Visualisation Tool to Support Fraud Detection

1st Pedro Silva
CISUC - Department of Informatics Engineering, University of Coimbra
Coimbra, Portugal
pedros@dei.uc.pt

2nd Catarina Maçãs
CISUC - Department of Informatics Engineering, University of Coimbra
Coimbra, Portugal
cmacas@dei.uc.pt

3rd Evgheni Polisciuc
CISUC - Department of Informatics Engineering, University of Coimbra
Coimbra, Portugal
evgheni@dei.uc.pt

4th Penousal Machado
CISUC - Department of Informatics Engineering, University of Coimbra
Coimbra, Portugal
machado@dei.uc.pt

***Abstract*—Automatic fraud detection and prevention are challenging problems that have attracted the attention of many researchers in academia and industry. Over the last few years, many improvements have been achieved, especially in predictive models based on Machine Learning. However, a considerable amount of these models only provide a prediction score and a short explanation which may not be enough to make informed decisions. This paper presents a visualization tool that aims to assist fraud analysts in making informed decisions and increase their effectiveness in the detection of fraud. To this end, we designed three visualisation models that apply state of the art techniques to support the analysis of fraudulent transactions. To demonstrate the analytic capabilities and benefits of the proposed tool, we discussed a real use case scenario and conducted user testing with real fraud analysts. Through the feedback from both studies, we were able to conclude that the tool is an asset to facilitate the detection of suspicious events as well to improve the analysis times of the fraud analysts' work process.**

***Index Terms*—Information visualization, Data Analytics, Fraud**

I. Introduction

With high amounts of data being generated every day, it is becoming increasingly difficult, or even impractical, to analyse all the collected information in an efficient and appropriate way. To overcome this, in the area of fraud detection, fraud analysts started to apply Machine Learning (ML) algorithms to support decision making and to assess data efficiently [1]–[3]. However, in domains where the consequences of bad decisions have a strong impact on the good functioning of a system, a single evaluation of a model may be not enough to make well-informed decisions and therefore additional information may be necessary. To overpass this, fraud analysts begin to employ data visualisation techniques [4].

In this paper, we present a three-visualisation tool, which combines various state of the art techniques. With the proposed tool, we aim to detect two types of fraud intrinsically connected—Bot Attack (BA) and Account Takeover (ATO), usually associated with periodic and behavioural patterns. ATO consists of the theft of personal or financial information and the usage of the stolen data to make unauthorised bank transactions or purchases [5]. BA is performed by a software or scripts (i.e., bot) that through automated processes attack e-commerce merchants or financial institutions. Typically, bots perform simple repetitive tasks faster than humans. However, these tasks are performed in a way that can resemble a human interacting with the system and, for this reason, can pass unnoticed.

The data-set provided from a world-leading company in fraud prevention for this project is fully anonymised and contains already annotated both non-fraudulent transactions and fraudulent transactions associated with the two fraud patterns previously described, ATO and BA. Each transaction contains attributes, such as personal details, billing and shipping addresses, IP address, and details associated with the device used in the transaction. Given the different types of data attributes present in our dataset, the goal of our visualisations is to highlight patterns that may exist in the data by combining and revealing relationships between key attributes which are more relevant for the detection of ATO and BA. With our approach, we intend to facilitate the work of the fraud analysts, who base their decisions through the exclusive analysis of tabular data, a process which can be a lengthy and time-consuming. Hence, the proposed visualisation tool aims at assisting fraud analysts in exploratory scenarios, with longer analysis times, by helping them to highlight and recognise new patterns in the data. In more detail, we aim to help the analyst to uncover new cases of fraud, false positives or false negatives, and detect new or more complex fraud cases that were not possible to detect previously. Our visualisations work as an auxiliary tool, to be integrated into their current system to facilitate information crossover and the performance of tasks in parallel. Finally, to test the efficiency of the proposed visualisation tool, we performed a use case study with the provided data-set and conducted a user test with fraud analysts.

The remainder of this article is organised as follows: Section II presents the related work; Section III details our visualisation tool, presents the data and its structure, and describes the functionalities and goals of its three visualisations; Section IV describes the application of the proposed visualisation tool on a real use case; Section V presents the user testing conducted to validate our tool and discusses its results. Lastly, Section VI presents the conclusions and future work.

II. Related Work

For the development of our fraud detection application, taking into account our data-set composition and the types of fraud to be detected, a study of art was conducted as follows: (i) a survey of the area visualisation tools was

2375-0138/21/$31.00 ©2021 IEEE
DOI 10.1109/IV53921.2021.00022

made to understand their functionalities and main aspects to be considered in our type of fraud detection. (ii) Then, we analysed the main visualisation techniques used for each purpose, namely periodicity detection and patterns in the data distribution.

Bolton and Hand published a survey of available tools used in the area of fraud detection and presented the most used technologies for the detection of the main types of fraud [1]. Later, Kou et al. [3] surveyed existing techniques for identifying the same types of fraud described by Bolton and Hand. Other researchers developed systems for fraud detection, such as VISFAN [6], that supports the analyst with effective tools to discover financial crimes, like money laundering; and EVA [7], an approach for supporting fraud investigation and fine-tuning fraud detection algorithms. In Financevis, a web repository of papers related to the visualisation of financial data, its users can see how the different types of data are most commonly visualised [8]. Among the analysed papers, we found considerable use of heat maps [9], [10], circular representations [11], [12], and graphs [13], [14]. Also, Dilla & Raschke [15] developed a theoretical framework to predict when and how researchers should use visualisation techniques to detect fraudulent transactions. Among the different visual approaches, we focused on heat maps and circular representations since these techniques revealed to be often used to detect patterns and periodic events in data, two principal objectives of our tool. Interaction techniques, such as the visual results exploration, selection and reconfiguration, data filtering and comparison of multiple queries results are also applied in our tool. Interactive visualisations also have used multiple linked views, connected by linking and brushing techniques, to provide visual connection and differentiation of the displayed elements in different views [16]. The application of multiple linked views can also be found in the field of fraud detection. A three visualisation system aimed at fraud detection is proposed by Argyriou and Symvonis [17]. Their system supports the identification of periodic activity in time-oriented data involving pairs of entities. Additionally, Argyriou et al., propose a similar approach in the field of fraud detection using a spiral representation for the same goals [18]. Other visualisation tools [7], [19], [20] also apply a multiple view design in fraud detection by allowing the representation of data through different perspectives.

A. Heat Maps

Heat maps have also been used in systems focused on fraud detection [21]–[23]. In this visual approach, higher values are commonly represented by darker colours and lower values by lighter ones. An early example of the application of this technique is the shading matrix developed by Loua to visualise social statistics across the districts of Paris [24]. The advantage of heat maps is that, within a relatively compact display area, this visual technique facilitates the inspection of large amounts of information and the observation of patterns in the data distribution [25]. Zieglerh et al. [10] verified that the use of an even more compact approach, a pixel-based approach, on financial time series data improved insights into the characteristics of assets such as the discovery of patterns, trends and correlations in different markets, strengthening the added value of this visual technique in time series analysis. In the present work, the heat map technique was seen as a compact representation of data for overview analysis, with a higher level of detail. Applying this technique on a calendar-like rectangular area could help to detect periodic patterns and trends on multiple time scales (days, weeks, seasons) simultaneously [26], operating as a first step to guide the analysis.

B. Circular Representations

We found diverse variations of circular representations in terms of design and applications in the detection of periodic events, including with the intent of intrusion and fraud detection systems. An example of these techniques is the spiral visualisation by Carlis and Konstan designed to reveal frequent occurrences for a given data type [27]. In their visualisation, the data is displayed along a spiral, accentuating serial attributes along the spiral axis and periodic ones along the radii. Another approach, by Weber et al., presents a visualisation system for time-series data based on spirals suited to visualise large data-sets and to support the identification of periodic structures for both qualitative and quantitative data types [28]. According to their approach, the spiral is the time axis where the remaining attributes are represented by other geometric symbols. SpiralView, proposed by Bertini et al., also used spirals to identify periodic patterns in the data [29]. Circular approaches are also common in network monitoring [30]–[32] to provide overall views of the data and thus perceive its similarities and periodic patterns. A historical review of this type of visualisations is presented to us in a survey by Draper, Livnat & Riesenfeld [33] where the domains of application, taxonomy and design considerations are proposed. Comparative studies between visualisations that make use of the Cartesian coordinate system and its variants with radial systems were made in [34], [35]. Burch et al. state that radial representations outperform Cartesian ones at data correlation analyses and do not cause the ”blinders effect”, an effect that occurs in Cartesian representations. Follow up works from Burch [36], [37] make use of radial diagrams to represent dynamic graphs to avoid animation problems and to ease the exploration and detection of trends or outliers due to the ease with which the locations of the elements are distinguished and remembered. The positive aspects associated with radial visualisations, such as the ease of memorisation, the non-neglected visualisation areas and the facilitated detection of trends and periodic events, were seen as pivotal to the presented tool use scenarios.

III. Three-visualisation Tool

The visualisation tool is composed by three views, each one representing different time scales and levels of detail: (i) a calendar view, where it is given an overview of the data; (ii) a monthly view, where the transactions of a specific

month are visualised in more detail; (iii) and a detail view, where an attribute-wise analysis can be performed for a set of transactions. The data used in each view results from previously performed queries. For each query, the analyst enters a string and chooses the attribute on which he wants to do the research. The query returns all transactions in which the values of the chosen attribute are equal to the string entered. Additionally, the analyst can introduce a time interval to view only the transactions contained inside that same interval. By navigating throughout the views, from the calendar View to the Detail View, the data is filtered and a new sub-set of data visualised. Each view enriches the analysis by providing additional details about the data.

In the following subsections, we present the data-set, the tasks and design requirements of the work, and the GUI and its components. In the end, we describe the three views and their main functionalities.

A. Data

The data-set provided for this project is fully anonymised and contains the two fraud patterns described previously, ATO and BA. The data-set contains more than 4 million online transactions, previously annotated as fraudulent or non-fraudulent. In this data-set, approximately 3% of the transactions are annotated as fraudulent, which, according to the company's analysts, is common for these types of data-sets. The transactions occurred within a period of four months, from November 2016 to February 2017.

Each transaction contains attributes, such as personal details, billing and shipping addresses, IP address, and details associated with the device used in the transaction. In short, the data-set consists of 32 attributes, of which 30 are categorical and only 2 continuous. All transactions associated to a fraud score and are annotated as fraudulent or non-fraudulent by a complex ML model. For each transaction, the model assigns a score, ranging between 0 and 1000, the latter being the maximum fraud value. Even though the ML model is the first instance of fraud analysis and validation, some transactions within a certain score range still need to be verified manually by the fraud company's analysts. For those transactions, the fraud analyst has to verify its authenticity based only on the given score, its brief explanation, and the transaction details in a tabular form. This task is performed manually and under a short period of time (approximately 30 seconds). As this task is focused on a single transaction, it may be difficult for the analyst to have an overview of all transactions and find relationships among them in the same place, leading to erroneous transaction classifications.With our approach, we aim to improve this aspect by giving a more global and contextualised view of the data.

B. Tasks

Through the partnership with the company's analysts, we were able to define the main objectives and the principal requirements for the developed tool. Given the described characteristics of ATO and BA fraud patterns along with the additional insights provided by the company's analysts, it was possible to define the key points such as, the patterns to look for, the most relevant data attributes and how to relate them. Knowing the analysis rationale of the fraud analysts for this kind of fraud patterns allowed us to establish the necessary tasks and design requirements for our tool.

T1 Identify activity patterns. Detection of atypical behaviours through the observation of the user's activities;

T2 Distinguish transaction types. Identification of two types of transactions, fraudulent and non-fraudulent;

T3 Detect periodic patterns between transactions. Infer periodic activities throughout time;

T4 Examine the similarity between transactions. The user should be able to analyse the similarities and patterns between and within the transactions details to recognise compositions associated with fraudulent and non-fraudulent transactions.

T5 Identify false-positive and false-negative fraudulent transactions. Through the transactions arrangement and the identification of areas associated with fraudulent and non-fraudulent transactions, the user should be able to detect false positives, false negatives or even new fraud case scenarios;

T6 Examine transaction details. The analyst should be able to consult the details of the transactions so that he/she can understand the relationship between its composition and its annotation. This way the analyst can detect attributes associated with fraudulent transactions;

T7 Refine and Compare Results. To enable deeper and more detailed analysis, the tool should allow the user to filter and compare the data being analysed, obtaining more specific and detailed data, without losing the previous visual results.

C. Design Requirements

Given the established tasks to be performed, we have identified 5 design requirements to include in our visualisation tool:

DR1 Visualisation of the transactions' temporal distribution. The tool should allow the analysts to visualise the temporal distribution of the transactions for different logical intervals (year, month, week). This way, they can perceive periodic patterns and outliers at those same temporal levels, indicating possible risk behaviours [T1, T3];

DR2 Distinguish types of transactions. The tool should provide visual arrangements, with a clear distinction between fraudulent and non-fraudulent transactions [T2].

DR3 **Detection of temporal and similarity relationships.** The tool should provide means to find temporal and similarity relationships among transactions, allowing the detection of outliers, false-positives or false-negatives [T3, T4, T5];

DR4 **Overview of transactional details.** The tool should support the analysis of the transaction details and enable their comparison between transactions [T6].

DR5 **Filter and comparison of transaction data.** The tool should allow the user to execute multiple queries and visualise the results. Also, the tool should allow the user to consult the results at any time within the same session [T7].

D. The GUI

The Graphical User Interface (GUI) of the proposed visualisation tool is divided into three main areas: (i) search area, where the user chooses the attributes that will serve as query filters to obtain the desired data (Figure 1.a); (ii) history panel, that contains all the visualised views and allows the user to navigate to any instance within the analysis session (Figure 1.b); and (iii) main visualisation area, that contains the visualisation views and their legends. (Figure 1.c). The search area is always visible in the screen so that the current query can be consulted or a new one executed whenever desired. The history panel is responsible for the navigation between previous results. Each element of this panel contains basic information about the visualisation view, providing a condensed history about the analysis steps made by the user. By clicking on a given element of the panel, the visualisation returns to the corresponding view. This approach is an analogy to browser tabs which are familiar to most participants, making it easier to interpret. The transition between views can also be made by clicking on the arrows placed vertically in the middle of the left and right margins of the main visualisation area (C).

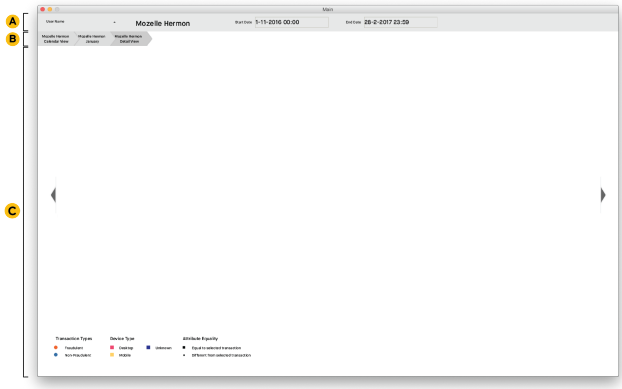

Fig. 1. Main elements of the graphical user interface consisted of (A) search area, (B) History panel, and (C) Main Visualisation area.

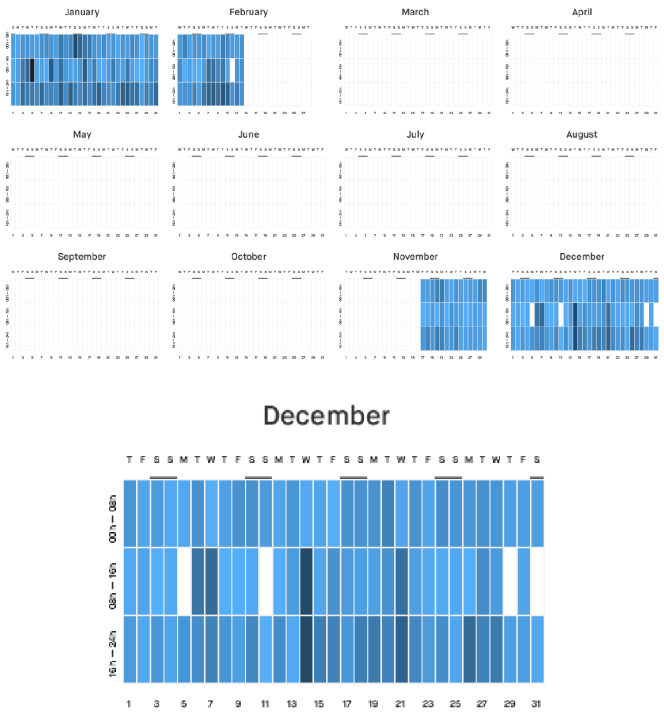

Fig. 2. Calendar view for a given user name. Painted rectangles display associated distribution of the transactions (top). Calendar view in detail for the month of December (bottom).

E. The Calendar View

The calendar view is based on a heat map technique and is structured as a typical calendar as shown in Figure 2. The visualisation is composed of twelve wide rectangular spaces, each one representing one month. Each day of the month is represented by a vertical bar, positioned from left to right in each rectangular space. The bars are further divided vertically into three parts, each one corresponding to an 8 hours interval. These three parts are coloured with a blue tone that varies depending on the number of transactions made in each period of time. The higher the number of transactions for a given hour interval, the darker its blue tone. The aggregated values are globally normalised to enable the comparison of blue tones over the different months. This view enables the analysis of daily, monthly, and annual patterns. It was developed to give an initial overview of the data-set, enabling the analysis at a higher level. To go from the present view to the next, the user should select a specific month and then click on the right arrow to proceed. In Figure 2, it is possible to perceive the transaction behaviour of a specific user, his/her most regular and active periods, and detect visually the outlier time intervals.

F. The Monthly View

The monthly view provides the analyst with the following: (i) a more accurate time distribution of the transactions; (ii) the distinction between fraudulent and non-fraudulent transactions; and (iii) the possibility to relate transactions through

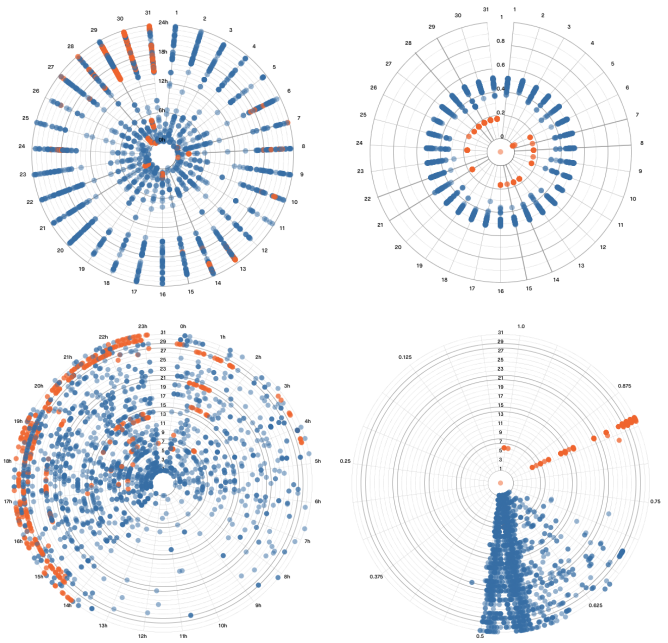

Fig. 3. Radial layout approach for the month of January in temporal arrangement, top-left image, and in similarity arrangement of the transactions, top-right image. Polar layout approach for the same data in temporal arrangement, bottom-left image, and in similarity arrangement of the transactions, bottom-right image.

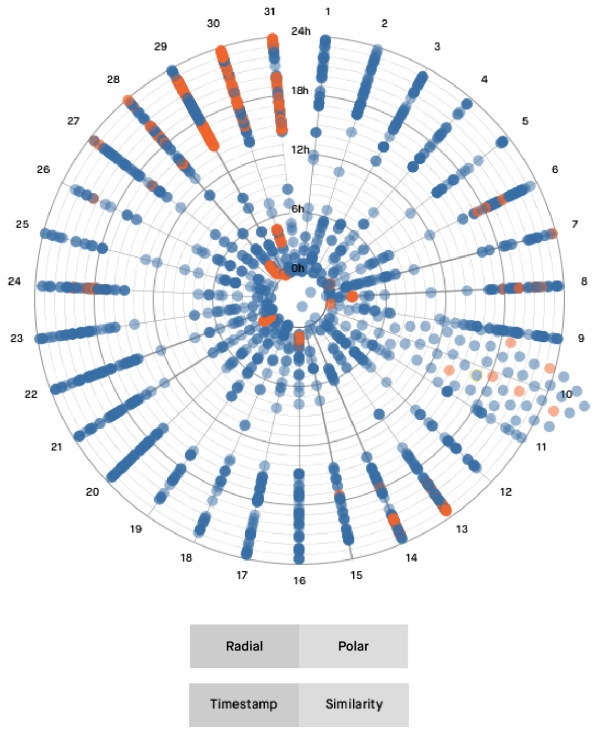

Fig. 4. Hovering a set of overlapping transaction circles located at the 10^{th} axis and the buttons.

their similarity or periodicity. To respond to these analytic objectives, the transactions were arranged on a circular layout.

We defined two layout approaches, radial and polar. Additionally, for each one of these layouts, we defined two types of arrangements to visualise: (i) the distribution of transactions over time; and (ii) the similarity between transactions. In the radial layout, and to visualise the transactions over time, each concentric circle represents an hour and the equally spaced lines arranged radially, represent each day of the month (Figure 3, top-left image). For the second arrangement, to represent the transactions' similarity to a given transaction of choice, 10 concentric circles are drawn to represent the similarity value—the more similar the transactions are, the closer to the centre they will be placed (Figure 3, top-right image). In the polar layout, to represent the transactions over time, each concentric circle represents one day and the hours are mapped in a clockwise direction along the equally spaced lines (Figure 3, bottom-left image). To represent the similarity between transactions, the similarity values are distributed also along the equally spaced lines in a clockwise direction (Figure 3, bottom-right image). For both layout approaches, when in the similarity arrangement, the transaction with which the remaining ones are compared is placed at the centre of the chart by previously clicking on it. This way, the selected transaction is separated from the rest. The user can switch to the desired layout approach or the desired arrangement by pressing the corresponding buttons, placed below the circular visualisation.

The lines representing the weekends are coloured with a darker grey, allowing a faster distinction between weekdays. Each transaction is represented by a filled circle, in blue (non-fraudulent) or orange (fraudulent), and with transparency to detect overlapping positions. Additionally, the user can hover any circle to consult the complete transaction details in the form of an organised list, placed on the right of the main visualisation. This way, the user can create correlations between the temporal or similarity mapping and the attributes that compose the transaction.

Regarding visual occlusion, multiple transactions may occur at the same time or may have an equal similarity value, leading to the overlap of different circles. To overcome this, we apply a mechanism that rearranges overlapping circles perpendicularly to the day line, in both directions. This mechanism is triggered only when hovering any overlapping circles (Figure 4). With this, we aim at preventing visual clutter and allowing the isolated analysis of each transaction in detail.

1) Similarity metric: Due to the diverse nature of the data-set's attributes, we defined a similarity measure that can be applied to mixed data, i.e., categorical and numerical features. For each data type, different similarity measures were defined and combined according to the equation 1. Also, and since each variable belongs to different space values, all the similarity measures were normalised.

$$d_W(i,j) = \frac{\sum_{k=1}^{p} W_k \cdot S_k(x_i, x_j)}{\sum_{k=1}^{p} W_K} \quad (1)$$

In the equation 1, W_k is the weight of variable k in the range $[0, 1]$, $S_k(x_i, x_j)$ is the similarity between x_i and x_j of variable k in the normalised form, and p is the number of variables (15 in the case of our data-set).

Regarding the different similarity measures we defined the following. For continuous variables, $S_k(x_i, x_j)$ equals to $|x_{ik} - x_{jk}|/R_k$, where R_k is the range equal to the difference between the maximum and minimum value present in the set of transactions at hand. For binary variables we defined similarity measure as equal (0) if $x_i = x_j$ or different (1) if $x_i \neq x_j$.

For nominal variables, domain-specific characteristics should be taken into account. Variables such as person's name, address, billing address, country among other similar attributes, were treated as real actions of writing, where people can commit spelling mistakes (e.g., John Doe can be mistakenly written as john Doe). Therefore, the similarity for such nominal variables were defined using Damerau-Levenshtein distance [38], [39], and normalised using the length of the biggest variable, i.e., $S_k(x_i, x_j) = d_L(x_i, x_J)/max(|x_i|, |x_J|)$, where $d_L(x_i, x_J)$ is the Damerau-Levenshtein distance between x_i and x_j.

For another special nominal variable, which is the *email address*, the similarity measure was defined as the combination of the similarity measures for each part of an email address. We considered two constituent parts (e.g., john.doe@example.com): (i) local-part (e.g., john.doe); (ii) domain (e.g., example.com). For the first part (i) we applied a similar metric as for special nominal variables described above. For the second part (ii) the values were treated as equal or different, applying the same logic as for binary variables.

Finally, for yet another special nominal variable, which is the *device*, a different similarity measure was defined. A simple observation of the data made it possible to distinguish three types of devices used in the transactions: (i) portable device platforms (e.g., Android, iOS, etc.); (ii) fixed device platforms (e.g., Windows, Mac OS, etc.); and (iii) others. Having this in mind, we defined similarity $S_k(x_i, x_j)$ to be equal to 0 if $x_i = x_j$, 0.5 if the type of x_i is equal to the type of x_j, and 1 of the type of x_i is different from the type of x_j.

G. The Detail View

For a better understanding of the relationships between transactions, a third view is made available for the analysts. In detail view, it is possible to analyse and compare the transactions attribute-wise. Figure 5 shows an example of the visualisation available in the tool's detail view.

In this view, the transactions are displayed in a grid. Each transaction is represented by a column where each row represents an attribute. The attribute associated with each row is indicated in the list on the left of the grid. To consult the details of a given transaction, the analyst can hover a column and see its details on the right side of the grid. The first column of the grid is, by default, the transaction with which the remaining transactions are compared. Transactions are temporally ordered from left to right. The visual elements encode the comparison between the attributes' values of the different transactions. If the attributes are equal, a black filled rectangle is used. If they are different, a filled rectangle with double the size is used, emphasising such cases. Finally, in the case of missing values, no visual mark is attributed. To emphasise the values and enable a faster interpretation of a given transaction, we used other representations for the attributes: score, fraud, amount handled, and device type. The score and amount attributes are represented by rectangles whose height is mapped according to their values, similarly to a bar chart visualisation. With this approach, higher values will be highlighted by bigger areas. To be visually consistent with the previous views, the fraud attribute is represented by an orange or blue circle to represent a fraudulent or non-fraudulent transaction, respectively. Finally, the device type is represented by a rectangle painted with one of three colours: yellow for portable devices, pink for fixed devices and blue for the remaining type of devices. This decision gives the analyst the possibility to quickly identify the type of device used and to detect unusual recurring changes between them.

In this view, two types of operations can be performed: (i) select the transaction to be compared with the rest; and (ii) highlight transactions of interest according to the attributes. For the first operation, the selection is made through the click on the checkbox element placed above the desired transaction column, which will update the grid representation of the remaining transaction attributes (Figure 6). Then, the second operation allows the selection of all the transactions which have a certain attribute in common. This operation enables the user to analyse and infer relations among the transactions that are interconnected through the selected attribute. To do so, the analyst has to click over the transaction column and row of the desired attribute. With this, the display is updated highlighting the transactions of interest. Also, a reset button is provided to return to the default state, displaying all the transactions. Finally, due to possible scale-related constraints, the size of the grid area is fixed and a horizontal scroll is implemented to enable the navigation between the various transactions.

IV. Use Case

To demonstrate the effectiveness of our approach, we have conducted a use case to describe a possible use scenario using the provided data-set.

A. Analysis Rationale

Since the focus of our exploratory tool is on the comparison between attributes, we started our exploration by searching for a unique identifier, i.e. the User ID attribute. To do so, we selected the attribute of interest in the drop-down list of the search area and typed a user ID. As a result, a set of transactions distributed over the months, was presented to the user in the calendar view. In this view, we easily detected an apparent daily behavioural pattern in November

Fig. 5. Detail visualisation for the month of January following the same search of Figure 2

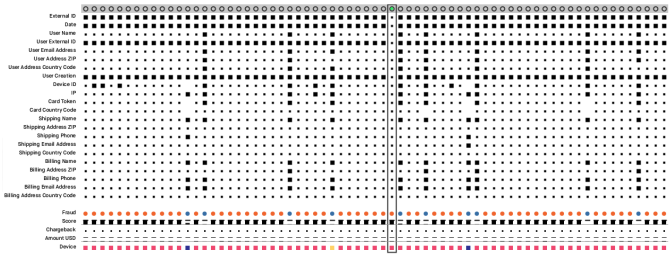

Fig. 6. Detail view of the 4th week of the January following the same search of Figure 3

and December, where an evenly distributed activity over the months is clear. By contrast, our attention was drawn to a set of days in January represented with darker colours at the end of the month (Figure 7, top image). By selecting the month of January and moving to the monthly view, we could see the presence of fraudulent transactions, which occurred mainly at the beginning and end of the days. By applying the similarity arrangement when in the radial layout, we detected a clear difference between the position of some fraudulent (closer to the centre) and non-fraudulent transactions (further away of the centre) (Figure 7, the middle-right). We also could detect the presence of fraudulent transactions among the area occupied mostly by non-fraudulent transactions which may indicate the presence of false-positives or false-negatives.By consulting the organised list, we were able to perceive that those fraudulent transactions had similar attributes to the ones considered as non-fraudulent. To further understand this distinction, we analysed the data through the detail view. In this view, we could perceive that a small group of IP addresses were common to the fraudulent transactions, as illustrated at the bottom of Figure 7.

To continue this exploratory analysis, a new search was performed using one of the most found IP addresses. Starting again with the calendar view, we could see darker areas in the first and third parts of multiple days bars in January. In the monthly view, we noticed a large number of both fraudulent and non-fraudulent transactions that were connected to the introduced IP address (Figure 8). In addition to the already known user associated with the fraudulent transactions, through a fast search by hovering the transactions circles, new users were found (Figure 8).

To better understand the connections between these users, we proceeded with our analysis with the detail view. At this stage, we verified that all the users associated with fraudulent transactions had consecutive ZIP codes (i.e., 100, 101, 102). However, we also noticed that a set of non-fraudulent transactions was also related to the same new users, but with different consecutive ZIP codes. Both findings are usual signatures of an ATO where the robbed user details are used but with some small changes or even omissions (Figure 9).

In terms of fraudulent transactions, the main difference to non-fraudulent transactions appeared to be related to the absence of values for some attributes, which can indicate that fraudulent transactions consistently contain empty fields regarding personal data. At this point, a set of other users and ZIP codes can be flagged as cases of risk, due to the unusual case of consecutive ZIP codes being associated with the same IP address. Consequently, the analyst may now apply a refined search for those ZIP areas to infer a possible hacked set of identities or a fraudulent operation in specific geographic areas. In this sense, a new search for one of the ZIP codes was carried out. Again, we have found that several new fraudulent transactions were associated with the chosen ZIP codes. Among them, a new set of users were found as well.

The scenarios of an absence of attribute values in the data or similarities between attributes, like similar ZIP codes, are common cases in the types of fraud we intent to detect. With this use case it was possible to: (i) exemplify the usage of the proposed system in a real situation; (ii) perceive the type

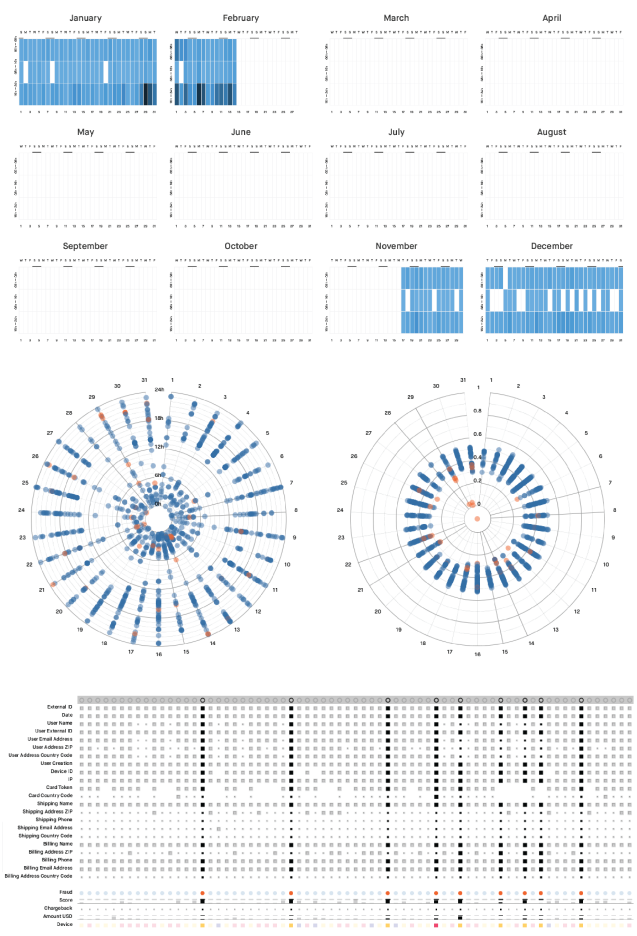

Fig. 7. Calendar view for the User ID search (top), monthly view (middle) and detail view (bottom).

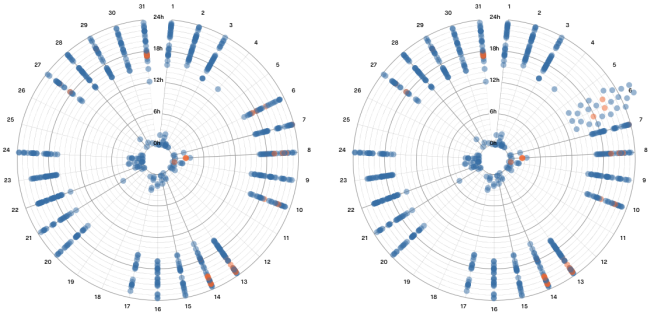

Fig. 8. Monthly visualisation of January in temporal arrangement for the input IP address, image on the left, and with hover over the axis of 5th day, image on the right.

of insights that can be discovered and explored, even with a reduced number of data; and (iii) demonstrate that potential cases of fraud can be detected.

B. Discussion

With this study, we concluded that the tool meets its main objectives and provides a more insightful understanding of the

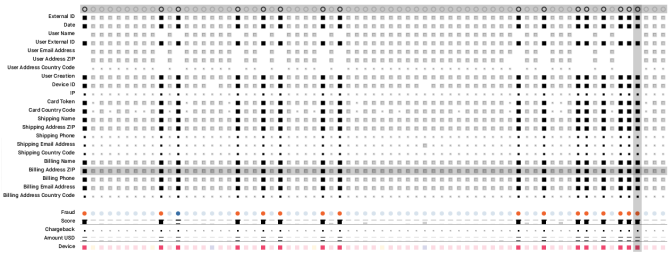

Fig. 9. Detail view highlighting fraudulent transactions with consecutive ZIP codes.

data, corroborating the added value of a multiple visualisation tool. With the presented case study we have shown how visual outliers in the data representation may lead the analysis rationale, as happened in the Calendar View. Also, further explorations may arise due to emerging data relationships, such as false positives or negatives, a scenario that occurred in the monthly view, when organising the transactions by their similarity. These relationships between transactions can be better understood by exposing the similarities between the values of their attributes, possible trough the Detail View.

V. Validation

To evaluate the developed tool, we conducted two different user tests to evaluate different aspects of the tool. In the first, we aimed to assess the ability of the tool to detect fraud patterns and evaluate the tool's interface. After improving the tool with the conclusions and feedback obtained from the first user test, we conducted the second test with the company's fraud analysts to assess the tool's performance through a simulation of a real fraud detection context.

A. User Test 1—non Fraud Analysts

The first user test was performed with a set of static images of the three views, without any sort of interaction, since we wanted to focus on the readability and effectiveness of the visualization models rather than on their interaction mechanisms. Before each test, the context of use and purpose of the tool was described and a general description of its three views was made using a set of screenshots. The same tasks were performed by 15 users, with diverse backgrounds and levels of expertise in visualisation tools, but without any knowledge in fraud detection.

For the calendar view, the users were asked to identify the daily periods, as well as the months, with the highest number of transactions. With this, we could evaluate the heat map approach at global and local levels, i.e. the efficiency of the heat map in what regards the identification and comparison of patterns among months, as well as within each day of the month, respectively.

In the monthly view, we intended to study whether or not the users could detect any type of periodic patterns and interpret the similarity arrangements. For this purpose, we showed two images in the timestamp arrangement, one with periodic transactions and the other without. With this,

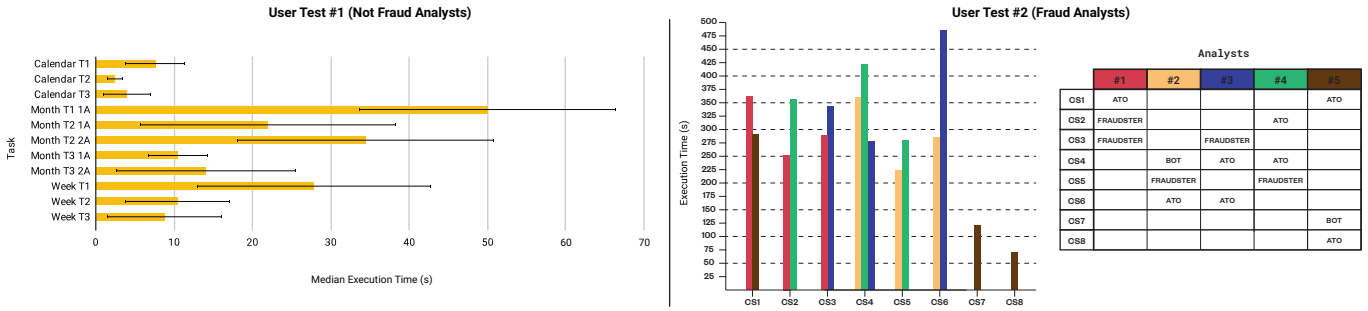

Fig. 10. Median execution time and standard deviation for each view tasks for the first user test(left). Execution time for each case study and corresponding user classification (right).

we aimed to verify whether the users could find periodic patterns. With the similarity arrangement, we asked the users to: (i) identify which transactions were the most similar to the transaction with which they were being compared; and (ii) if they could perceive a difference between the fraudulent and non-fraudulent transactions regarding their positions. These questions were made for both layout approaches to understand which one is the most appropriate layout, depending on the intended analysis. The order of presentation of the images was alternated to understand whether switching between the two layout approaches and their two arrangements was easily decoded and interpreted by the users.

For the detail view, the tasks were related to the interpretation of the attributes, specifically their multiple representations. The users were asked to identify the most changed attributes among transactions and which transactions present a given attribute different from the transaction being compared.

On the left side of Figure 10, we show the medians and standard deviations for the execution time of each task. As one can perceive, the tasks of the calendar view have reduced execution times as well as small standard deviations, with 100% assertiveness rate for all tasks. This indicates that the heat map approach, as well as the day partition, were easily interpreted by the users. For the monthly view, we obtained higher execution times since the tasks were more subjective. Comparing both layout approaches, the radial layout showed the best results in detecting daily periodicity, while the polar layout outperforms the first one in perceiving hourly periodicity. This demonstrates that placing the transactions along the circumference, in a polar layout with a temporal arrangement, improves the readability of periodic events over the radial layout for the same arrangement. In the detail view, it is clear that the analysis of the visualisation requires higher attention of the user given the number of visual elements. Nevertheless, the tasks were executed with 100% of success since the users quickly distinguished the multiple visual representations and could focus their attention on the necessary elements to finish the given tasks correctly.

B. User Test 2—Fraud Analysts

As previously mentioned, we conducted a second user test to assess the tool performance and suitability to achieve its established objectives. The test was conducted with five fraud analysts at the company with an average experience of five years and very little or no contact with visualization tools. Similarly to the first user test, the entire tool was presented to the fraud analysts before the test itself. All analysts visualised the entire activity of three anonymous users previously selected from a group of eight. The anonymous users were selected according to our perception of their different behaviours. Through our analysis, we detected three users which may have suffered an ATO (CS1, CS2, CS6), two users which may have suffered bot attack (CS4, CS7), and with three other users we could not classify them as trustworthy or fraudulent given the composition of some transactions (CS3, CS5, CS8). For the three selected users in each test, the analysts were asked to freely explore the tool and categorise the type of user they were analysing (e.g., trusted user or fraudster and, if possible, the type of fraud performed). The analysts were encouraged to share their rationale and thoughts out loud so that we could better understand how they used the tool, namely which components they relied on to support their decisions and the reasons behind the actions carried out during their exploration. The duration of each analysis was timed.

Although we could not make a direct comparison between our tool and the tools commonly used by the analysts (due to limited access to the data for security and privacy reasons), the analysts reported that the time needed to analyse the users data was perceived as shorter when compared to their current tools (i.e., spreadsheets and web databases) to perform similar tasks. The analysts also reported that their analysis process was facilitated by the tool. One analyst also referred that the similarity metric was especially suitable for bot fraud detection and the detail view was the most suitable for detecting ATO patterns, since this type of fraud is closely linked to changes in attributes between transactions.

The average time of analysis was approximately 5 minutes and, in many cases, the final decision was reached before the analysts finished their analysis process. The analysis times and final decisions of the analysts can be consulted in Figure 10 (right). The calendar view, although it was referred to be functional by the analysts, was less used since the analysts usually resort to shorter time intervals of data when making

more detailed analyses. Although the analysts had some initial difficulties in reading the monthly view visualization, with the additional mental effort when switching between layouts, all of them referred to it as very promising view to detect patterns such as periodicity and activity speed. However, the representation of large quantities of transactions in the monthly view can hinder this type of pattern analysis as well as make it difficult to consult individual transactions. This is a difficulty that is solved with the detail view, which through a much more controlled exploration between transactions, allows to perceive changes in attributes between transactions. All analysts suggested that in addition to the ability to compare all transactions to a given selected transaction, it would be interesting to also compare two consecutive transactions.

The interconnected visualisations proved to be a good approach, as they function as complementary representations to each other, allowing to draw more confident conclusions. For example, one analyst quickly concluded that the user under analysis had suffered an ATO only through the calendar view. Nonetheless, the analyst continued the exploration with the monthly and detail views to better analyse the transactions, namely their annotated classification. Another example of the added value of combined views could be perceived when an analyst initially considered the user activity to be legitimate, but stated that the detail view was decisive in concluding that it was an ATO fraud (due to the constant changes of Device ID that until then would be very difficult to detect).

C. Discussion

The user test results revealed that the visualisation techniques used in the different views turned out to be useful when combined for a better understanding of the data and efficiency in response to the imposed tasks. In situations where there is a lot of activity in a given month, the data exploration can be hampered due to excessive visual cluttering (monthly view) or due to a very long scroll (detail view). A solution may involve allowing the users to more precisely define the time intervals they want to explore (e.g. weeks or a varying number of days). Based on the analysts' feedback, another improvement for the detail view is to add the possibility of comparing transactions consecutively. The analysts revealed that although the tool's learning curve may be a little longer as it presents less common approaches (monthly and detail view), after this extra initial effort, the tool proves to be quite useful and intuitive to use. It is also important to note that the analysts mentioned that they generally use data-sets with time intervals longer than four months to be able to classify these types of fraud with confidence, which may have somehow conditioned their assessments. At the same time, the tool has enabled, in several cases, to reach these confidence levels with fewer data.

VI. Conclusion and Future Work

In this paper, we presented a visualisation tool that incorporates three interconnected views, each one displaying different levels of detail and pursuing different goals. The three visualisations operate as a decision-supporting tool once they enable the user to make more informed and accurate decisions when in exploratory scenarios. The three views were designed based on the tasks and design requirements defined in collaboration with fraud detection experts. Their design and interactions have also been thought based on the state of the art analysis and in order to respond to the user needs in the most intuitive and efficient way possible. With the calendar view, the analyst can obtain a first insight about the data distribution over the year. Then, the monthly view provides the analyst with more accurate details for each transaction within a specific month and distinguishes fraudulent and non-fraudulent transactions. Finally, the detail view allows a more detailed attribute-wise analysis of the transactions.

We also conducted a use case to demonstrate the functionalities of the presented visualisation tool in a real scenario of fraud detection, revealing the advantages of the different visualisations. With this, we were able to show the value of the proposed tool in terms of helping to understand the relationships between transactions and their attributes in fraudulent transactions. This enables the analysts to make more informed decisions and, if necessary, reclassify transactions.

As future work, regarding the tool functionalities, we intend to allow the definition of time intervals for the visualisations and the comparison between multiple views simultaneously. With this, the user will be able to make more accurate and appropriate analysis for the desired subset of data without the need to perform a new query. Regarding the data analysis, the addition of statistical values indicating, for example, the number of different values for each transaction attribute, may provide insightful aspects that would otherwise be quite difficult to detect, specially when analysing large data-sets.

Acknowledgment

This work is funded by national funds through the FCT—Foundation for Science and Technology, I.P., within the scope of the project CISUC—UID/CEC/00326/2020 and by European Social Fund, through the Regional Operational Program Centro 2020. The first author is funded by the Fundacao para a Ciˆencia e Tecnologia (FCT), Portugal under the grant SFRH/BD/144283/2019

References

[1] R. J. Bolton and D. J. Hand, "Statistical fraud detection: A review," *Statistical science*, pp. 235–249, 2002.

[2] T. Fawcett and F. Provost, "Adaptive fraud detection," *Data mining and knowledge discovery*, vol. 1, no. 3, pp. 291–316, 1997.

[3] Y. Kou, C.-T. Lu, S. Sirwongwattana, and Y.-P. Huang, "Survey of fraud detection techniques," in *Networking, sensing and control, 2004 IEEE international conference on*, vol. 2. IEEE, 2004, pp. 749–754.

[4] E. W. Ngai, Y. Hu, Y. Wong, Y. Chen, and X. Sun, "The application of data mining techniques in financial fraud detection: A classification framework and an academic review of literature," *Decision support systems*, vol. 50, no. 3, pp. 559–569, 2011.

[5] B. Baesens, V. Van Vlasselaer, and W. Verbeke, *Fraud analytics using descriptive, predictive, and social network techniques: a guide to data science for fraud detection*. John Wiley & Sons, 2015.

[6] W. Didimo, G. Liotta, F. Montecchiani, and P. Palladino, "An advanced network visualization system for financial crime detection," in *2011 IEEE pacific visualization symposium*. IEEE, 2011, pp. 203–210.

[7] R. A. Leite, T. Gschwandtner, S. Miksch, S. Kriglstein, M. Pohl, E. Gstrein, and J. Kuntner, "Eva: Visual analytics to identify fraudulent events," *IEEE transactions on visualization and computer graphics*, vol. 24, no. 1, pp. 330–339, 2018.

[8] M. Dumas, M. J. McGuffin, and V. L. Lemieux, "Financevis. net-a visual survey of financial data visualizations," in *Poster Abstracts of IEEE Conference on Visualization*, vol. 2, 2014.

[9] C. S. Merino, M. Sips, D. A. Keim, C. Panse, and R. Spence, "Task-at-hand interface for change detection in stock market data," in *Proceedings of the working conference on Advanced visual interfaces*. ACM, 2006, pp. 420–427.

[10] H. Ziegler, T. Nietzschmann, and D. A. Keim, "Visual analytics on the financial market: Pixel-based analysis and comparison of long-term investments," in *Information Visualisation, 2008. IV'08. 12th International Conference*. IEEE, 2008, pp. 287–295.

[11] S. T. Lei and K. Zhang, "A visual analytics system for financial time-series data," in *Proceedings of the 3rd International Symposium on Visual Information Communication*. ACM, 2010, p. 20.

[12] M. Ankerst, D. A. Keim, and H.-P. Kriegel, "Circle segments: A technique for visually exploring large multidimensional data sets," in *Visualization*, 1996.

[13] T. Dwyer and P. Eades, "Visualising a fund manager flow graph with columns and worms," in *Information Visualisation, 2002. Proceedings. Sixth International Conference on*. IEEE, 2002, pp. 147–152.

[14] M. L. Huang, J. Liang, and Q. V. Nguyen, "A visualization approach for frauds detection in financial market," in *Information Visualisation, 2009 13th International Conference*. IEEE, 2009, pp. 197–202.

[15] W. N. Dilla and R. L. Raschke, "Data visualization for fraud detection: Practice implications and a call for future research," *International Journal of Accounting Information Systems*, vol. 16, pp. 1–22, 2015.

[16] D. A. Keim, "Information visualization and visual data mining," *IEEE transactions on Visualization and Computer Graphics*, vol. 8, no. 1, pp. 1–8, 2002.

[17] E. N. Argyriou and A. Symvonis, "Detecting periodicity in serial data through visualization," in *International Symposium on Visual Computing*. Springer, 2012, pp. 295–304.

[18] E. N. Argyriou, A. A. Sotiraki, and A. Symvonis, "Occupational fraud detection through visualization," in *Intelligence and Security Informatics (ISI), 2013 IEEE International Conference on*. IEEE, 2013, pp. 4–6.

[19] R. A. Leite, T. Gschwandtner, S. Miksch, E. Gstrein, and J. Kuntner, "Visual analytics for fraud detection: Focusing on profile analysis." in *EuroVis (Posters)*, 2016, pp. 45–47.

[20] ——, "Visual analytics for event detection: Focusing on fraud," *Visual Informatics*, vol. 2, no. 4, pp. 198–212, 2018.

[21] E. N. Argyriou, A. Symvonis, and V. Vassiliou, "A fraud detection visualization system utilizing radial drawings and heat-maps," in *Information Visualization Theory and Applications (IVAPP), 2014 International Conference on*. IEEE, 2014, pp. 153–160.

[22] R. Chang, A. Lee, M. Ghoniem, R. Kosara, W. Ribarsky, J. Yang, E. Suma, C. Ziemkiewicz, D. Kern, and A. Sudjianto, "Scalable and interactive visual analysis of financial wire transactions for fraud detection," *Information visualization*, vol. 7, no. 1, pp. 63–76, 2008.

[23] S. McKenna, D. Staheli, C. Fulcher, and M. Meyer, "Bubblenet: A cyber security dashboard for visualizing patterns," in *Computer Graphics Forum*, vol. 35, no. 3. Wiley Online Library, 2016, pp. 281–290.

[24] T. Loua, *Atlas statistique de la population de Paris*. J. Dejey & cie, 1873.

[25] L. Wilkinson and M. Friendly, "The history of the cluster heat map," *The American Statistician*, vol. 63, no. 2, pp. 179–184, 2009.

[26] J. J. Van Wijk and E. R. Van Selow, "Cluster and calendar based visualization of time series data," in *Proceedings 1999 IEEE Symposium on Information Visualization (InfoVis' 99)*. IEEE, 1999, pp. 4–9.

[27] J. V. Carlis and J. A. Konstan, "Interactive visualization of serial periodic data," in *Proceedings of the 11th annual ACM symposium on User interface software and technology*. ACM, 1998, pp. 29–38.

[28] M. Weber, M. Alexa, and W. Müller, "Visualizing time-series on spirals." in *Infovis*, vol. 1, 2001, pp. 7–14.

[29] E. Bertini, P. Hertzog, and D. Lalanne, "Spiralview: towards security policies assessment through visual correlation of network resources with evolution of alarms," in *Visual Analytics Science and Technology, 2007. VAST 2007. IEEE Symposium on*. IEEE, 2007, pp. 139–146.

[30] Y. Zhao, F. Zhou, X. Fan, X. Liang, and Y. Liu, "Idsradar: a real-time visualization framework for ids alerts," *Science China Information Sciences*, vol. 56, no. 8, pp. 1–12, 2013.

[31] F. Zhou, W. Huang, Y. Zhao, Y. Shi, X. Liang, and X. Fan, "Entvis: a visual analytic tool for entropy-based network traffic anomaly detection," *IEEE computer graphics and applications*, vol. 35, no. 6, pp. 42–50, 2015.

[32] Y. Shi, Y. Zhao, F. Zhou, R. Shi, and Y. Zhang, "A novel radial visualization of intrusion detection alerts," *IEEE computer graphics and applications*, vol. 38, no. 6, pp. 83–95, 2018.

[33] G. M. Draper, Y. Livnat, and R. F. Riesenfeld, "A survey of radial methods for information visualization," *IEEE transactions on visualization and computer graphics*, vol. 15, no. 5, pp. 759–776, 2009.

[34] M. Burch, F. Bott, F. Beck, and S. Diehl, "Cartesian vs. radial–a comparative evaluation of two visualization tools," in *International Symposium on Visual Computing*. Springer, 2008, pp. 151–160.

[35] M. Adnan, M. Just, and L. Baillie, "Investigating time series visualisations to improve the user experience," in *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 2016, pp. 5444–5455.

[36] M. Burch and S. Diehl, "Timeradartrees: Visualizing dynamic compound digraphs," in *Computer Graphics Forum*, vol. 27, no. 3. Wiley Online Library, 2008, pp. 823–830.

[37] M. Burch, M. Hoferlin, and D. Weiskopf, "Layered timeradartrees," in *2011 15th International Conference on Information Visualisation*. IEEE, 2011, pp. 18–25.

[38] F. J. Damerau, "A technique for computer detection and correction of spelling errors," *Communications of the ACM*, vol. 7, no. 3, pp. 171–176, 1964.

[39] V. I. Levenshtein, "Binary codes capable of correcting deletions, insertions, and reversals," in *Soviet physics doklady*, vol. 10, no. 8, 1966, pp. 707–710.