On the other hand, this interface allows an easy process to mark down the fraud, with a button located in the right side of the pane. Also, in order to attend the user requirement of reporting and information dissemination, there is the explain button, which will give details related to a transaction.

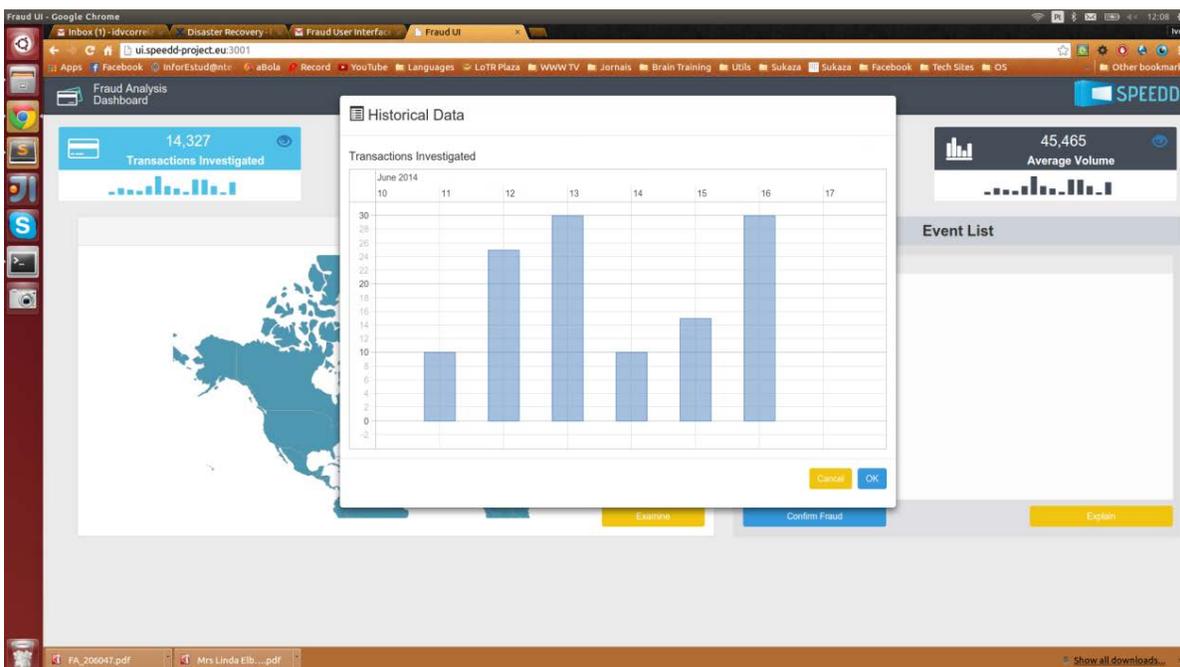


Figure 26: Screen for historical data.

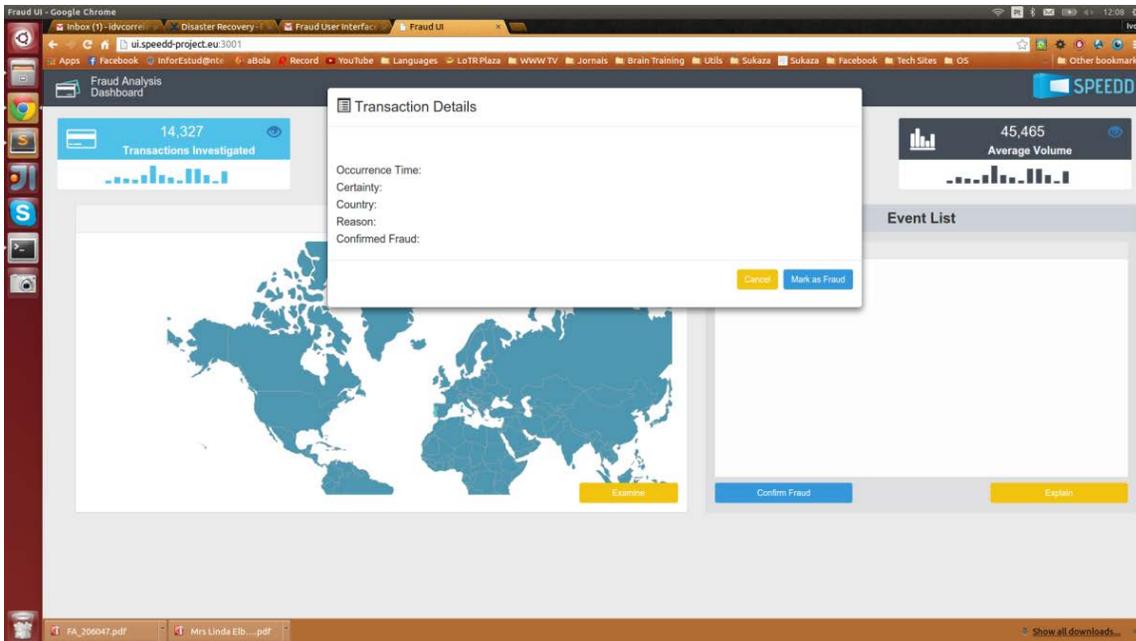


Figure 27: Screen for transaction details.

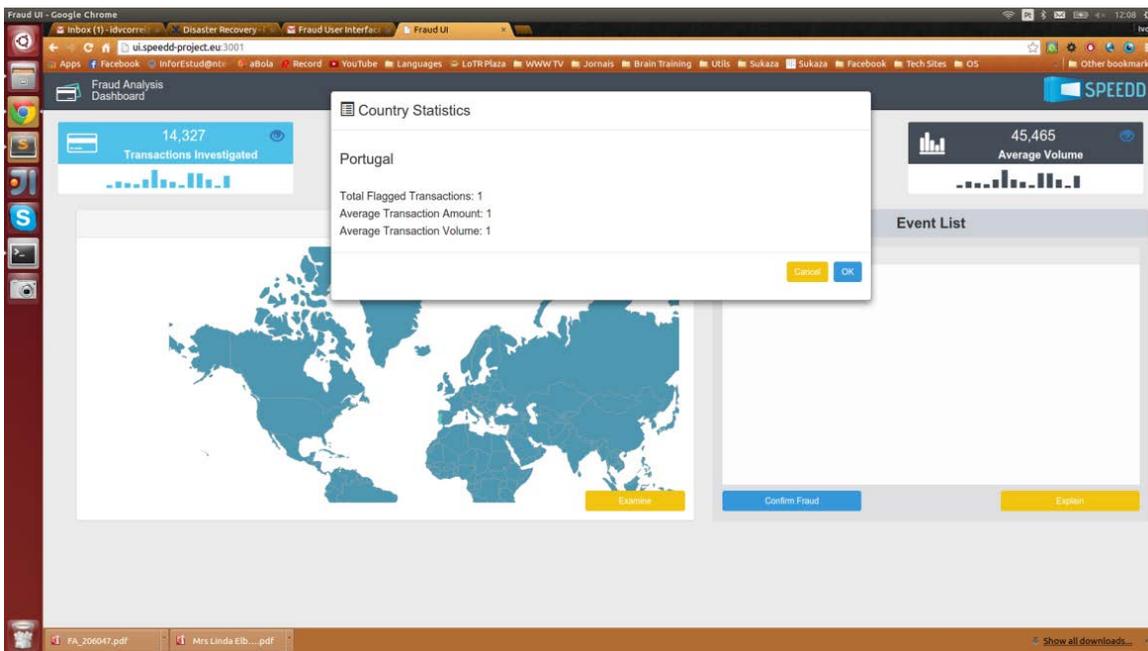


Figure 28: Screen for the details of a country.

4.2 Usability of Evaluation of the Two Versions

As in D8.3, a usability evaluation of the UI design was performed using the Software Usability Scale questionnaire (Brookes, 1988). The questionnaire was translated into Portuguese (the English and Portuguese versions are in Appendix A). As noted in D8.3, SUS scale consists of 10 simple questions concerning the potential usefulness and benefit that users feel that the User Interface might provide them. Each statement is rated on a scale of 0 to 4. The scoring of responses then involves subtracting 1 from odd-numbered questions and subtracting scores of even-numbered questions from 5. This is because the questions alternate between positive and negative connotations. Scores are then summed and multiplied by 2.5, to give a final score out of 100. As a rule of thumb, scores in excess of 65 are deemed 'acceptable'.

Four employees of FeedZai participated in this evaluation (1 female and 3 male). All employees had knowledge of fraud analysis and one was specifically employed to analyse fraud patterns. While none of the participants were professional financial fraud analysts, it was felt that their knowledge of the domain provided sufficient experience to allow them to make informed evaluation of the prototype designs.

The participants were presented with the UIs shown in figures 25-28 and asked to rate these using the SUMI scale. Figure 29 the results of this exercise.

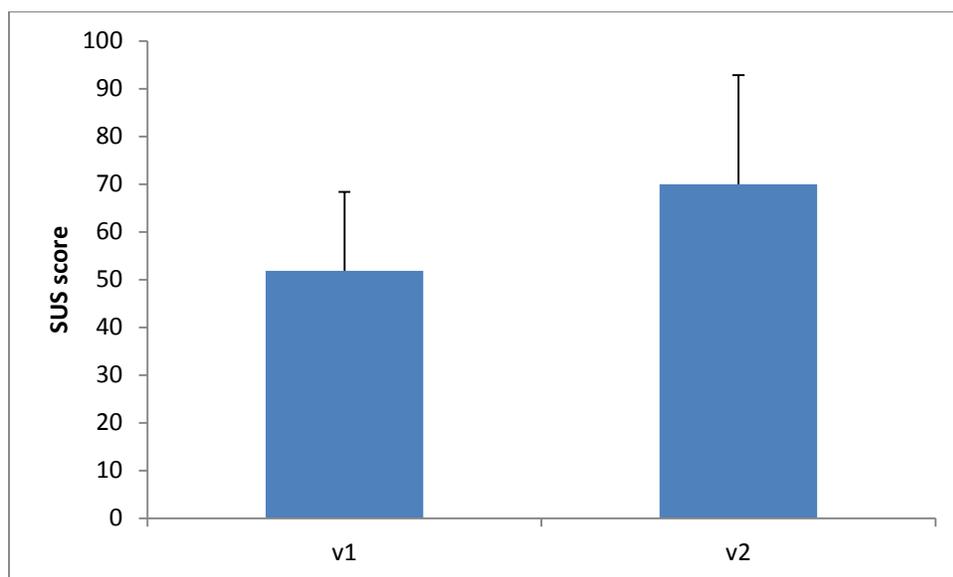


Figure 29: SUMI evaluation of initial prototypes

From this evaluation, the design of V2 exceeds the 'acceptable' level. This implies that, in terms of usability of the UI, V2 is meeting an acceptable standard. The next question to ask is whether the content of this dashboard can inform decision making.

4.3 Additional Comments

During the usability evaluation, FeedZai staff made a number of comments and queries on each version, which will be considered for future iterations of the prototype:

Prototype #1:

- Query regarding what the size of each rectangle represents
- The re-location of countries within the tree-map as values change caught attention
- Query regarding the colour scheme and association with continents
- Query regarding the insertion of Boolean operators for customised fraud patterns
- Comment that it would be beneficial to be able to search for specific characteristics, such as amount > x

Prototype #2:

- Query regarding the lag between clicking a button and response of the world map through resizing (will have to be addressed in terms of system response lags)
- Query regarding what the bars underneath the top row dashboard panel represent (they represent months)
- Query whether the event list could be filtered by country etc. (this can be incorporated in future and is partially implemented already)
- Query on how information shown in the world map and interaction with the world map affects the event list
- Suggestion to display the name of the country on mouse hover over a country, as an operator may not be able to identify a country after it changed size or if it is from a less well known region
- Query regarding the facility of labelling a transaction by an operator
- Comment regarding the function to move window panels around (at present the UI includes the ability to re-arrange the user workspace based on personal preference; however, this function was not made explicit)

5. Understanding the Fraud Analysis Process

5.1 Introduction

In D5.1, we presented an initial sketch of the strategy that we speculated could be followed by fraud analysts. This is shown in figure 30. The aim was to present a 'best-guess' description of how analysis might be undertaken.

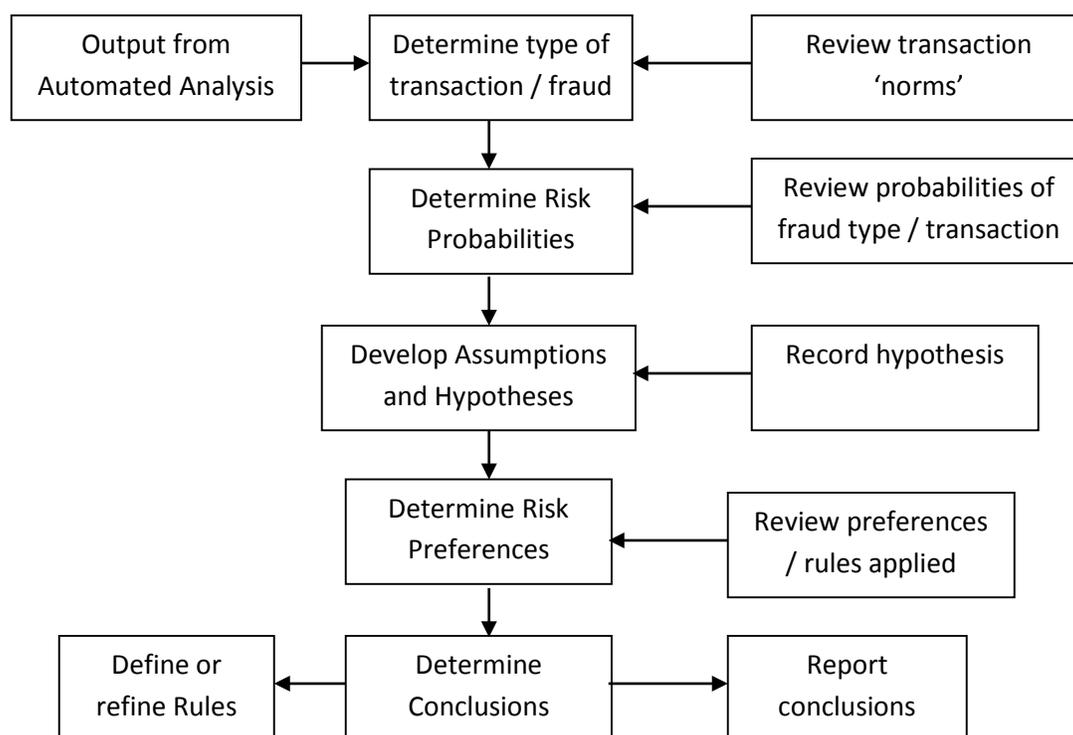


Figure 30: Possible strategy for fraud analysis

This workflow analysis was used to guide assessment of operator workflows and to design an abstract lab experiment investigating feedback-based fraud pattern recognition and information integration. In the abstract lab experiment, operators were presented with more than 25 transactions in sequence. The attributes of the transactions were chosen according to defined statistical distributions of what would constitute 'normal' and fraudulent. These rules were chosen narrowly for the lab experiment to investigate whether operators quickly learn to detect a few, very distinct, fraud patterns and associated rules. This lab experiment relied on implicit risk probabilities of four staff at FeedZai, who were familiar with the process of fraud detection. We were then able to track whether participants updated risk probabilities after receiving feedback on each evaluation on whether their classification of

transactions was correct or not, and we were able to track the learning effects over time. Interviews appended to this lab experiment allowed us to directly ask for the rules developed and patterns spotted in the experiment. These could then be compared to the known injected fraud patterns.

5.2 Using Fraud Analysis Software: a case study conducted at FeedZai

Operators usually work on workstations with two monitors. One monitor contains the FeedZai UI (figure 31), the other is used for other, supplementary, software packages, such as Excel or SQL.

The FeedZai UI contains several panels with various information sources (see screenshot below for the general layout).

The information sources available include:

- Cued transactions
- Transaction information
- Global location of card issuer and purchase made (displayed on world map)

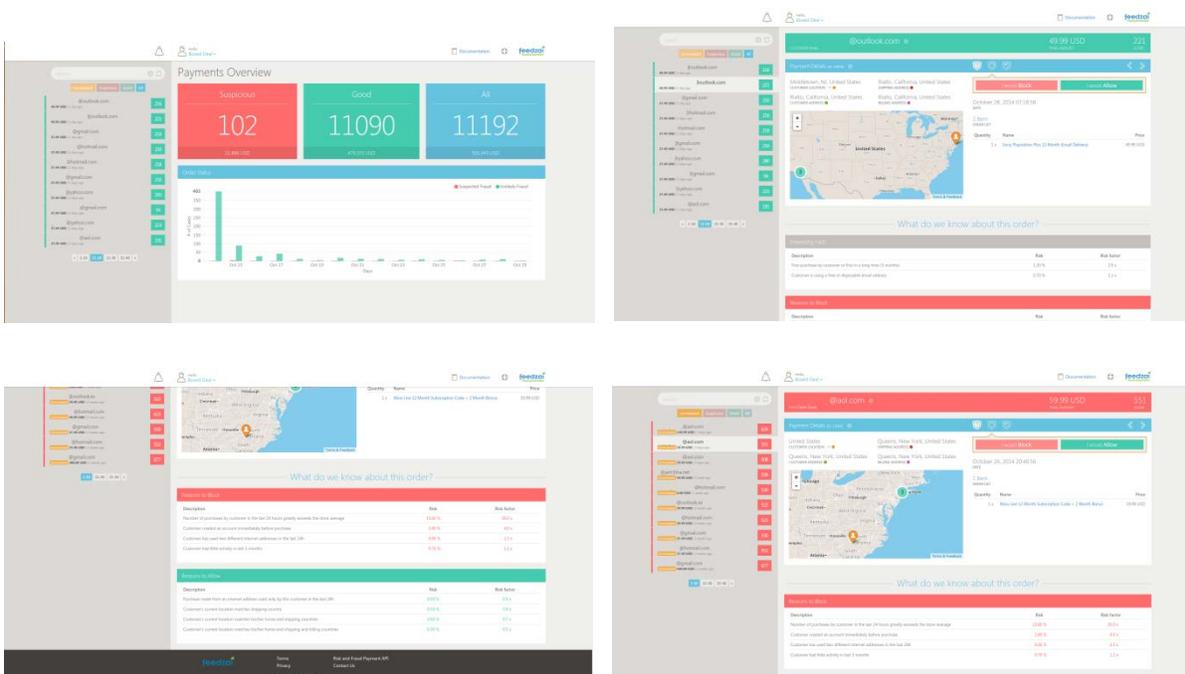


Figure 31: Screenshot of FeedZai UI for fraud management.

5.2.1 Fraud analysis workflow: interviews following work through the 10 sample transactions

Participant #1

As the first step, transactions are ordered by date, so that the older ones are reviewed first. For the first transactions, the participant then looked whether there have been other transactions by the same user / customer to see if there was any pattern. Then transactions are looked at in the map; whether the card information matched the customer information (whether e.g. addresses matched); the amount of the purchase; the information provided by FeedZai. Also, specific countries are risky countries, or activities such as a card in China being used in Japan to ship things to the Philippines.

Participant #2

The participant checked whether the locations were convergent, with a particular interest in where the current location of the customer, his address, and the place where something was being ordered to; where these three were all different, the transaction was classed as fraud. Also it was checked whether there were multiple transactions with the same IP or from the same customer.

Participant #3

As the first step, transactions are ordered by score, from the highest score to the lowest score. For each transaction, as a first step the IP /user ID is checked to see whether it recurs in other transactions that are suspicious of fraud. Then the card country and user country are verified; a mismatch between the two would be very suspicious. Also, the reasons for blocking or for allowing the transaction given by the computer are read.

Participant #4

After looking at the number of transactions and seeing that there were not that many, the transactions were ordered by score from the most suspicious one to the least suspicious one. Then transactions were worked on one by one, mostly looking at differences in shipping and billing country and address (including the IP address). Two transactions were not suspicious at all. One transaction appeared not suspicious, but then the participant saw other transactions with a similar pattern, where purchases were made around a similar area but the name on the cards were different; so all these were flagged as fraud.

5.2.2 Fraud analysis workflow: detailed workflow explained by dedicated fraud analyst

In an example of a transaction with a high risk score that could not be easily classified as genuine or fraudulent, the workflow was as follows: in the user interface, all **transactions by a specific user** are called up. This shows that the user has a lot of transactions and a high risk score, which makes the user very suspicious. Then, all the transactions from this IP address are looked at; if there are several users making purchases from the same IP, that would also be very suspicious. But in the present case, there was only one user for the IP and one purchase from this user ID. If there were more, it could be checked whether they matched each other, the amounts were similar, or what the pattern was of the

user IP. As a next step, the **customer information** is called up as a pop up from the UI. The customer address and card's country are checked; if the two are the same, the behaviour is not that suspicious. The pop-up is then closed and the **transaction information** is called up as a new pop-up, where the addresses related to the purchase are shown. Following this, the **reasons given by the software** for accepting the transaction or declining the transaction are examined based on the machine learning output. All the gathered information is then combined to tag the transaction as fraud or normal. In this specific case, the analyst suspected fraud, but would have preferred more information to make the final decision. These steps are illustrated by figure 32.

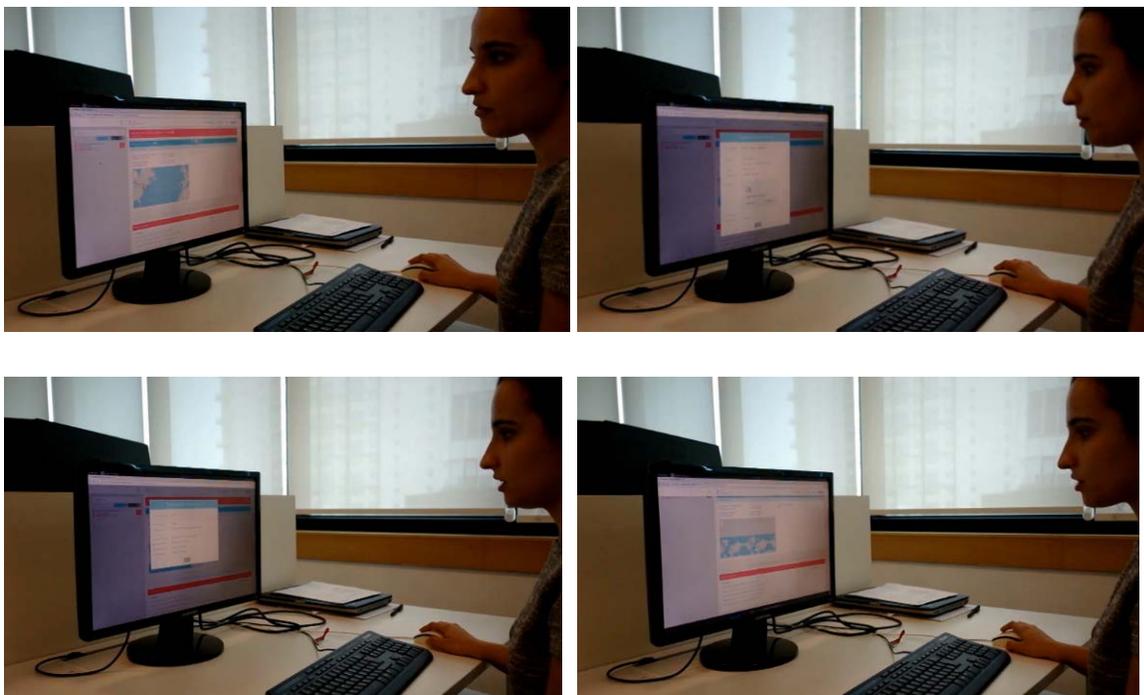


Figure 32: Subset of stages within a workflow during examination of a suspicious transaction.

In an example of a transaction with a lower risk score, the reasoning was as follows: as a first step, location information was checked and it was observed that shipping address, customer address and billing address all were different. On examining the risk score, the analyst noted that according to the score this transaction would be considered suspicious (risk score above the 500 threshold), but a person tagged it as allowed. In the next step, the computer generated reasons for allowing the transactions are examined. Since the only aspect of the transaction that appears suspicious is the different locations, the analyst would hence allow this transaction, too.

5.2.3 Gaze behavior while using the FeedZai UI

To quantify baseline gaze behavior while using a proprietary software suite, the four FeedZai employees working through an identical set of 10 suspicious transactions were equipped with eye tracking glasses (Tobii glasses v.1). Participants were instruction to follow their usual routine. The distribution of fixations is shown in figure 33. It is apparent that the space of the monitor is used in a strongly unbalanced way, with most relevant information viewed towards the top left corner of the screen.

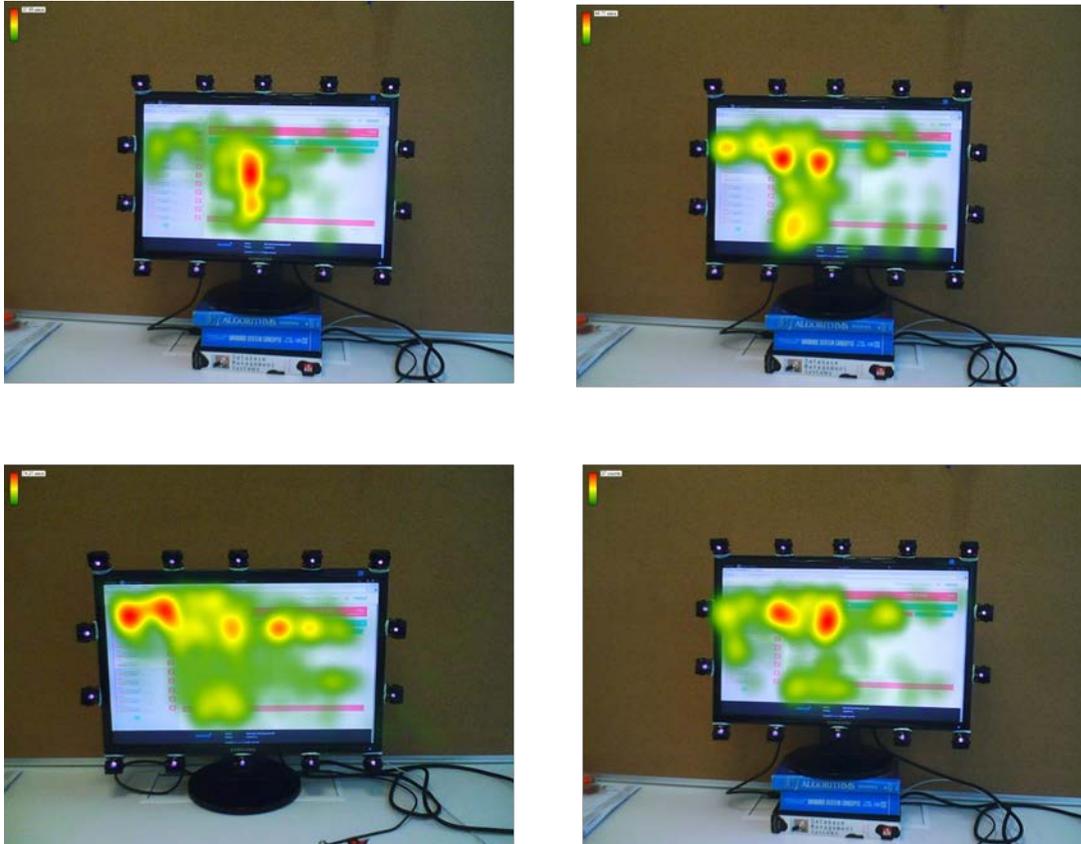


Figure 33: Distribution of fixations across the monitor while working with the FeedZai UI.

5.3 Laboratory Task to Investigate the Fraud Analysis Process

5.3.1 Background for the work

In this lab experiment we abstracted a credit fraud management user interface for a pilot study into user interaction with information relating to credit card fraud. We performed this work in order to examine how fraud analysts scan different information sources, which sources they attend to, whether they develop systematic scan patterns and whether they pick up on fraud patterns injected into the system. This experimental study was designed with cognitive modeling in mind, which will allow comparing user behaviour to optimal behaviour and ultimately design user interfaces that avoid redundant information and allowing for efficient and effective human analysis of complex systems.

Real time visual analytics interfaces for SPEEDD will show multiple information sources, some with updating content and varying rates of these updates. However, making multiple information sources available does not guarantee usage by a user, as we have already shown in our analysis of operator behaviour in the road traffic control room study performed at DIR-CE. User interfaces will also have to take into account and individual's preferences for specific content / presentation style and the weighting of information reliability.

In both the fraud management and road traffic use case, developing an understanding of a 'true' situation and necessary action requires integration of information across sources. This is constrained by limits to (working) memory while also relying on top-down attentional control of a user. Hence, and understanding is needed how people develop an information sampling strategy and how this compares to optimal solutions.

5.3.2 Experimental design

The experiment was based on an abstracted / simplified user interface displaying nine information sources relating to a credit card transaction created in Matlab using custom written code by Sandra Starke. The content of each source was based on transaction attributes frequently used in the literature on data mining in credit card fraud based on the literature [Bhattacharyya et al. 2011; Caldeira et al. 2012; Duman et al. 2011; Hand et al. 2008; Krivko 2010; Leonard 1993; Leonard 1995; Sahin et al. 2013; Whitrow et al. 2009; Wong et al. 2012] as well as UoB's discussion with UK Cards Association and FICO. The nine information sources contained:

- Transaction amount
- Transaction History
- Card present
- CVV entered
- Location info: card
- Location info: purchase
- Purchase category *[from literature]*
- Mean category purchase value *[from literature]*
- Merchant ID

The user interface further contains a panel showing account information and times, buttons which the participant interacts with as well as feedback on each decision (figure 34).

The target of the task was to correctly evaluate 25 transactions as fraud or normal. No information on patterns was given. The percentage of fraudulent transactions averaged 20%. Two fraud patterns were injected into the system, each making up on average half of the fraudulent transactions:

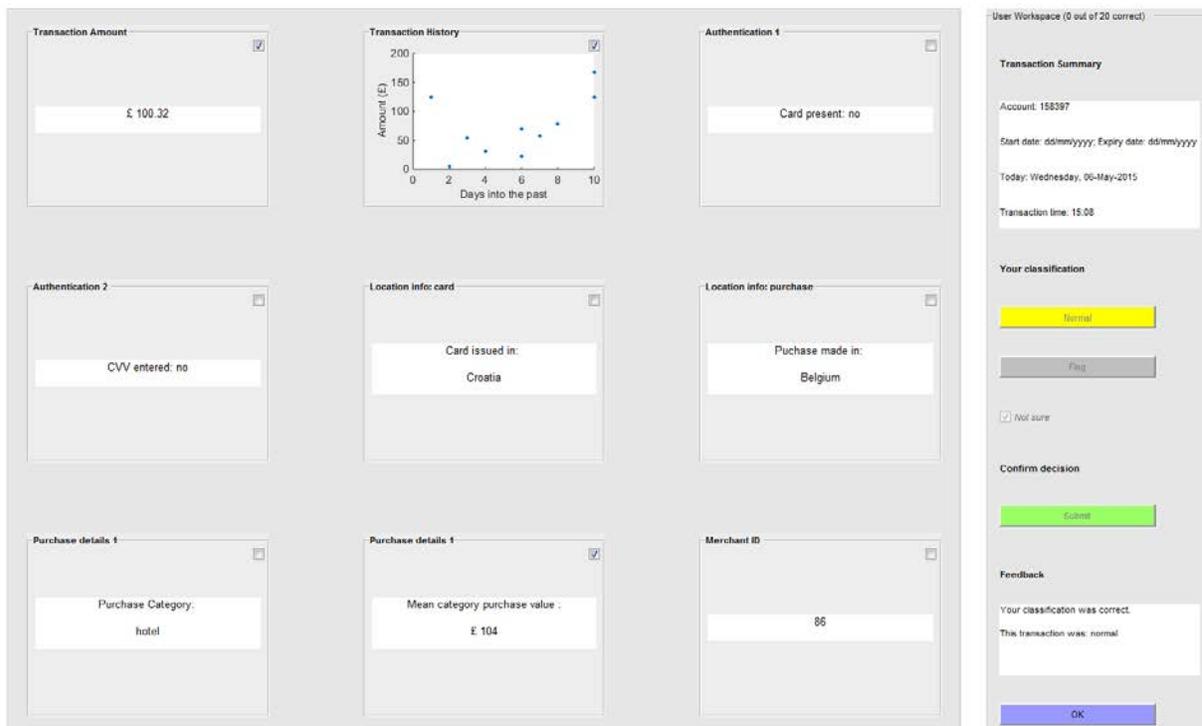


Figure 34: Simplified UI displaying nine information sources related to a credit card transaction.

1. *Pattern #1* consisted of several small transactions made on the same day (figure 35). The value of the corresponding multiple historical transactions and the current transaction was drawn from a uniform distribution, where a random number from -20 to 20 was added to an average value of 30.

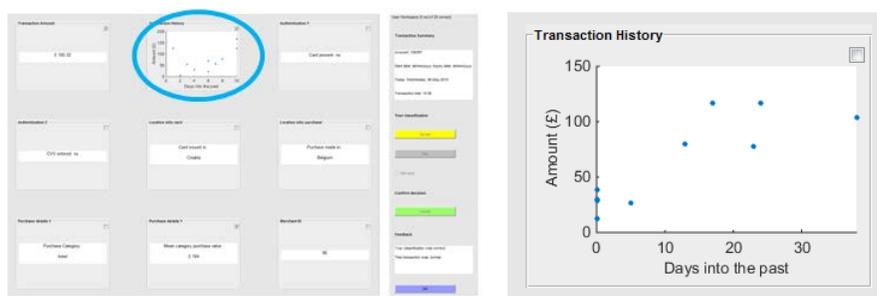


Figure 35: Information showing fraud pattern #1, several small transactions on the day of the current transaction.

2. *Pattern #2* consisted of a single large transaction (figure 36), drawn from a uniform distribution, where a random number between -500 and 1000 was added to an average value of 1000. Hence, fraudulent transactions ranged from £500 to £2000. This was just touching on the normal distribution of the highest mean + 3 SD purchase category value (£531). For this pattern, it was (at least initially) beneficial to compare the amount to the mean purchase value for the purchase category, shown in panel #7 and #8.

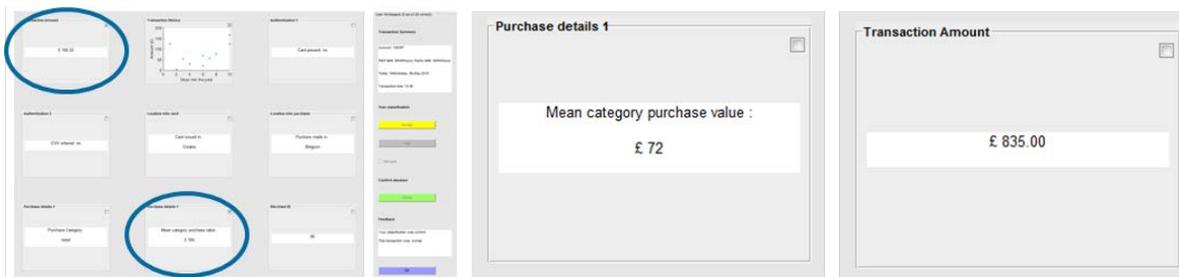


Figure 36: Information showing fraud pattern #2, a single large transaction.

Based on these patterns, the different panels or ‘cues’ had different use to the participant, which can be separated in main cues, weak cues and irrelevant cues:

Main fraud cues

- Transaction amount (compared to mean purchase value)
- Transaction history
- Merchant ID (always the same 2 fraudulent merchants)

Weak fraud cues

- In fraud, always both CVV and card present or both not present

No correlation with fraud

- Relationship between ‘card issued’ and ‘purchase made’ location
- Purchase category

For each trial, each field in the UI was populated with data drawn for specified normal or uniform distributions. Values for the purchase category and mean purchase category value were taken from the published literature [Hand et al. 2001]. Participants were able to learn based on feedback given on each transaction (figure 37), enabling them to test hypotheses about suspected fraudulent behaviour.

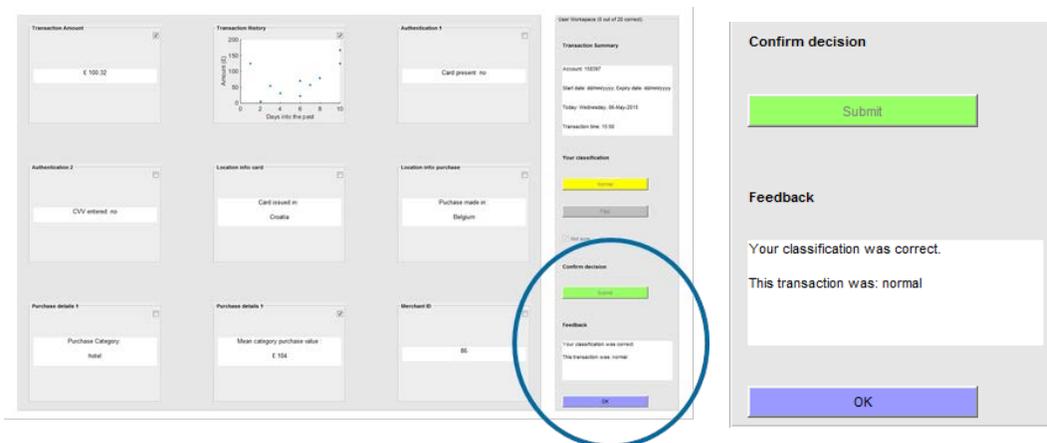


Figure 37: Feedback given on each transaction evaluation.

5.3.3 Participants

Four staff from FeedZai participated in the experiment. The task took approximately 10 to 20 minutes to complete. Participants were equipped with eye tracking glasses (Tobii glasses v.1) while recording all their interactions with the UI through Matlab. After the task was completed, interviews were conducted to query what pattern participants had observed.

5.3.4 Preliminary results

Decision times averaged around 15 seconds per transaction (figure 38), with large variation noted for participant 3, which was the most experienced fraud analyst. Misclassifications were observed for both normal and fraudulent transactions, and there was no obvious decline in misclassification with time. We assume that the period for learning was rather short, as typical experiments of this kind usually include an extended training phase and around 200 to 400 tasks in the testing phase. However, this was not possible given time constraints and the participants work requirements.

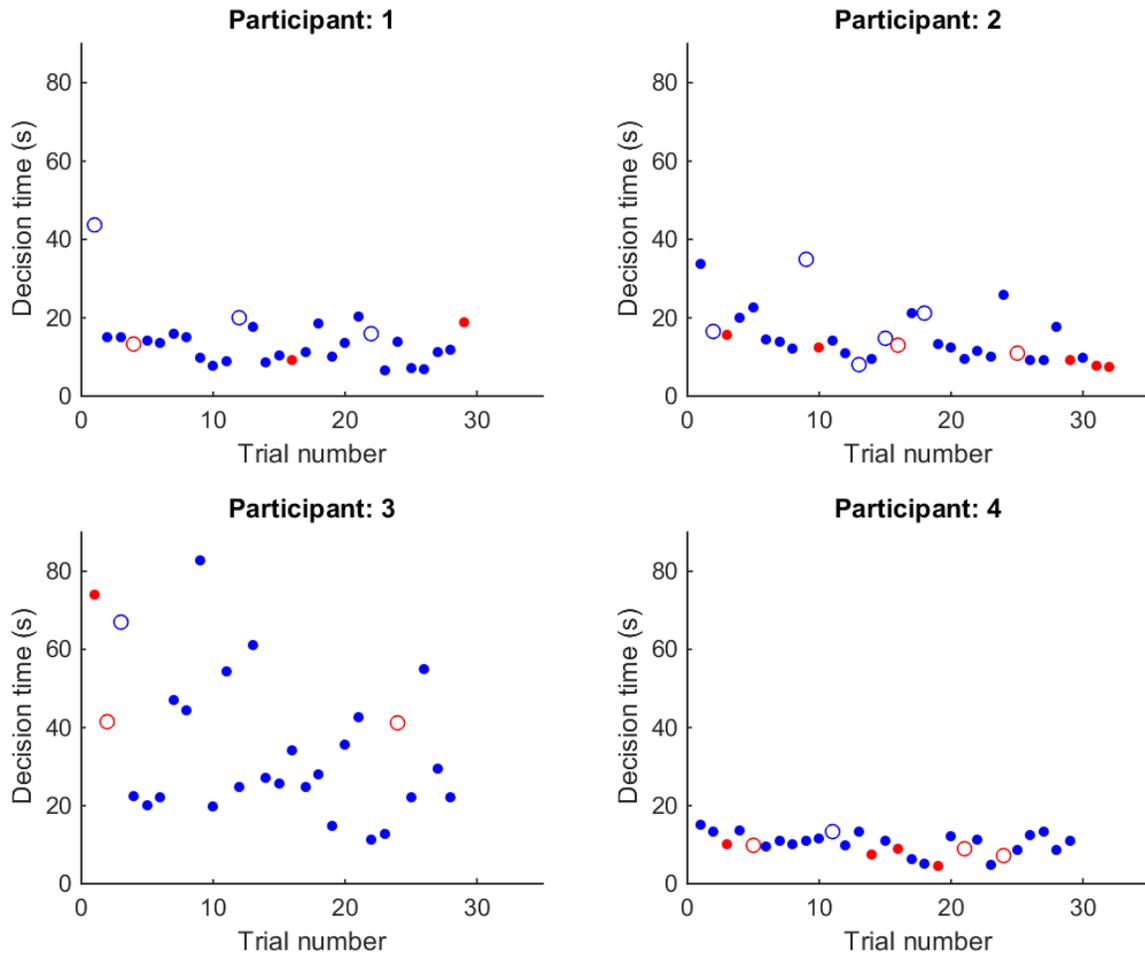


Figure 38: Responses to individual transactions by four participants.

[Blue solid: normal transaction, classified correctly. Red solid: fraudulent transaction, classified correctly. Blue empty: normal transaction, classified incorrectly. Red empty: fraudulent transaction, classified incorrectly. Note the large variation and extended decision time for participant 3.]

Participants looked at all available information sources across all trials (figure 39), however the weighting differed between sources. Sources that were look at the most were:

- Transaction amount (panel #1)
- Transaction History (panel #2)
- Card present (panel #3)

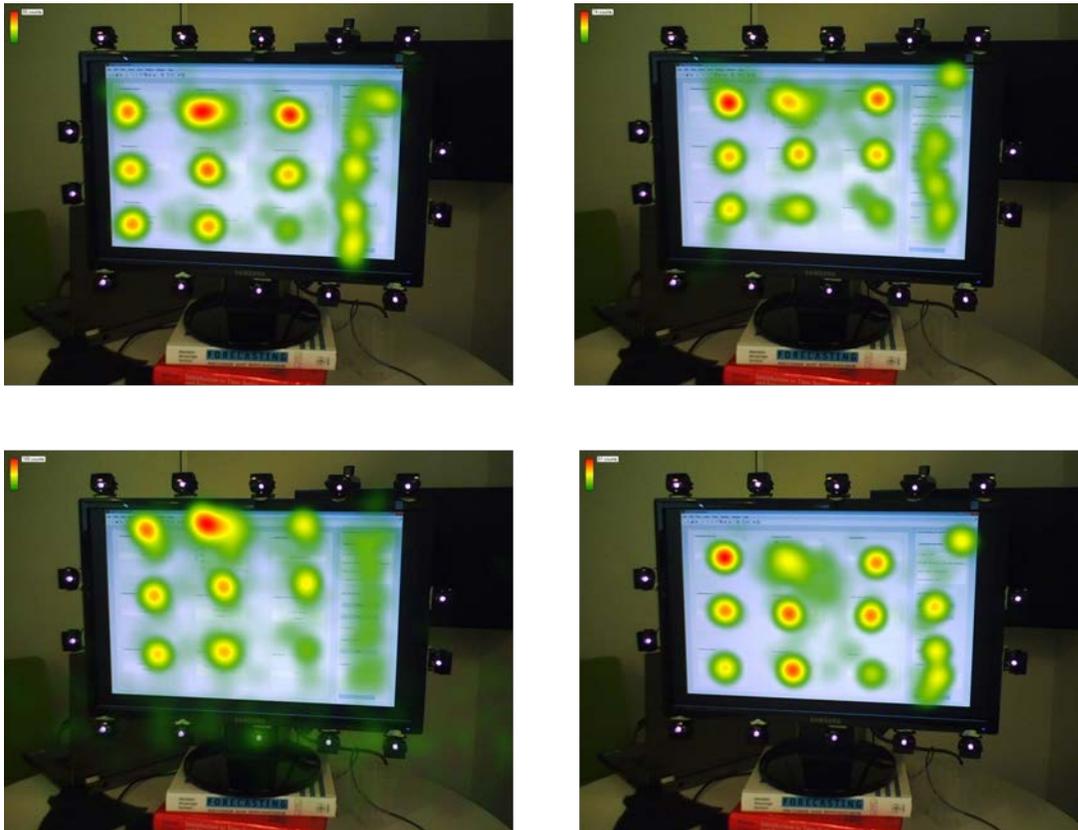


Figure 39: Visual attention allocation to the different information sources across all trials.

Preliminary analysis of dwell times and the total viewing time across all trials showed that the mean dwell time (calculated across median dwell times per participant) ranged from 0.40 to 0.67 s and the mean total viewing time ranged from 11.07 s (panel #9) to 102.74 s (panel #2).

Table 1: Dwell times and cumulative viewing times.

Dwell times and cumulative viewing times for all ten panels (regions of interest, ROI); panel 10 is the panel on the right of the UI which the participant interacts with.

	Dwell time (s)	Total time (s)
ROI 01	0.48 ± 0.04	49.69 ± 23.83
ROI 02	0.58 ± 0.11	102.74 ± 82.46
ROI 03	0.45 ± 0.02	42.85 ± 19.81
ROI 04	0.46 ± 0.02	32.70 ± 16.23
ROI 05	0.41 ± 0.04	32.93 ± 6.12
ROI 06	0.42 ± 0.04	25.92 ± 3.41
ROI 07	0.67 ± 0.11	42.03 ± 14.67
ROI 08	0.51 ± 0.08	40.45 ± 18.48
ROI 09	0.40 ± 0.05	11.07 ± 2.07
ROI 10	1.22 ± 0.37	151.38 ± 22.70

The interviews revealed the following observed patterns:

Participant #1: This participant discovered after several rounds of feedback that several transactions on the same day of the Transaction History (panel #2) were a good indicator of fraud, so this panel was judged as a very good indicator and was looked at a lot. The participant also noted that it was good to look at the purchase amount (panel #1) and the average amount for the purchase category (panel #8). The participant further noted that it was predictive whether the purchase was a card present or card not present purchase (panel #3), which was coincidence and not injected as part of the fraud patterns; and that if the time of day was 'fishy', it was also likely fraud (a pattern not injected either, but merely reflecting coincidence).

Participant #2: At the very end of this participant noticed that it would be predictive to pay attention to how many transactions were made on the day of the purchase; the participant had for some time read the history in reverse, before noticing that the axis was labelled as days into the past. The participant also noticed that very high transactions of around £2000 were completely over board in context of the transaction history. The participant also considered 'card not present' (panel #3) to be a strong indicator of fraud, which did not match the injected fraud pattern; neither was entering the CVV code (panel #4) on its own. The participant also mentioned differences between the country where the card was issued and where the card was used, which bore no correlation to fraud.

Participant #3: This participant first checked whether the user had made several transactions over the past days (panel #2) and compared the average purchase category amount (panel #8) and the amount used for that purchase (panel #1). If the amount was very high, it was also checked whether the card was present or not present (panel #3). Other checks included whether the card was issued and used in different countries (no injected correlation) and whether the purchase type made sense (no injected

correlation). This participant considered the amount, card present / not present and country information most helpful.

Participant #4: This participant felt that no patterns for fraud were obvious. The first region of interest was the current amount (panel #1). The countries were compared (panel #5 and #6); if countries were the same, the transaction was deemed normal (no injected correlation) unless the purchase value was very high. This participant noticed that in some transactions where countries were the same and the amount matched the average amount (panel #8), when the transaction was tagged as normal it was in fact fraudulent; however, the participant did not notice the several small amounts in the transaction history and did not indicate looking at the history a lot (as data were considered always evenly spread), matching the cumulative gaze data shown above. Also the category was sometimes considered suspicious (no injected correlation).

5.3.5 Summary and outlook

In summary, the pilot study showed that most participants spotted injected fraud patterns, however they also observed several false correlations and no participant spotted that each fraudulent transaction always went through one of the same two merchants. Gaze data showed different attentional weighting of information sources between participants, and response times indicated that participants spent substantial time (more than several seconds, usually around 15 seconds) to make a decision.

5.4 Defining Baseline Performance Metrics

On the basis of the review of the analysis process, we will refine these tasks and measure performance (in terms of speed, accuracy, types of error, information search etc.). Qualitative assessment of the gaze data showed that participants revisited panels and cycled through many panels before making a decision. Current work involves the use of these data to apply the eye tracking metrics proposed for UI evaluation, and we will compare human scanning behaviour to model predictions of optimal scanning created by Xiuli Chen and Andrew Howes. This will provide objective measures of performance which can be used to compare with the observed performance for participants using different forms of visual display.

6. References

- Bhattacharyya, S., Jha, S., Tharakunnel, K. and Westland, J.C. 2011. Data mining for credit card fraud: A comparative study. *Decision Support Systems* 50, 602-613.
- Brooke, J., 1996, SUS: a quick and dirty usability scale. In: P.W. Jordan, B. Weerdmeester, B.A. Thomas and I.L. McLelland (eds) *Usability Evaluation in Industry*, London: Taylor and Francis, 189–194
- Caldeira, E., Brandao, G., Campos, H. and Pereira, A. 2012. Characterizing and Evaluating Fraud in Electronic Transactions. In *Web Congress (LA-WEB), 2012 Eighth Latin American*, 115-122.
- Duman, E. and Ozcelik, M.H. 2011. Detecting credit card fraud by genetic algorithm and scatter search. *Expert Systems with Applications* 38, 13057-13063.
- Hand, D.J. and Blunt, G. 2001. Prospecting for gems in credit card data. *IMA Journal of Management Mathematics* 12, 173-200.
- Hand, D.J., Whitrow, C., Adams, N.M., Juszczak, P. and Weston, D. 2008. Performance criteria for plastic card fraud detection tools. *J. Oper. Res. Soc.* 59, 956-962.
- Krivko, M. 2010. A hybrid model for plastic card fraud detection systems. *Expert Systems with Applications* 37, 6070-6076.
- Leonard, K.J. 1993. Detecting credit card fraud using expert systems. *Computers & Industrial Engineering* 25, 103-106.
- Leonard, K.J. 1995. The development of a rule based expert system model for fraud alert in consumer credit. *European Journal of Operational Research* 80, 350-356.
- Sahin, Y., Bulkan, S. and Duman, E. 2013. A cost-sensitive decision tree approach for fraud detection. *Expert Systems with Applications* 40, 5916-5923.
- Whitrow, C., Hand, D.J., Juszczak, P., Weston, D. and Adams, N.M. 2009. Transaction aggregation as a strategy for credit card fraud detection. *Data Mining and Knowledge Discovery* 18, 30-55.
- Wong, N., Ray, P., Stephens, G. and Lewis, L. 2012. Artificial immune systems for the detection of credit card fraud: an architecture, prototype and preliminary results. *Information Systems Journal* 22, 53-76.

7. Appendix 1 – SUS questionnaires in English and Portuguese

System Usability Scale

© Digital Equipment Corporation, 1986.

	Strongly disagree				Strongly agree
1. I think that I would like to use this system frequently	<input type="checkbox"/>				
	1	2	3	4	5
2. I found the system unnecessarily complex	<input type="checkbox"/>				
	1	2	3	4	5
3. I thought the system was easy to use	<input type="checkbox"/>				
	1	2	3	4	5
4. I think that I would need the support of a technical person to be able to use this system	<input type="checkbox"/>				
	1	2	3	4	5
5. I found the various functions in this system were well integrated	<input type="checkbox"/>				
	1	2	3	4	5
6. I thought there was too much inconsistency in this system	<input type="checkbox"/>				
	1	2	3	4	5
7. I would imagine that most people would learn to use this system very quickly	<input type="checkbox"/>				
	1	2	3	4	5
8. I found the system very cumbersome to use	<input type="checkbox"/>				
	1	2	3	4	5
9. I felt very confident using the system	<input type="checkbox"/>				
	1	2	3	4	5
10. I needed to learn a lot of things before I could get going with this system	<input type="checkbox"/>				
	1	2	3	4	5

Escala de usabilidade do sistema

	Discordo muito				Concordo muito
1. Acho que usaria este sistema frequentemente					
	1	2	3	4	5
2. Achei o sistema demasiado complexo					
	1	2	3	4	5
3. Achei o sistema facil de usar					
	1	2	3	4	5
4. Achei que precisaria de assistência de um técnico para					
	1	2	3	4	5
5. Achei as diversas funções do sistema bem integradas					
	1	2	3	4	5
6. Achei que há demasiadas inconsistências no sistema					
	1	2	3	4	5
7. A maioria das pessoas irá aprender a usar este sistema					
	1	2	3	4	5
8. Achei o sistema muito complicado de usar					
	1	2	3	4	5
9. Senti-me confiante a usar o sistema					
	1	2	3	4	5
10. Preciso de aprender muitas coisas antes de conseguir usar este					
	1	2	3	4	5