IEEE Access
Multidisciplinary Rapid Review Open Access Journal

Received September 9, 2019, accepted September 30, 2019, date of publication October 21, 2019, date of current version November 1, 2019.

Digital Object Identifier 10.1109/ACCESS.2019.2948490

Anomaly Detection Using Pattern-of-Life Visual Metaphors

JASSIM HAPPA[1], THOMAS BASHFORD-ROGERS[2], IOANNIS AGRAFIOTIS[3], MICHAEL GOLDSMITH[3], AND SADIE CREESE[3]

[1]Information Security Group, Royal Holloway, University of London, London TW20 0EX, U.K.
[2]Department of Computer Science and Creative Technologies, University of the West of England, Bristol BS16 1QY, U.K.
[3]Department of Computer Science, University of Oxford, Oxford OX1 3QD, U.K.

Corresponding author: Jassim Happa (jassim.happa@rhul.ac.uk)

This work was supported by the UK Defence Science and Technology Laboratory (DSTL).

ABSTRACT Complex dependencies exist across the technology estate, users and purposes of machines. This can make it difficult to efficiently detect attacks. Visualization to date is mainly used to communicate patterns of raw logs, or to visualize the output of detection systems. In this paper we explore a novel approach to presenting cybersecurity-related information to analysts. Specifically, we investigate the feasibility of using visualizations to make analysts become anomaly detectors using *Pattern-of-Life Visual Metaphors*. Unlike glyph metaphors, the visualizations themselves (rather than any single visual variable on screen) transform complex systems into simpler ones using different mapping strategies. We postulate that such mapping strategies can yield new, meaningful ways to showing anomalies in a manner that can be easily identified by analysts. We present a classification system to describe machine and human activities on a host machine, a strategy to map machine dependencies and activities to a metaphor. We then present two examples, each with three attack scenarios, running data generated from attacks that affect confidentiality, integrity and availability of machines. Finally, we present three in-depth use-case studies to assess feasibility (i.e. can this general approach be used to detect anomalies in systems?), usability and detection abilities of our approach. Our findings suggest that our general approach is easy to use to detect anomalies in complex systems, but the type of metaphor has an impact on user's ability to detect anomalies. Similar to other anomaly-detection techniques, false positives do exist in our general approach as well. Future work will need to investigate optimal mapping strategies, other metaphors, and examine how our approach compares to and can complement existing techniques.

INDEX TERMS Cyber security, visualization, anomaly detection, feasibility study, human factors.

I. INTRODUCTION

Data sources (e.g. network packet, CPU, process, RAM logs etc.) form complex *patterns of life* (dependencies and activities on machines), and intrusion detection methods can be used to detect attacks [1], [2]. Misuse detection relies on signatures, and fails when not being able to match signatures to attacks (e.g. zero-day attacks). Anomaly detection relies on a baseline to identify how newly observed activities deviate from the norm. Anomaly detection can suffer from inadequate baselines, with benign behaviour appearing as a matter of concern, and actual concerns to appear within normal tolerance. We postulate that it is possible to use the cognitive processes and lateral thinking of the human mind to detect anomalies using transformations of complex data sources with mapping strategies to create simpler, real-time, navigable and procedural virtual environments.

In this paper, we investigate whether dependencies and activities on complex systems can be represented as simpler metaphors for anomaly detection applications. Inspired by 1980s and 1990s cyberpunk literature and media, we investigate alternatives to visualization paradigms by visualizing anomalies on computer systems as part of pattern of life. Unlike existing literature, our approach shows how activities and dependencies in the technology estate can be transformed into an anomaly detection system using visualization. We then use these metaphors to communicate the pattern-of-life in a manner that speaks to the observer. For instance,

The associate editor coordinating the review of this manuscript and approving it for publication was Giacomo Verticale.

instead of displaying network behaviour in time series visualizations, transform the network behaviour as a modern-day city, with roads, cars, weather and buildings. In this paper, we study whether people can identify anomalies in the machine activities by spotting odd behaviours in a city metaphor as well as a galaxy metaphor.

The hypothesis is that this novel, general approach can complement existing anomaly detection and data visualization techniques by helping users gain insight into complex dependencies and activities in the technology estate with *different associations and perspectives of the same data*. The purpose of this paper is not to investigate whether our approach is more effective than more traditional techniques. Instead, we examine their uses by studying whether this general approach of communicating data insight is capable, and worth investigating further. Presently, we do not have a complete understanding of what makes appropriate transformations in visual metaphors. We believe this is dictated by *technical factors* (i.e. what types of mapping strategies are possible) and *human factors* (i.e. what types of mapping strategies, visualizations and associations make sense to the human observer – we expect these to be dictated by user experience and personal associations).

A. PAPER CONTRIBUTION

This work combines computer graphics, computer security and psychology into a novel method to present insight about anomalies in complex computer systems. These results are novel, since prior research articles have not suggested the use of pattern of life visual metaphors specifically to facilitate anomaly detection. The wider question that we wish to address, and that this paper contributes evidence towards, is the utility of pattern of life visual metaphors to aid anomaly detection. We present evidence towards this, and a discussion on recommendations and what future work into this topic.

We label this general approach a subclass of anomaly detection. We use a single host to demonstrate two designs and their implementation, but our approach can be applied to a domain in which monitoring of real-time data takes place. Our general approach is intended to be a first-pass mechanism to detect anomalous behaviour. More traditional visualizations can be used to investigate detected anomalies further.

We overview our general approach by discussing theoretical and mapping-strategy considerations (Section III), then design and implementation (Section IV) before presenting a study in which we tested its feasibility (Section V). Our paper demonstrates feasibility through in-depth case studies and discussing lessons learned of our investigation (Sections VI and VII). The contributions of our paper include:

1) A novel approach to transforming complex dependencies and activities in systems (pattern-of-life) into simpler visual metaphors from which analysts and lay people can identify anomalous behaviour;
2) A novel approach to limit amount of data necessary to process activities;
3) An implementation of two metaphors using real-time data;
4) A mixed-method use-case study trialling our system on three security researchers, and finally;
5) An in-depth discussion on lessons learned, recommendations and the feasibility of pattern-of-life visual metaphors for anomaly detection more broadly.

Our findings suggest that users can detect suspicious behaviour using our new paradigm, however, much like traditional anomaly detection methods, our method is also subject to false positives. This stems from misinterpretation due to sub-optimal mapping strategies. More research will be necessary to refine the general approach.

II. RELATED WORK

A. ANOMALY DETECTION

Anomaly detection can be broken down into three subclasses of methods, these are: 1) statistical methods (e.g. univariate and multivariate analysis) [2]–[6]: requiring no prior knowledge about normal behaviour. These are typically straightforward to implement, but if attackers are sufficiently capable they may be able to avoid detection altogether. 2) knowledge-based systems (e.g. making use of a finite-state machine, heuristics or rulesets) [5]: encode an understanding about a system's normal activities prior to detection. These are intelligent in their design, but knowledge-based systems may have an incomplete understanding about a monitored system, and may be difficult to update. 3) machine learning methods establish normal using automated approaches with supervised or unsupervised machine learning [5]–[8], but may identify malicious behaviour as normal or normal behaviour as malicious. Minimising false positives and false negatives remains a major challenge in anomaly detection. Glass-Vanderlan *et al.* [2] provide a more up-to-date discussion on the state of host-based intrusion detection.

B. CYBERSECURITY VISUALIZATION

Visualizations often communicate raw-log insights for detection purposes, or output patterns of detection systems so analysts can prioritise their actions. Common visual variables in the literature include [9]–[14]:

- **Colour**, e.g. nominal data such as TCP or UDP traffic, or ordinal data such as severity levels);
- **Position on screen** represents data that are unique, e.g. IDs, IP address, port number, GPS coordinates;
- **Motion, opacity or time on screen** can indicate freshness or throughput intensity of data;
- **Shape** shows data belonging to a same category such as subnets, hosts or type of connections;
- **Size** presents amount of data in the same category;
- **Abstraction**, shows a summary of more complex data using for instance graphs or hierarchies.

Our approach makes use of the aforementioned conventions, however, we serve the metaphor first, and convention second. We deem this important to best address common associations for lay people as well as the security analysts. The inclusion and exclusion criteria for each pattern-of-life visual metaphor will depend on the scope of the metaphor and the system being monitored.

Staheli *et al.* [14], Harrison and Lu [13], and Shiravi *et al.* [12] all present in-depth discussions on the state of the art and trends in cybersecurity visualizations. Staheli et al. make a noteworthy point that from 10 years of VizSec literature: "*(...) no papers used physiological methods for evaluating security visualizations (...). Yet given the sustained focus in the security community on topics such as situational awareness and information overload, existing research in physiological techniques from visualization and human-computer interaction present valuable new dimensions for security visualization evaluation.*" Future work needs to address how well analysts are able to consume and process information provided to them through visualization.

Harrison and Lu [13] note that security visualizations is limited in several areas: scale (many dimensions and throughput), and the lack of explicit representations of network topology and heterogeneous network data, which (we believe) is an area of research that pattern-of-life visual metaphors can address. Finally, Shiravi *et al.* [12] outline the most common types of cybersecurity visualization in a table, including: scatter plots, graphs, tree maps and parallel coordinate plots. We will discuss these different approaches further.

1) DATA MODEL VISUALIZATIONS

Visualizations can be broken down into *data models* (mathematical abstractions such as time series, scatter plots or parallel coordinate plots etc.), or *semantic models* (visualizations with reasoning structures that rely on data to form), with the analyst not necessarily viewing the raw data itself, but the output of some reasoning system with the input being the raw logs. Network statistics and graphs [9]–[11], geographic map overlays [15], plotting of activities (e.g., time series, histograms etc.) are other common techniques.

Visualizations that mainly focus on communicating anomaly detection include [16], which focuses on showing the anomaly data with respects to normal, expected data. Zhang *et al.* [17] provide automated anomaly detection in the observed network activities through probabilistic reasoning of the causal relations in traffic as a radial map. Other examples include: FlowTag [18], NVisionIP [19], IDS RainStorm [20], Spinning Cube of Potential Doom [21], Visual Firewall [22], Netvis [23] among others. The segmented views of activity logs can reveal anomalous behaviour, but not straightforwardly for complex systems where dependencies and expected behaviour is not mapped. Multidimensional data is challenging to visualize for dimensions greater than three, and while glyphs and scatter plots are often used, they often require a steep learning curve [24].

When reviewing metaphors in visualization, the term *metaphor* is typically used to describe a specific type of glyph: a visual variable that is also intrinsically associative [24], such as the use of a broken heart symbol. Metaphoric glyphs simplify more complex ideas through associations. A pattern-of-life visual metaphor on the other hand is here (similarly) a simpler representation (transformation) of a more complex system into a simpler one through the visualization itself: enabled by systematically mapping raw data to visual metaphor variables [25]–[27]. The real-time updating of the visual metaphor forms the pattern of life.

2) SEMANTIC MODEL VISUALIZATIONS

Semantic models are mainly driven by underlying reasoning structures, and can be used to improve understanding for domain experts by relating the raw (low level) data to more high-level interpretations of that same data. Examples include: VisAlert [28] aimed at situational-awareness and decision-making and consists of multiple co-centric circles. Securescope [29] addresses business impact by mapping enterprise units to geographic workspace-locations, Tenable 3D [30] which summarises vulnerabilities in networks, and CyberVis [31] visualizes the potential impact of attacks by showing how Intrusion Detection System (IDS) alerts relate to business processes.

Risch [26] explores the role of metaphors in information visualization and discusses the specific distinctions between analogies and metaphors, as well as how variables can be mapped. Risch also provides an in-depth discussion on the importance of how visuals can "*feel wrong*" and can affect semantic comprehension and abstract reasoning processes.

Ziemkiewicz and Kosara [27] present a discussion on how the structure of a visualization influences how we process it. They discuss how pattern-of-life visual metaphors influence the representation of information in the mind.

Averbukh [32]–[34] explores the notion of visualization languages and metaphors describing them as figurative similarities of concepts, suggesting metaphors will have their own vocabulary, syntax, semantics and pragmatics, that visualization languages are built on the idea of similarities between application domains, that user evaluation is necessary, and finally, that it is necessary to understand the adequacy of the visualization itself. Averbukh also highlight a number of factors that affect interpretation of metaphors, including: psychophysical state (e.g. age, sex, emotional state), knowledge, familiarity, incentives and motivation of the user and finally national and professional culture.

Santos *et al.* [35], Santos *et al.* [36], Santos *et al.* [25], Abel *et al.* [37], Abel *et al.* [38] work present mapping strategies of network information. The work particularly explores visualizing large volumes of network data. Their work emphasise on file system and network, but not on the interactions between several components on host systems.

Brown *et al.* [39] demonstrate use of animations to show network performance. Their work focuses on delivering interactive network data and projects the activities as terrains and vertex edge graphs. These demonstrate patterns of network, but are restricted to the network domain.

Outside the visualization domain, there has been work on mapping network traffic to other domains, particular of interest is the sonification of network security research domain [40]. Finally, outside the security domain, pattern-of-life visual metaphors have for instance been used to communicate multivariate information using a magnet metaphor [41] and software production visualization [42].

C. VISUAL METAPHORS

Using metaphors in visualizations is not novel. Gershon and Page [43] discuss the application of re-framing information into stories and highlight the challenge that scientific and information visualization often lack natural and obvious physical representation by remaining abstract. A key research problem for any visualization designers is identifying how new pattern-of-life visual metaphors can represent information and understanding the analysis tasks they support.

Recently, Latvala *et al.* [44] presented a network monitoring tool in which a 3D fish tank shows different kind of fish that represent network nodes. The authors themselves specify that: "*As this is still a work in progress, more development is needed; especially adding functionality to visualize normal network traffic besides Snort events is crucial*". How the fish move is derived from misuse detection alerts (Snort) only. This is an example of pattern-of-life visual metaphors used for misuse behaviour, as opposed to our approach which focuses on both visualizing normal as well as anomalous behaviour.

Carroll *et al.* [45] study the design of a *cyber satellite navigation* system to improve situational awareness for non-expert users. The core tasks they facilitate include understanding current, past and likely future locations in cyber space. This work focuses on navigational aspects, and not anomaly detection. Our work communicates location in systems as well, but its primary purpose is anomaly detection.

While uses of metaphors is not novel, the use of metaphors to transform complex systems into simpler ones using a mapping strategy to detect anomalies is.

III. THEORETICAL AND PRACTICAL CONSIDERATIONS

At the outset of our research, we identified key areas necessary to consider from a theoretical perspective, before moving onto design and implementation. These include: *Data collection* (Section III-A), identifying what data sources are meaningful to collect; *Data management* (Section III-B), minimising the host-monitoring footprint and amount of data necessary; *Establishing requirements* (Section III-C), identifying what makes a capable pattern-of-life visual metaphor. Designing these is challenging because individual associations need to be general enough so a large audience can understand and make use of them meaningfully; *Metaphor mapping*(Section III-D), identifying strategies to map data on a machine to simpler metaphor.

A. DATA COLLECTION

Many exist that collect data about networks and system activities. Examples of some of the more popular data-capturing tools include: nmap, wireshark, ngrep, top, iotop, traceroute or tcpdump, netflow and Security Information and Event Management (SIEM) tools. When considering the landscape of data sources, we can assume they exist within one of four key layers as described by Legg *et al.* [46], describing the purpose of the machine, users (who they are and what they are doing), logical content (software and their behaviour), or physical content (hardware information and behaviour).

We defined a list glossary that is able to express particular concepts within the context of the pattern-of-life visual metaphors. Below is a list of them we have found necessary to define to date:

- **Property**: an aspect of the Machine that can be measured on a regular or irregular basis, denoted as F_P.
- **Value**: the string or number associated with a Property, denoted as F_V.
- **Event**: a change (delta) in a Value that is recorded as a row in a CSV log file.
- **Category**: a generic term for a physical or logical domain of a Machine. Our system currently supports ten Categories: CPU, MEMORY, HDD, NETWORK, PROCESS, PERIPHERAL, FILE, FOLDER, USER, OTHER, denoted as F_{Ca}.
- **Subcategory**: the sub-domain of a Machine's Category. We make this distinction to classify detailed information about the activity of interest. For instance; RAM activity could related to physical or virtual RAM. Subcategories allow us to make these distinctions, denoted as F_{Cs}.
- **Component**: a physical part necessary to run a Machine. A component can for instance be the CPU, RAM, NETWORK (card), MOTHERBOARD, USB, PERIPHERAL or a USER (who operates a Machine). We deem the distinction between Category and Component necessary to separate the concepts of logical and physical information about a machine. This is denoted as, denoted as F_{Co}.
- **Significance**: a measured Property is *Significant* if the difference between the old and new Value is greater than or equal to a specified threshold. Thresholds can be manually specified or derived computationally on a per Property-basis. Significance answers: *how much a Property needs to change before we deem it to be noteworthy?*
- **Importance**: a Property can differ in terms of how much it intrinsically matters to an analyst (akin to *severity* in intrusion detection system). We consider three levels of priority. For example, occasional CPU spikes are expected to happen, so these are deemed to be of low importance, but any spike of more than 50% are deemed of high significance. Importance answers the question: *what priority does a Property have?*

B. DATA MANAGEMENT

Data management can be split into two separate concerns: the *collecting and storing data* of monitored systems, and *visualizing the stored data*. Our implementation deals with these

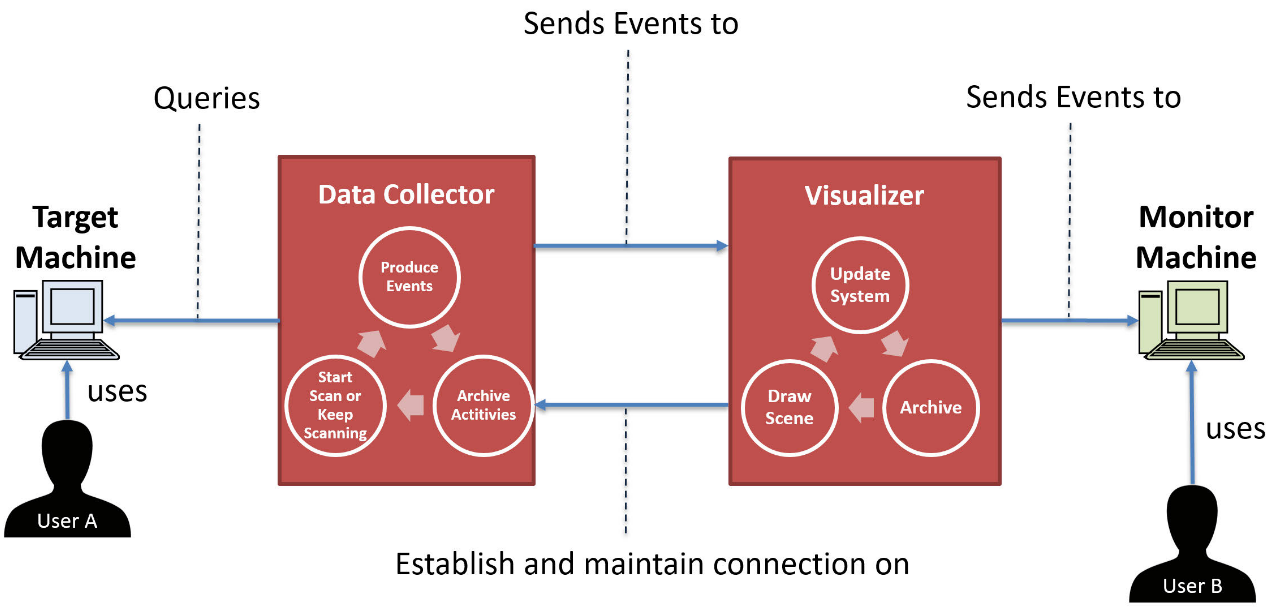

FIGURE 1. High-level view of data flow in our tool, showing the data collector and the visualizer.

two challenges as two separate modules, the *Data Collector* and *Visualizer*, as shown in Figure 1.

Instead of continually monitoring Values (akin to how a system monitor records data), the Significance threshold is used to describe the criteria necessary to trigger an Event (a delta that is deemed significant enough from the last Value). By storing the deltas as Events, as opposed to the raw values themselves, the system monitors key changes on a host.

1) DATA FORMAT

The data collector creates a CSV file on a regular interval (set by an analyst), which is sent to the visualizer. Each column in an event corresponds to the following in an Event (a row in the log file):

- **timestamp** of the Event time.
- **name** of the Property being reported in the Event. Presently, the data collector monitors 120 different types of Properties.
- **category** is the Category reported in the Event.
- **subcategory** is the Subcategory reported in the Event.
- **importance** is the severity of the Event.
- **old_value**: the previously recorded value of a Property.
- **new_value**: the current value of a Property. Together the old and new value make up the delta.

Significance is a static list of thresholds that the analyst is in control over and configured before running the tool. We record both the new and old value to provide additional assurances that the next Event is indeed a follow-up from the previous event. Any inconsistencies are reported, but still visualized, although any Events in the past are dropped. In our study, no inconsistencies were present. Component is derived at the Visualizer from an Event's Category through a static lookup-table.

C. ESTABLISHING REQUIREMENTS

To the best of our knowledge, there are no best practices for scoping and designing new pattern-of-life visual metaphors. As a first attempt, our key requirements included:

- **Driven by existing visualization literature**. Design and implementation of any pattern-of-life visual metaphor must be informed by existing knowledge in psychology (esp. visual perception and cognition), graphics, visualization and usability principles and literature. We built our designs on the literature that relates to usability and understanding data visualization, with an understanding of visual variables theory [47], levels of realism [48], metaphors as visualization concepts [32] and usability principles in mind [49].
- **Compatibility**. Metaphors must first and foremost serve their *natural purpose*, while not deviating from its *anomaly-detection purpose* in any destructive way. If there are any naturally competing forces in the metaphor, we deem them *incompatible*.
- **A mapping strategy** must exist to:
  - **deterministically create and maintain a pattern-of-life** system which exhibit similar characteristics to that of the original system, and ensure that
  - **users are able to predict** the actions in the virtual scene (heuristically and associatively).
- **Associativity**. We assume that metaphor associations either generalisable or personal. The mappings that we propose are intended to work according to the principles found in the aforementioned related-work literature. We also recognise that a single developer cannot identify all common associations. The metaphor scoping exercise mentioned later in this section was (and should be) a team exercise in peer reviewing candidate metaphor

analogies to ensure that associations are appropriate for a large number of people.
- **User-testing is essential**. Assessing performance and refine any design and implementation is vital due to associativity.

1) SCOPING EXERCISES

We identified a wide range of candidates by prioritising intrinsic purpose, usefulness, and building on existing literature enabled us to propose a first set of metaphors for anomaly detection. Metaphors considered, but not selected included:

- **Bloodvessel**: A snippet of bloodvessel in the human body in which different types of cells (white, red) and other items of vital to the human body transported could be metaphors for how much activity we are seeing.
- **Electrocardiogram** (ECG). This metaphor would show the steadiness of a heartbeat pulse (of the Machine). The data collected could be a heartbeat function that updates every heartbeat, or show heartbeats per Component.
- **Electroencephalogram** (EEG): Similar to the heartbeat monitor, but the metaphor shows brainwave activity.
- **Human Body**: A representation of the human body and its overall health over time is expressed by; changes in skin tones, facial expressions, gait, etc. The Metaphor should could express subtle clues about how healthy a system is.
- **Nature**: A virtual environment composited by e.g. trees, landscapes, rain, storms, thunder, grass, winds, wildlife and cabins to represent aspects of a Machine.
- **Transportation Network**: A transport metaphor may be useful to explore activity usage of fixed components, transportation networks would also able to express expected paths and regularities and deviations from those.

D. METAPHOR MAPPINGS

In order to map features to the metaphors, we first have to define the elements which make up a pattern-of-life metaphor. We denote these as *visual variables* [47], and they consist of some aspect of the pattern-of-life visual metaphor.

Each visual variable can be seen as a function which maps features to some quantity Q associated with a rendered representation, for instance the number or movement of objects in the environment. If we start with a space of features $F \in \{F_P, F_V, F_{Ca}, F_{Cs}\}$, then we can define a mapping from feature space to the quantity associated with the visual variable $M : F \longmapsto Q$. We first define an indicator function:

$$\mathbb{I}(F_x, F_{Ve}) = \begin{cases} 1 & \text{if } F_x \in F_{Ve} \\ 0 & \text{otherwise,} \end{cases} \tag{1}$$

where F_x is any feature, and F_{Ve} is the set of allowed features which can be represented by the visual variable. In our context, we also define $\mathbb{I}_1 = \mathbb{I}(F_P, F_{Ve})\mathbb{I}(F_{Ca}, F_{Ve})\mathbb{I}(F_{Cs}, F_{Ve})$, as a shorthand for the case that all conditions are met. Here, M is defined as:

$$Q = \sum_{i=1}^{N} w_i \mathbb{I}_1 g_i(F_V(i)). \tag{2}$$

This is a weighted sum of N values, where the weights obey the following conditions $\sum_{i=1}^{N} w_i = 1, \quad w_i > 0 \quad \forall w_i$. The function $g_i(x)$ maps the value stored in $F_V(i)$ to the range of values that Q requires for display of the visual variable. This function is specific to each visual variable and Value. For example, $g_i(x)$ can be used to re-scale or non-linearly map numerical values, or map text values to numerical values. Note that in many cases a visual variable will perform a 1 : 1 mapping where the mapping simplifies to $Q = \mathbb{I}_1 F_V$.

We demonstrate 1 : 1 and 1 : N mapping in pattern-of-life visual metaphors in this paper, although extensions such as 1 : N ("*One-to-Many*"), N : 1 ("*Many-to-One*"), $N_a : N_b$ ("*Many-to-Many*") and *Mixed Mapping*, where any combination of 1 : 1, 1 : N, N : 1, $N_a : N_b$ may exist.

To illustrate a simple example of a mapping, an application may want to change the brightness of an object (the visual variable) proportional to CPU usage. In this case, a 1 : 1 mapping would be applied where F_{CPU} corresponds to measured CPU usage as a percentage, and Q corresponds to brightness $\in [0..1]$. As there is only one value, $w_1 = 1$ in Equation 2, and the function $g_1(\cdot)$ re-scale the range [0..100] to [0..1], i.e. $g_1(x) = x/100$. Therefore, the resulting mapping is $Q = \mathbb{I}_1 g_1(F_{CPU})$.

Another example is mapping the speed of animation of some object ($Q \in [0..1]$) to an equal weighting of a combination of the number of files modified per second F_{Files} and whether any of a set of certain processes $\{TargetProcesses\}$ are running, i.e. a N : 1 mapping. Here $N = 2$ and $w_i = 0.5$. Therefore, Equation 2 would be written as $Q = 0.5g_1(F_{Process}) + 0.5g_2(F_{Files})$. As $Q \in [0..1]$, the function $g_1(F_{Process})$ needs to map set membership to this range (for example, this could return 1 if any of the $\{TargetProcesses\}$ are running), and $g_2(F_{Files}) = \frac{S \cdot F_{Files}}{1+|S \cdot F_{Files}|}$ is a non-linear mapping with an unknown maximum and scale factor S of the number of files modified; the non-linear mapping is used as there is an unbounded maximum of number of files modified.

IV. DESIGN AND IMPLEMENTATION

To demonstrate the feasibility of pattern-of-life visual metaphor for anomaly detection, we down-selected candidate metaphors to two metaphors: *a city landscape* and a *cluster of galaxies*. In this section we outline how both were designed.

A. DESIGN: CITY

We opted for a US-like city metaphor given their wide open streets and grid-like structures. Our city attempts to match each Property to an analogy (1:1) by specifying:

- **Buildings** represent files (white buildings) and folders (grey buildings) in different districts. Their sizes are determined by file and folder sizes (scaled logarithmically). By default, the buildings represent a selected folder (and its subfolders). Our system does allows for monitoring of the whole OS – at the cost of a much slower start-up. For the purpose of our concept demonstrator, we monitor a pre-selected Documents folder deemed to be *sensitive* in the attack scenarios.

- **Rain** represents CPU and RAM usage $> 50\%$ and lasts until both are below 50%.
- **Snow** represents CPU and RAM usage $> 75\%$ and lasts until both are below 75%.
- **People** are Processes and Users. Once a new process is created a new person would be created and walk down the main street. People keep walking as long as they exist.
- **Cars** are network connections. Once a new connection has been made a new car is created and drive down the main street. Cars are destroyed at the end of the road.

The city landscape demonstrates a 1:1 mapping, where M maps a specific combination of F_P, F_{Ca}, F_{Cs} to a visual variable, where the value $Q = \mathbb{I}_1 g_1(F_V)$, where $g_1(F_V)$ is specific to each visual variable (i.e. Buildings, Rain, Snow, People and Cars) and combination of F_P, F_{Ca}, F_{Cs}.

B. DESIGN: GALAXIES

The galaxy metaphor focuses on spirals that look and behave similarly to how astronomers describe the appearance of the Milky Way galaxy. Each hardware Component is a new spiral galaxy, with stars orbiting the galaxy cluster centre (appearing and fading away) being a metaphor for recent Events on that hardware Component.

The purpose of this mapping is to demonstrate how single Events can exist on several hardware components, giving the end-user some idea about distribution of Events at a hardware level. One important practical decision was taken in the interest of usability: as real stars have distinct temperature colours: red, orange, yellow, white and blue, this limits the number of colours available to use for any photorealistic visualizaion. There are more hardware Components than distinct colours. We therefore chose to use all distinct colours. Below follows a list of metaphor participating actors:

- **Centre**. The existence of a galaxy cluster centre shows that the component exists.
- **Stars**. Each star is a new Event. Its speed relates to Importance. Its colour relate to its corresponding Component.
- **Dwarves**. As Events age they become dwarves to show that they are old. All other properties about the dwarves are the same as the Stars.

The Galaxy metaphor demonstrates a $1 : N$ mapping, where M maps multiple combinations of F_P, F_{Ca}, F_{Cs} to a visual variable, where the value $Q = \sum_{i=1}^{N} w_i \mathbb{I}_1 g_i(F_V(i))$. Again, each $g_i(F_V(i))$ is implemented specific to the visual variable and F_P, F_{Ca}, F_{Cs} combination. Other factors contribute to the galaxy mapping. These basic ruleset includes:

- **The existence of a galaxy centre** shows which Components are being monitored. Each Component is given a different categorical colour (not to be confused with Category). In the examples shown in the figures in this paper for instance: yellow is the hard drive and white is the network card.
- **The radius of each galaxy centre** is determined by the volume of data relating to Component. This means that if many Events are generated that relate to a Component, the larger the centre becomes. It scales logarithmically to prevent it from becoming too large.
- **Each star created** represent a new Event observed.
- **The star lifespan** is 30 seconds as a bright star, then 30 second as a dwarf, before fading out.
- **Distance of the star** to galaxy centre is based on a normalised value for each of the Properties.
- **Speed of the star** indicate Importance, meaning that stars can take one of three speeds with more Important Properties going faster.
- **Colour of the star** indicates its corresponding Component. For instance, the white galaxy is the Ethernet card, but white stars also appear in other galaxies.

C. IMPLEMENTATION

Our tool supports real-time data collection and visualization, as well as a playback feature. The tool consists of two key modules: a *Data Collector* and a *Visualizer* as shown in Figure 1. The Data Collector monitors for significant changes in Values. In our Data Collector, we record: *Timestamp, Category, Subcategory, Property, Importance, Old Value*, and *New Value*. These are recorded as CSV (for playback purposes). The Data Collector was built entirely in Python on top of the *psutil*[1] library python functions and OS-specific calls to retrieving information on running processes and system utilisation.

Events are sent to the Visualizer whose task is to render the mappings. The visualizer was also built in python to manage data to be sent to the visualization via a basic webserver. The rendering rulesets of the scene were written in javascript and implemented using *three.js*[2] and *dat.gui*.[3]

The visualizer uses the *Model-View-Controller* software architectural pattern [50]. Users can navigate the scene with the mouse and keyboard, and swap between metaphors by clicking the GUI in the top-right corner. Our system keeps a per-Machine profile in the Data Model, with which the tool can maintain data from many machines, although in our attack scenarios we have focused our efforts on single-machine pattern-of-life visual metaphors.

Our approach simply appends Events as differences that allow us to update the machine-relevant data. The visualiser does not keep track of the history of Events. In our implementation, we supported monitoring of 120 different Properties across ten different Categories. Both pattern-of-life visual metaphors run simultaneously, and rely on the Data Collector

[1] https://pypi.python.org/pypi/psutil
[2] https://threejs.org/
[3] https://workshop.chromeexperiments.com/examples/gui/

producing Events before invoking the Visualizer. If the Data Collector stops, the Visualizer stops producing more Events, in this means that stars eventually fade away with small galaxies present, and empty streets in the city metaphor.

V. STUDY

A three use-case study was conducted to assess the feasibility of using the tool for anomaly detection using human participants. The study had a focus on three key areas: **detection capability**, **user comprehension of metaphors**, and **usability**. We believe a study to examine the feasibility of the general approach itself is necessary, before conducting extended user studies on specific pattern-of-life visual metaphors and their false positive rate of detection. As the general approach is a novel concept, the key purpose was to identify whether analysts are be able to link abnormal activity in visualizations through metaphor associations and link these to potential malicious host activity. We therefore limited the number of participants in favour of taking significantly more time per participant to obtain in-depth feedback to get some indication about the performance of visual metaphors. From a study design perspective, we attempt to answer three broad questions:

1) *Are participants able to identify attacks correctly?*
2) *What are the participants' opinion about the approach?*
3) *Do participants use the tool as designed, or are there usability issues that prevent visual metaphors from being used as intended?*

A. STUDY DESIGN

The pilot study used a mixed-method cross-sectional study around four main parts: *introduction, training, scenarios* and *reflection*. The study had three researchers with more than five years of experience in cybersecurity research, go through a *demographic questionnaire*, a *video tutorial*, *basic training* to get first-hand experience of using the system, then review *two attack scenarios* were presented to the participants to explore. During the training and attack scenarios we used eyetracking to obtain foveal vision patterns of participants. Finally, after the attack scenarios, a follow-up *questionnaire* and *semi-structured interview* were conducted during which we obtained their feedback. Figure 2 shows the running order of the study. The selection criteria of participants was that they have had to research in security for at least two years prior to the study, with judgement sampling recruitment.

Each of the scenarios were selected pseudo-randomly using a random number generator. The training used one of the three scenarios, while the main scenarios made use of the two remaining scenarios. Participant were not told which attack scenario related to which metaphor, or what attacks to expect. We asked them to give commentary at pre-selected intervals (before, during and after an attack had been executed). Participants sat in front of a single computer and used the tool with a mouse and keyboard.

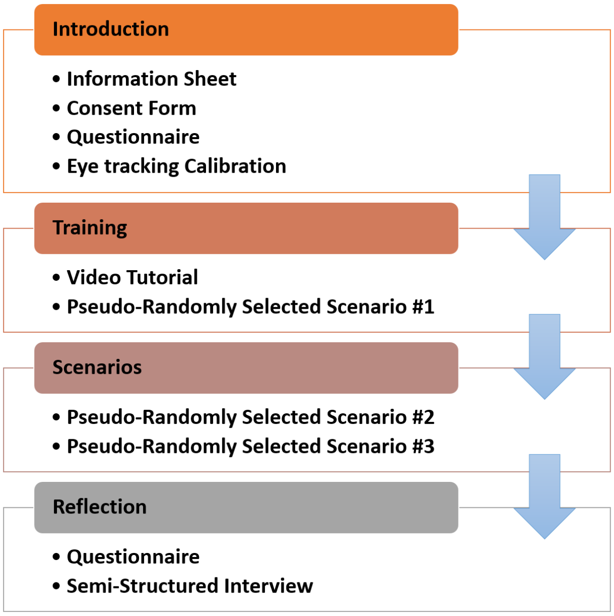
Introduction
• Information Sheet
• Consent Form
• Questionnaire
• Eye tracking Calibration
Training
• Video Tutorial
• Pseudo-Randomly Selected Scenario #1
Scenarios
• Pseudo-Randomly Selected Scenario #2
• Pseudo-Randomly Selected Scenario #3
Reflection
• Questionnaire
• Semi-Structured Interview

FIGURE 2. The running order of the study.

Our assessment relies more on the qualitative approach due to novelty of the pattern-of-life metaphor concept. Eye-tracking was used in our assessment to cross-check participant answers with viewing patterns to provide assurance that their answers matched viewing patterns during analysis of interview notes. A single coder and a single interviewer were involved in asking questions and coding and analysing the interviews.

1) INTRODUCTION

The introduction included reading a project-information sheet, signing a consent form, eyetracking calibration and filling in a demographics questionnaire. The questions have been shortened to fit the table in the paper.

2) TRAINING PERIOD

The Training Period has participants being presented with a video that summarised the tool, and they were given hands-on experience of the tool. The participants were given specific tasks to complete – on how to use the tool and how to interpret the visuals. After the training tasks were completed, the participants were free to familiarise themselves with the tool, and ask any questions they may have about navigating and operating the tool. We eye-tracked and voice recorded them during the whole training period. Once the user felt comfortable with operating our tool, the session ended.

3) SCENARIOS

This Scenarios part was intended to assess how well a participant is able to use our tool. During this part, we assess their ability to detect anomalies. To minimise disruption (and let the participant explore the tool as much as possible

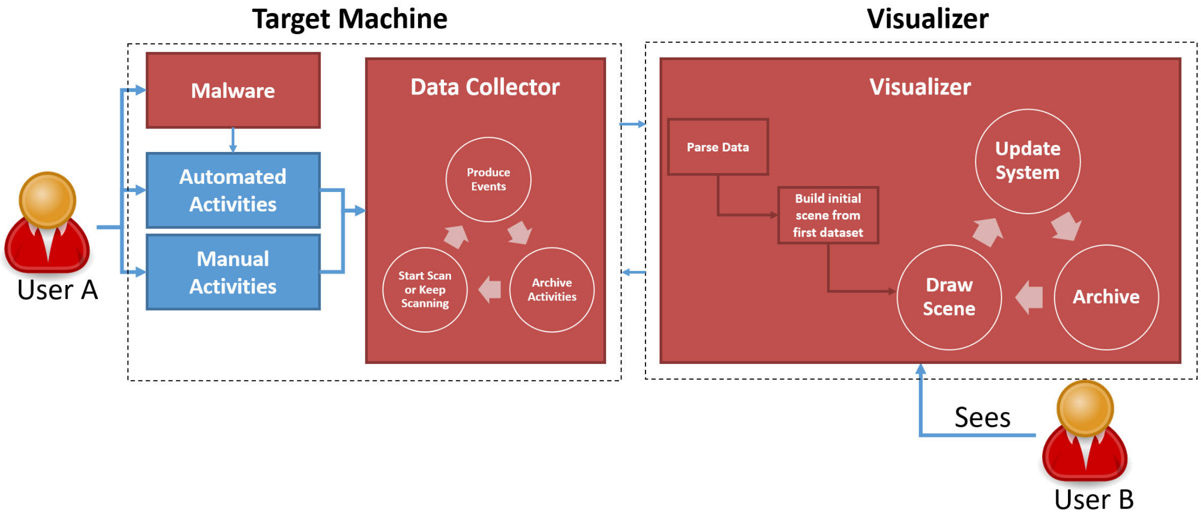

FIGURE 3. Attacks as executed in the study. The attacks were recorded in the Data Collector and replayed to participants during the study.

without disruptions from interviewer), we only asked them to occasionally tell us "*when you find something noteworthy or suspicious, please tell us, and point out what that is in the visuals*". We took note of those observations, and asked about them later in the reflection period. Training and attack scenarios were recorded with voice recording, screen capturing as well as (non-invasive) eyetracking [51]. Due to the novelty of the concept of pattern-of-life visual metaphors, we deemed it necessary to employ a mixed-method approach to obtain both qualitative and quantitative data insight.

4) REFLECTION

The Reflection included a follow-up questionnaire and semi-structured interview. The purpose of the questions asked was to obtain feedback about the positive, negative and neutral aspects of their experience with our tool, as well as identifying which future features would be of most use to them. The reflection questionnaire included questions on: how well the tool is able to accommodate for a variety of *usability features* (incl. abstraction of data, detailed data, ease of use, emphasis on pertinent information, exploratory abilities, situational awareness, real-time performance and ability to predict), *rating concerns* (incl. learning curve, situational awareness, real-time performance, scale of volume, other), rating future possible features, opinion on limitations, expected frequency of use if the tool was available as a mature production-line tool, their overall interest in the tool, and any concerns related to the study.

The interview questions focused on having the participants articulate their opinions about the tool, their over experience, what features could improve tool, which elements they think should remain the same, and have them critique the metaphor mappings. The interview questions were the following:

1) *Describe in your own words what you think of the tool.*
2) *What was your overall experience with the tool like?*
3) *What can the tool improve?*
4) *What should the tool keep the same?*
5) *What other features do you think the tool should include?*
6) *Describe what you would consider normal activity?*
7) *Is there anything in this experience or experiment setup you found particularly problematic and would like to highlight? (If yes, then ask: "what"?)*

B. SCENARIOS IN THE STUDY

Attacks were designed to be simple in order to: 1) allow for controlled laboratory condition testing of attacks and 2) allow for testing the feasibility of pattern-of-life visual metaphors as an detection tool, not how it compares against other visualization methods. We deemed it necessary to simplifying the attacks in order to assess the viability of the general approach.

For each of the three attacks, we recorded a session of malicious behaviour using a combination of automated and manual activities invoked by an actual human (prior to the study), see Figure 3. Each of the three attack datasets were recorded once and played back (using the aforementioned playback feature) to ensure each participant would receive the same stimuli over the course of the study session. Each scenario was pseudo-randomly selected to prevent order effect biases as each participant had to view all three scenarios exactly one time each. The attack scenarios (at a high level) involve:

1) **sabotaging of local file stores** through creation and deletion of sensitive files and folders
2) **botnet scanning activities**
3) **resource flooding** (CPU and network).

Each scenario lasted five minutes in length and consists of three instances of attacks being executed, began at different time intervals, and lasting between five to 20 seconds at a time. It is important to note that we are interested in the

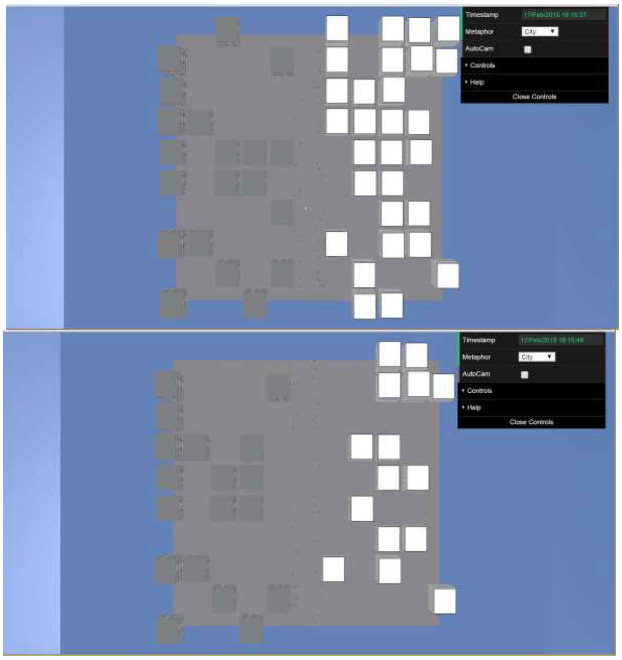

Timestamp 17/Feb/2015 16:15:37
Metaphor City
AutoCam □
Timestamp 17/Feb/2015 16:15:49
Metaphor City
AutoCam □

FIGURE 4. Scenario 1: (top) before the sabotage attack (top-down perspective). (bottom) An insider threat has run a piece of malware on a restricted system to damage its files and folders. This attacker aims to sabotage integrity of a sensitive system. Note the disappearance of specifically monitored files and folders (buildings).

number of true positives and negatives rather than the number of false positives.

As our implementation uses metaphors that have not been used in detection before, our focus is on whether participants can understand these metaphors and point out abnormal activity in the pattern-of-life visual metaphors that can be indicative of malicious host activity. We believe any false positive alerts would depend on how well optimised mapping strategies are or, more importantly, on whether participants have understood the metaphor. We asked participants during the interview session why they believed anomalous activity took place and we try to establish which aspects of the metaphors confused them. We believe it is important to determine if participants understand the mapping strategies from host activity to metaphors, before trying to optimise by obtaining statistics about the false positive ratio.

1) SCENARIO 1 – SABOTAGE

The attack script is built on the idea that an insider has run a piece of malware on a restricted system to sabotage integrity of a sensitive files and folder system. During the attack, the local host continually writes to disk (arbitrarily re-writes to files and folders at script-specified intervals). Structures of machine home directory change more often than they should, and the user can no longer trust the integrity of the file structure. Figures 4 and 5 show the attack affecting the pattern of life:

- **City**: buildings appearing and disappearing throughout, while the patterns and behaviour of people, weather and cars stay (largely) the same throughout.

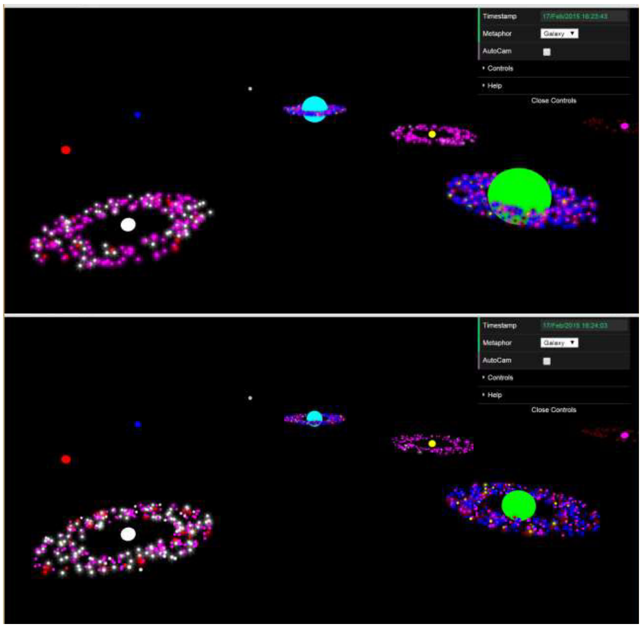

Timestamp 17/Feb/2015 16:23:43
Metaphor Galaxy
AutoCam □
Timestamp 17/Feb/2015 16:24:03
Metaphor Galaxy
AutoCam □

FIGURE 5. Scenario 1: (top) before the sabotage attack. (bottom) An insider threat has run a piece of malware on a restricted system to damage its files and folders. This user aims to sabotage integrity of a sensitive system. Note the increase in yellow stars.

- **Galaxy**: a large increase and unusual behaviour in yellow stars. These are generally rare as they relate to activities that have to do with file and folder creation and deletion.

2) SCENARIO 2 – BOTNET SCANNING

The attack scripted mimics malware that scans the LAN and reports results out to an IP address on the Internet. This scenario aims to show a confidentiality attack. During the attack the Local Network is periodically scanned by malware with legitimate user credentials, and increases network activity in bursts. Minor CPU and RAM increases, see Figure 6:

- **City**: Bursts of cars appearing in intervals, each being a new connection made. Minor CPU disruptions by the bursts of network scans, which also affect the weather. The number and behaviour of people and buildings remain the same throughout.
- **Galaxy**: A large number of Network card Events (white stars) appearing on three of the components (CPU, RAM, Network Card) compared to other machine activities.

3) SCENARIO 3 – RESOURCE FLOODING

In this attack script, we assume an insider has run malware on a safety-critical system to deliberately affect its performance with high resource utilisation - akin to flooding attacks. This malware aims to disrupt service of a system that is crucial to remain available. During the attack intervals, a resource utilisation cripple the performance of several of the Components of the machine, with a large volume of participating actors present at seemingly haphazard intervals,

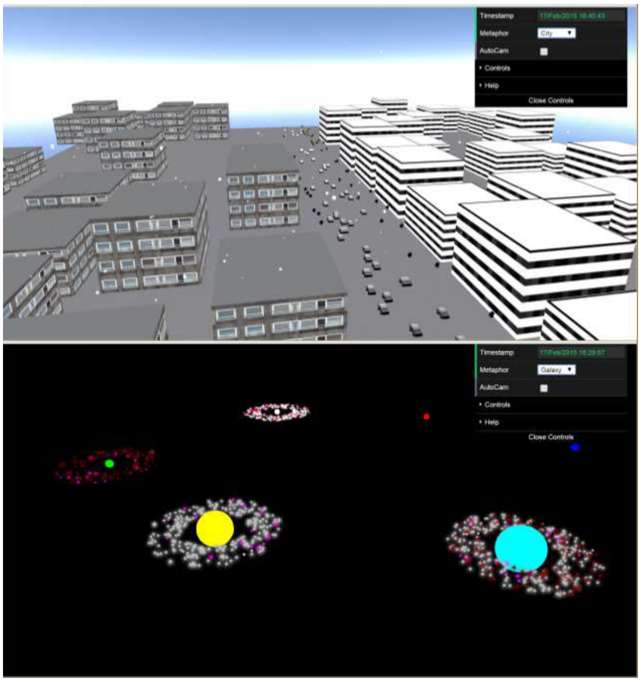

Timestamp 17/Feb/2015 16:40:43
Metaphor City
AutoCam □

Timestamp 17/Feb/2015 16:29:57
Metaphor Galaxy
AutoCam □

FIGURE 6. Scenario 2: Botnet scan activity from a target machine. A local machine has been compromised through a phishing attack, and now belongs to a botnet. A piece of malware scans the LAN using credentials of a legitimate user and reports results out to the Internet. (top, city) Significant increase and unusual behaviour in car traffic patterns. (bottom, galaxy) Significant increase and unusual behaviour in white stars (network related activities).

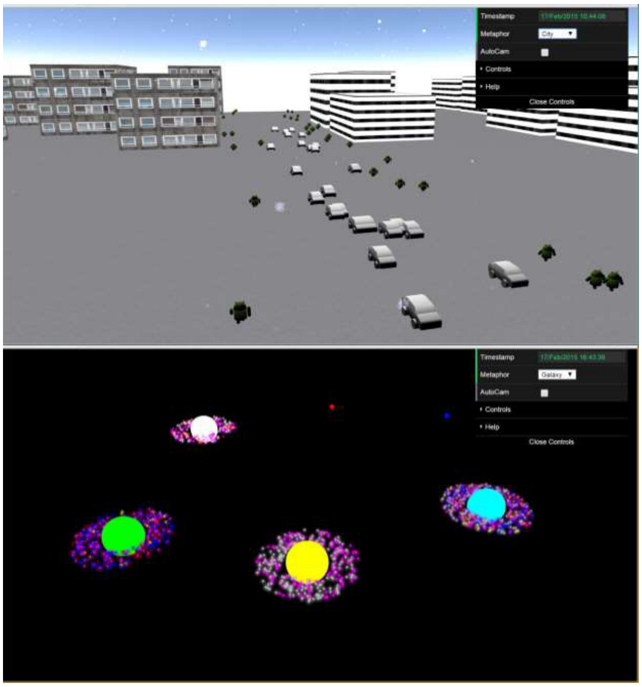

Timestamp 17/Feb/2015 10:44:06
Metaphor City
AutoCam □

Timestamp 17/Feb/2015 10:43:36
Metaphor Galaxy
AutoCam □

FIGURE 7. Scenario 3: An APT running on a safety-critical system that impacts its performance, aiming to disrupt s service that must remain available. (top, city) It snows and rains continuously. (bottom, galaxy) A significant number of CPU and RAM events compared to other machine activities.

including CPU, RAM, number of processes active and network activities. Flooding of resources affect the availability of the system. In Figure 7, the attack affects the pattern of life in the following ways:

- **City**: continuously snowing with irregular patterns of volume of participating actor.
- **Galaxy**: CPU and RAM stars dominate the stars being generated.

VI. OBSERVATIONS AND FEEDBACK

A. PARTICIPATION SUMMARY

Each session took between 1.5 and 2 hours each to complete. All participants were security researchers. In total three participants completed the study; two male and one female, with 2hours, 58min and 21 seconds of voice recording (interviews, audio records from practice period and main scenario) and 66min and 54 seconds of eyetracking data and video (no audio). Table 1 shows the demographics of the participants.

B. STUDY RESULTS

In this section we review our qualitative assessment. Eyetracking recordings were used to cross-check statements by the participants. As pattern-of-life visual metaphors for anomaly detection is a novel concept, our participants had no reference point to compare with pre- and post exposure other than expectations. This is why we employed a qualitative approach to assessment. Participant-specific observations and statements about tool usage included:

TABLE 1. Study demographics.

Participant No.	P1	P2	P3
Age?	30-39	18-29	18-29
Gender?	M	F	M
Adolesence location?	Europe	Europe	North America
Security background?	Yes	Yes	Yes
Security aspect?	Technical	Technical	Technical
Visualization background?	No	Yes	Yes
Which visualization aspects?	N/A	Graphics	N/A
Work sector?	Academia	Academia	Academia
Work sectors in security?	Academia	Academia	Academia
How many years?	5-10 years	5-10 years	5-10 years
Visualization importance?	8/10	6/10	3/10
Play Video Games?	< Once a month	Never/ rarely	< Once a week
Colour Blind?	No	No	No
Corrected Eyesight?	Yes	Yes	Yes

Participant 1 (P1) focused on navigating around in the virtual environment. His navigation patterns were the most volatile (continually moving with the mouse and keyboard) of the participants. P1 had a more difficult time verbalising his thought process during the main attack scenarios and kept mostly quiet.

P2 focused on zooming in and out of each Metaphor in the effort to obtain both the bigger picture and all details within the metaphor. Like P1, she also assumed people entering buildings meant file access. P2 often placed the camera beneath the city. As the renderer performs back-face culling for performance reasons this makes the city floor to become invisible when the camera is beneath (and camera pointing upwards). In the interview P2 stated this was because she had trouble seeing the rain on the grey road in the city. This suggests we need to ensure visibility of the rain is guaranteed from any angle above the city. In the galaxy metaphor she would attempt to zoom out in the galaxy to get an overview of the scene, but kept zooming back into each galaxy centre when large amounts of new events would be generated. During the interview P2 expressed that the galaxy distances should be made configurable as it was difficult to view the whole galaxy scene at once and still be able to get a good view of the star distributions.

P3 zoomed out often in the effort to get a complete top down view of the scene, but unlike P2, P3 stayed zoomed out for a majority of the time (in both metaphors). The eyetracker pattern suggests he was able to view participating actors in the metaphor despite being zoomed out. During the interview P3 compared the city metaphor to the video game Sim City. P3 expressed an interest in being able to hover the mouse above actors in the scene to understand what they represent (at a raw log level). Both P1 and P3 discussed the possibility of higher-level Events (e.g. several Events can be related to a higher-level activity, e.g. opening up a browser can net new CPU, RAM and NETWORK Events – but the user is doing a single activity). P3 also wanted an easy way of changing the Importance of Events during run-time.

All participants preferred the city metaphor over the galaxy metaphor, and stated that the amount of data presented on screen in the galaxies was difficult at times to make sense of compared to the city.

Positive feedback and observations across participants:

- **Ease of detection**. Participants were able to identify all anomalous activities at the correct times straightforwardly.
- **Ease of navigation**: they stated they found navigation straightforward and easy to pick up and said the tool helped them focus on the relationship of the data types, as opposed to the data types themselves.
- **Ease of reasoning with data transformations**. Participants said that the strong aspect of the tool lies in its ability to transform information insight into more engaging visuals than traditional data visualizations.
- **Detection potential**. Participants expressed that pattern-of-life visual metaphors have potential, and that our tool is a good concept demonstrator, but that further development is necessary to make it reach its potential.
- **Using intuition to detect anomalies**. Participants guessed that people represented processes and users, but found it challenging to interpret the meaning of how people behaviour relates to buildings. All stated that they found rain and snow easy to interpret.

Negative feedback and observations across participants:

- **Presence of false positives**. P1 and P2 assumed that in the rare instance when they saw people entering buildings in the city metaphor that this meant "file access". In the case of network activities, all three assumed it was due to be large transfers of data as opposed to network scanning (with many connections being made). This could be corrected by giving users access to more information about what a participating actor represents, or allowing users to change what the participating actor represents. In this case we could show more cars of different colours or change vehicle types to signal other types of network related data.
- **Lack of confidence in detection**. Participants expressed confidence concerns with regards to their own ability to know "*what is anomalous?*" This could be attributed to the fact that pattern-of-life visual metaphors is a novel concept or that the tool is new to them. Further testing would be necessary to determine this.
- **Cross-examining different metaphors is not straightforward**. Participants said separate view of the same data were difficult to discern, as well as identifying how the two views relate to each other. P2 stated that multiple windows to view each metaphor in tandem may help.
- **Identifying relationships between participating actors can be challenging**. Participants found it challenging to interpret the meaning of:
  - how people behaviour relates to buildings.
  - stars in the galaxy metaphors, esp. when many stars are clustered together.
- **Visual perception performance**. P3 stated he had difficulties seeing dark-red stars in a black universe in the galaxy cluster metaphor.

Neutral feedback and observations across participants:

- **Generalisation vs. personalisation**. Two of the participants believed aspects of metaphors can be generalised, but specified that individual needs are more important to overcome. P2 highlighted this should be modifying the galaxy-star creation parameters (e.g. sizes of stars). P3 said that if the user can introduce other visual variables into the metaphor.
- **Movement may dominate foveal vision patterns**. Movements of participating actors appeared to dominate viewing patterns, and forced participants to look at certain areas of the screen. For instance, focus always shifted to incoming cars in the city metaphor each time there was a sudden influx of cars.
- **Duration of animations**. If objects are removed from the scene with the blink of an eye (e.g. files deleted → buildings disappearing), the participants' may not register a change due to inattentional blindness [52].

Figure 8 suggests that the tool is able to address the features listed below: delivering metaphors that resonate well; delivering detailed data to show fine information about activities; ease of use; easy to navigate; prediction capabilities;

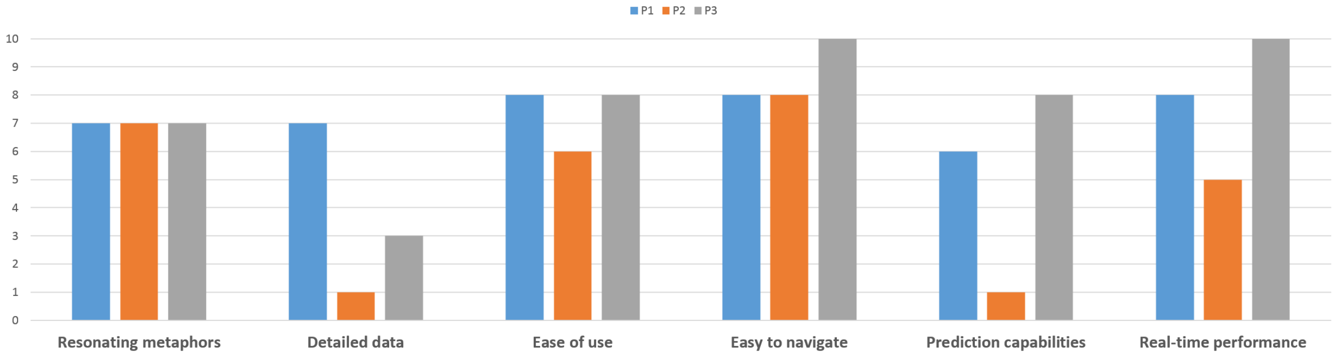

FIGURE 8. Ratings of opinions on feature aspects of the tool (0 = not well, 10 = very well).

real-time performance. The key improvement is access to detailed data on demand in the metaphors.

VII. DISCUSSION

A. REFLECTION

All participants stated they were able to use the city metaphor straightforwardly, but struggled to discern and extrapolate the real-world meaning from the galaxy cluster metaphor. Participants were able to correctly identify when the attacks happened, but found it difficult to discuss why and how they were able to identify the attacks. They stated they believe that pattern-of-life visual metaphors has potential to reliably detect anomalies, but suggested that further research and development will be necessary to reach its potential. We believe this to be the case as well.

Our findings indicate that pattern-of-life visual metaphors need to able to accommodate for individual needs (e.g., personalisation of the metaphor itself or optimisation of usability). All participants stated that metaphor need to be richer (i.e. more participating actors in them), while minimising false positive (i.e., minimise aspects of the metaphor that can be misinterpreted. For instance, participants found it somewhat challenging to interpret the behaviour of people in the metaphor, with P1 and P2 believing that people going into buildings meant "file access"). P2 also suggested elimination of metaphor actors that do not provide a direct mapping. Specifically, P2 suggested that if two datasets are strongly linked (e.g. network CPU process and network packets), it may also very well be that only one of the data sources ought to be required. Whether 1:1 mappings are the only pattern-of-life visual metaphors that can provide meaningful interpretations of complex behaviours, or whether other $1 : N$, or $N : 1$ and $N_a : N_b$ also yield merit remains to be seen. We have identified three areas for room for improvement in studying pattern-of-life visual metaphors for anomaly detection:

- **Novelty in pattern-of-life visual metaphors makes it challenging to assess them**. As anomaly detection using our general approach is a novel concept, *participants do not have a reference point for this approach* (to investigate security issues). We identified two false positives: with P1 and P2 interpreting a person entering a building as file access. It is challenging to give an exact false/true positive rate from real-time visualizations as we do not have access to how participants think, and how often they deemed an attack to have occurred. Continually pausing the tool to enquire the thought processes of participants might be a way to obtain this insight, but frequent pausing will break the flow of tool usage. Pausing also does not reflect the tool's intended use. Our current opinion is that in order to determine effectiveness of pattern-of-life visual metaphors, we measured how many appropriate incident response decisions are made as a consequence of viewing and interacting with our general approach – instead of reviewing every possible visual interpretation.
- **Verisimilitude of Attacks**. The study was created in laboratory conditions, and the attacks were simulated. Real attacks are unlikely to happen in the conditions imposed by the study. Ideally, a longitudinal study ought to be conducted without synthetic and manual created attacks, and instead using real attacks as input data.
- **Duration of Experiment Affecting Performance**. 1.5-2 hours for a study can be long for any participants to have to sit through. Between part 2, 3 and 4 we asked participants whether they wanted a break between session parts. All participants were comfortable to continuing to the end of the study. We do not believe this affected participant performance in any significant way, but we have no evidence to the contrary either.

It is important to point out that we designed the study for laboratory conditions to limit factors from affecting results and assess the idea behind our general approach, including: simulated, attack scenarios (with more control of factors), a small pool of participants (to get more in-depth information from each participant) and focus on a qualitative approach to analysis to identify broader issues and benefits of using metaphors for anomaly detection because of its novelty. A full user study will be necessary to make any generalisable claims about pattern-of-life visual metaphors.

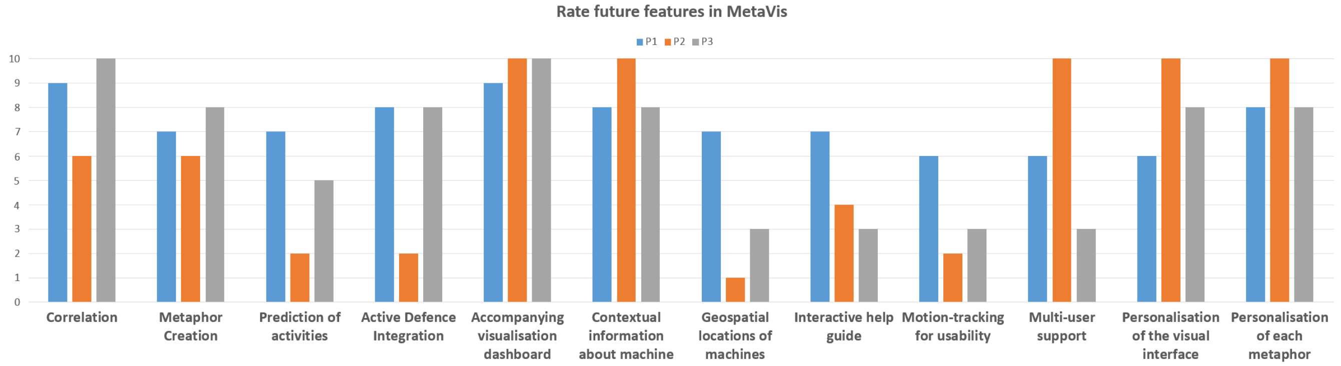

FIGURE 9. Rating preferences of possible future features of the tool (1 = not important, 10 = very important).

B. RECOMMENDATIONS

As mentioned, the purpose of this paper is to explore the feasibility of the concept of pattern-of-life visual metaphors. From our concept demonstrator and study, we have the following recommendations:

- **Pattern-of-life visual metaphors should complement, but not replace, other visualization and detection methods**. We envisage our general approach as a first-pass detection mechanisms, e.g. on large screens in security operation centres, and used collaboratively between analysts.
- **Allow users to access the raw logs or alert events** from within the metaphor if possible, and integrate with other visualization tools. We suspect this will *support internal locus of control* [49] for users, and allow users to use pattern-of-life visual metaphors as an investigative tool.
- **Make use of temporal and dynamic elements to show how the deltas are impacting the virtual scene** (perhaps animations to shows *cause and effect*). Small animations may help emphasise when noteworthy deviations are happening, or when transitioning between states. An example of this may be buildings collapsing to signify file deletion (rather than disappearing immediately).
- **Employ a well-defined approach to mapping variables to visuals**. We took Bertin's *visual variables* [47], Ferwerda's three varieties of realism [48] and Averbukh's discussions on metaphor visualizations [32] as a starting point. We believe Shneiderman's eight usability principles can be a useful guide to ensure usability further [49].
- **Ensure visuals metaphors follow easy-to-understand reasoning structures**, perhaps an underlying formal-semantic reasoning structure.
- **Mapping strategies need special attention, as they**:
  - **Can be strongly linked**. If two Properties as strongly linked (i.e. they correlate often, and there is a causal reason for that correlation), it may be that only one of them or some combination of them should be used.
  - **Must resonate well with users**. If the mappings work well on paper, but users find them challenging to work with, there is little value with the mapping as seen from a usability perspective.
- **Understand that unexpected behaviour in the virtual scene may be interpreted by the user as intentional**. As mentioned. in early iterations of the city metaphor, the virtual people would on occasion walk into buildings, which could be interpreted as "file access" as the buildings represent the file or folder structures on a computer.
- **Playback or rewind features** are likely to engage users to think laterally about the data in question and build hypotheses.
- **Design the workflow pattern** in order to be able to effectively communicate and identify behaviour and interpretations by the user. Iterate on this description and refine its design before, during and after user testing.
- **Understand that metaphors can be both general and highly personal**. While most people understand concepts such as city landscapes or galaxy clusters, users can interpret details in pattern-of-life visual metaphors differently.
- **Understand which type of analogies map well and which do not**. During metaphor scoping, there was significant disagreement between the researchers. We therefore opted for the lowest common denominator approach, and thus only added analogies that everyone could agreed on.

C. FUTURE WORK

To guide our future work, we asked participants to rate our ideas for future work, see Figure 9. The most agreed-on features include: access to detailed data, accompanying traditional data visualization dashboard, personalisation options, and contextual information about the scene itself. Below follows a further reflection on future work:

- **Improve assessment methodologies for pattern-of-life visual metaphors**. We have not investigated effectiveness, including to what degree metaphors can be compared against more traditional anomaly detection systems or data visualization techniques. We deemed it necessary to investigate the underlying feasibility of "anomaly detection using pattern-of-life metaphors" first. Future work will need to investigate assessment methodology-related challenges, including how to obtain and compare: false positive rates, response times, true positive rates, and false negative rates to existing threat detection and data visualization methods.

- **Attack Vectors**. It will be necessary to scale the complexity of attacks as attack vectors so far have remained simple. We will need to determine whether our approach concept can compete with traditional misuse and anomaly detection.
- **Metaphor Resonance**. There are no best practices that are able to determine how well-aligned a visualization is to an arbitrary observer. We recognise that a user's perception and understanding of a metaphor is likely to be based on their prior experiences with metaphors, their attentiveness, reasoning skills, cultural background, personal preferences, among other factors. A metric could be developed to provide some indication of how well metaphors are likely to resonate with a user.
- **Recording States**. We currently do not store the state of the Data Model, however, this could be useful to explore if users wish to quickly compare system states at different points in history or compare two different datasets and identify their similarities.
- **Variability Across Devices**. Hardware and software can differ significantly across devices. We envisage possibility of pattern-of-life visual metaphors for mobile devices or of networks.
- **Persistence and Anomalies**. Shifting baselines is a recurring problem in anomaly detection. This might be addressed with an adaptive approach to modifying the Significance threshold. In our tool, we also envisage that uses of deltas (differences only) and sigmas (aggregation of changes only) can also be used to make up metaphor scenes, differently to how we composite our scenes today.
- **Longitudinal Studies**. It would be useful to investigate how pattern-of-life visual metaphors perform over longer periods of time, under different attack conditions and in production environments.
- **Multiple-Tool Instantiations**. Our concept demonstrator showed the tool running on a single machine with a single Data Collector and a single Visualizer. Other configurations can exist, including multiple:
  - Visualizers connecting to the same Data Collector,
  - Data Collectors connecting to a single Visualizer,
  - Data Collectors connecting to multiple Visualizers,

  all of which may have different levels of access to the data collected.

VIII. CONCLUSION

In this paper, we have demonstrated the use of pattern-of-life visual metaphors for anomaly detection in a cybersecurity context. We outlined theoretical, design and implementation considerations. The tool shows a real-time, navigable virtual environment based on Events from a Target Machine and is capable of showing activities on a target Machine. We conducted a feasibility study to assess the tool and obtained initial feedback and detection data. Our findings, while indicative only, suggest that pattern-of-life visual metaphors is able to help end-users detect anomalies, but like other anomaly detection methods, metaphors are also subject to false positives (misinterpretation of visuals).

We envisage pattern-of-life visual metaphors being useful for teaching, accessibility and training purposes. In order to better understand the potential of these metaphors, we aim to investigate other metaphors and expand the two we already have with more scenarios, more analogies and mappings, investigate how to link metaphor data back to statistical data for further analysis, and run a larger experiment to assess usefulness with technical and lay users.

It is important to point out that our study assessed the feasibility of our proposed approach. This is why we designed the study with very specific laboratory conditions in mind including: simulated, simple attack scenarios (for more control of factors); a small pool of participants (to get more in-depth feedback from each participant) with a qualitative approach to analysis to identify broader issues and benefits of using pattern-of-life visual metaphors for anomaly detection.

Future research would need to investigate how to minimise false positives by optimising mapping strategies. The most difficult research challenge of pattern-of-life visual metaphors, we believe, is to assure that any designed metaphors will work well for an arbitrary user because no best practices currently exist. A full user study with the tool will be necessary to make any generalisable claims.

ACKNOWLEDGMENT

The authors obtained ethical approval from the Oxford Central University Research Ethics Committee (CUREC) to conduct our study.

REFERENCES

[1] O. Depren, M. Topallar, E. Anarim, and M. K. Ciliz, "An intelligent intrusion detection system (IDS) for anomaly and misuse detection in computer networks," *Expert Syst. Appl.*, vol. 29, no. 4, pp. 713–722, 2005.

[2] T. R. Glass-Vanderlan, M. D. Iannacone, M. S. Vincent, and R. A. Bridges, "A survey of intrusion detection systems leveraging host data," 2018, *arXiv:1805.06070*. [Online]. Available: https://arxiv.org/abs/1805.06070

[3] D. E. Denning and P. G. Neumann, *Requirements and Model for IDES—A Real-Time Intrusion Detection Expert System*, document A005, SRI Int., 1985, vol. 333.

[4] N. Ye, S. M. Emran, Q. Chen, and S. Vilbert, "Multivariate statistical analysis of audit trails for host-based intrusion detection," *IEEE Trans. Comput.*, vol. 51, no. 7, pp. 810–820, Jul. 2002.

[5] P. García-Teodoro, J. Díaz-Verdejo, G. Maciá-Fernández, and E. Vázquez, "Anomaly-based network intrusion detection: Techniques, systems and challenges," *Comput. Secur.*, vol. 28, nos. 1–2, pp. 18–28, 2009.

[6] V. Chandola, A. Banerjee, and V. Kumar, "Anomaly detection: A survey," *ACM Comput. Surv.*, vol. 41, no. 3, p. 15, 2009.

[7] A. Lazarevic, L. Ertöz, V. Kumar, A. Ozgur, and J. Srivastava, "A comparative study of anomaly detection schemes in network intrusion detection," in *Proc. SIAM Int. Conf. Data Mining*, 2003, pp. 25–36.

[8] C.-F. Tsai, Y.-F. Hsu, C.-Y. Lin, and W.-Y. Lin, "Intrusion detection by machine learning: A review," *Expert Syst. Appl.*, vol. 36, no. 10, pp. 11994–12000, 2009.

[9] G. Conti, *Security Data Visualization: Graphical Techniques for Network Analysis*. San Francisco, CA, USA: No Starch Press, 2007.

[10] R. Marty, *Applied Security Visualization*. Boston, MA, USA: Addison-Wesley, 2009.

[11] R. Marty. (2007). *SecViz*. [Online]. Available: http://secviz.org/

[12] H. Shiravi, A. Shiravi, and A. A. Ghorbani, "A survey of visualization systems for network security," *IEEE Trans. Vis. Comput. Graph.*, vol. 18, no. 8, pp. 1313–1329, Aug. 2012.

[13] L. Harrison and A. Lu, "The future of security visualization: Lessons from network visualization," *IEEE Netw.*, vol. 26, no. 6, pp. 6–11, Nov./Dec. 2012.

[14] D. Staheli, T. Yu, R. J. Crouser, S. Damodaran, K. Nam, D. O'Gwynn, S. McKenna, and L. Harrison, "Visualization evaluation for cyber security: Trends and future directions," in *Proc. 11th Workshop Vis. Cyber Secur.*, 2014, pp. 49–56.

[15] K. Gancarz and K. Prole, "Visual techniques for analyzing wireless communication patterns," in *Proc. IEEE Conf. Technol. Homeland Secur.*, Nov. 2012, pp. 341–347.

[16] I. Davidson, "Anomaly detection, explanation and visualization," SGI, Tokyo, Japan, Tech. Rep., 2007.

[17] H. Zhang, M. Sun, D. D. Yao, and C. North, "Visualizing traffic causality for analyzing network anomalies," in *Proc. ACM Int. Workshop Int. Workshop Secur. Privacy Anal.*, 2015, pp. 37–42.

[18] C. P. Lee and J. A. Copeland, "Flowtag: A collaborative attack-analysis, reporting, and sharing tool for security researchers," in *Proc. 3rd Int. Workshop Vis. Comput. Secur.*, 2006, pp. 103–108.

[19] K. Lakkaraju, W. Yurcik, and A. J. Lee, "NVisionIP: Netflow visualizations of system state for security situational awareness," in *Proc. ACM Workshop Vis. Data Mining Comput. Secur.*, 2004, pp. 65–72.

[20] K. Abdullah, C. P. Lee, G. J. Conti, J. A. Copeland, and J. T. Stasko, "IDS rainstorm: Visualizing IDS alarms," in *Proc. VizSec*, 2005, pp. 1–6.

[21] S. Lau, "The spinning cube of potential doom," *Commun. ACM*, vol. 47, no. 6, pp. 25–26, Jun. 2004.

[22] C. P. Lee, J. Trost, N. Gibbs, R. Beyah, and J. A. Copeland, "Visual firewall: Real-time network security monitor," in *Proc. IEEE Workshop Vis. Comput. Secur.*, Oct. 2005, pp. 129–136.

[23] J. Nicholls, D. Peters, A. Slawinski, T. Spoor, S. Vicol, J. Happa, M. Goldsmith, and S. Creese, "NetVis: A visualization tool enabling multiple perspectives of network traffic data," in *Proc. Comput. Graph. Vis. Comput. (CGVC)*, London, U.K., 2013. [Online]. Available: https://diglib.eg.org/handle/10.2312/LocalChapterEvents.TPCG.TPCG13.009-016

[24] R. Borgo, J. Kehrer, D. H. Chung, E. Maguire, R. S. Laramee, H. Hauser, M. Ward, and M. Chen, "Glyph-based visualization: Foundations, design guidelines, techniques and applications," in *Proc. Eurographics (STARs)*, 2013, pp. 39–63.

[25] C. R. Dos Santos, P. Gros, P. Abel, D. Loisel, N. Trichaud, and J. Paris, "Metaphor-aware 3D navigation," in *Proc. IEEE Symp. Inf. Vis.*, Oct. 2000, pp. 155–165.

[26] J. S. Risch, "On the role of metaphor in information visualization," 2008, *arXiv:0809.0884*. [Online]. Available: https://arxiv.org/abs/0809.0884

[27] C. Ziemkiewicz and R. Kosara, "The shaping of information by visual metaphors," *IEEE Trans. Vis. Comput. Graph.*, vol. 14, no. 6, pp. 1269–1276, Nov./Dec. 2008.

[28] Y. Livnat, J. Agutter, S. Moon, and S. Foresti, "Visual correlation for situational awareness," in *Proc. IEEE Symp. Inf. Vis.*, Oct. 2005, pp. 95–102.

[29] SecureDecisions. *Securescope*. Accessed: Feb. 30, 2012. [Online]. Available: http://www.securescope.com

[30] Tenable. *Tenable3D*. Accessed: Aug. 30, 2019. [Online]. Available: http://www.tenable.com/blog/3d-tool-version-20-released

[31] S. Creese, M. Goldsmith, N. Moffat, J. Happa, and I. Agrafiotis, "CyberVis: Visualizing the potential impact of cyber attacks on the wider enterprise," in *Proc. IEEE Int. Conf. Technol. Homeland Secur.*, Nov. 2013, pp. 73–79.

[32] V. L. Averbukh, "Visualization metaphors," *Program. Comput. Softw.*, vol. 27, no. 5, pp. 227–237, 2001.

[33] V. Averbukh, M. Bakhterev, A. Baydalin, D. Ismagilov, and P. Trushenkova, "Interface and visualization metaphors," in *Human-Computer Interaction. Interaction Platforms and Techniques*. Beijing, China: Springer, 2007, pp. 13–22.

[34] V. L. Averbukh, M. O. Bakhterev, A. Y. Baydalin, D. Y. Gorbashevskiy, D. R. Ismagilov, A. Y. Kazantsev, P. V. Nebogatikova, A. V. Popova, and P. A. Vasev, "Searching and analysis of interface and visualization metaphors," in *New Developments*. Berlin, Germany: Springer, 2008, p. 49.

[35] C. R. Dos Santos, P. Gros, P. Abel, D. Loisel, N. Trichaud, and J.-P. Paris, "Mapping information onto 3D virtual worlds," in *Proc. IEEE Conf. Inf. Vis.*, Jul. 2000, pp. 379–386.

[36] C. R. dos Santos and P. Gros, "Multiple views in 3D metaphoric information visualization," in *Proc. 6th Int. Conf. Inf. Vis.*, Jul. 2002, pp. 468–473.

[37] P. Abel, P. Gros, C. R. Dos Santos, D. Loisel, and J.-P. Paris, "Automatic construction of dynamic 3D metaphoric worlds: An application to network management," *Proc. SPIE*, vol. 3960, Feb. 2000, Art. no. 378908.

[38] P. Abel, P. Gros, D. Loisel, C. R. Dos Santos, and J.-P. Paris, "CyberNet: A framework for managing networks using 3D metaphoric worlds," in *Annales des Télécommunications*, vol. 55. Springer, 2000, pp. 131–142.

[39] J. A. Brown, A. McGregor, and H. Braun, "Network performance visualization: Insight through animation," in *Proc. PAM Passive Active Meas. Workshop*, Apr. 2000, pp. 33–41.

[40] L. Axon, J. Happa, A. Janse Van Rensburg, M. Goldsmith, and S. Creese, "Sonification to support the monitoring tasks of security operations centres," *IEEE Trans. Dependable Secure Comput.*, to be published.

[41] J. S. Yi, R. Melton, J. Stasko, and J. A. Jacko, "Dust & magnet: Multivariate information visualization using a magnet metaphor," *Inf. Vis.*, vol. 4, no. 4, pp. 239–256, 2005.

[42] T. Panas, R. Berrigan, and J. Grundy, "A 3D metaphor for software production visualization," in *Proc. 7th Int. Conf. Inf. Vis.*, Jul. 2003, pp. 314–319.

[43] N. Gershon and W. Page, "What storytelling can do for information visualization," *Commun. ACM*, vol. 44, no. 8, pp. 31–37, 2001.

[44] O.-M. Latvala, T. Keränen, S. Noponen, N. Lehto, M. Sailio, M. Valta, and P. Olli, "Visualizing network events in a muggle friendly way," in *Proc. Int. Conf. Cyber Situational Awareness, Data Anal. Assessment*, Jun. 2017, pp. 1–4.

[45] F. Carroll, A. Chakof, and P. Legg, "What makes for effective visualisation in cyber situational awareness for non-expert users?" in *Proc. Int. Conf. Cyber Situational Awareness, Data Anal. Assessment*, 2019.

[46] P. Legg, N. Moffat, J. R. Nurse, J. Happa, I. Agrafiotis, M. Goldsmith, and S. Creese, "Towards a conceptual model and reasoning structure for insider threat detection," *J. Wireless Mobile Netw. Ubiquitous Comput. Dependable Appl.*, vol. 4, no. 4, pp. 20–37, 2013.

[47] J. Bertin, *Semiology of Graphics: Diagrams, Networks, Maps*. Madison, WI, USA: Univ. Wisconsin Press, 1983.

[48] J. A. Ferwerda, "Three varieties of realism in computer graphics," *Proc. SPIE*, vol. 5007, Jun. 2003, Art. no. 473899.

[49] B. Shneiderman and C. Plaisant, *Designing the User Interface: Strategies for Effective Human-Computer Interaction*, 4th ed. Boston, MA, USA: Addison-Wesley, 2004.

[50] G. E. Krasner and S. T. Pope, "A description of the model-view-controller user interface paradigm in the smalltalk-80 system," *J. Object Oriented Program.*, vol. 1, no. 3, pp. 26–49, 1988.

[51] Seeing Machines. (2013). *Facelab*. Eyetracking System. [Online]. Available: https://www.eyecomtec.com/3132-faceLAB

[52] A. Mack and I. Rock, *Inattentional Blindness*. Cambridge, MA, USA: MIT Press, 1998.

JASSIM HAPPA is a Lecturer in Information Security at Royal Holloway, and a Visiting Lecturer in the Department of Computer Science at the University of Oxford. He obtained his BSc (Hons) in Computing Science at the University of East Anglia in 2006. After a year of working as an Intrusion Detection System (IDS) analyst, he began his PhD in Engineering at the University of Warwick in October 2007 in Computer Graphics. He defended his PhD in January 2012, and worked as a Research Fellow from December 2011- August 2019 at Oxford. In September 2019, he joined Royal Holloway. In recent years he has spent his research efforts on cybersecurity topics such as: analytics, visualization, threat modelling, situational awareness, risk propagation, resilience, decision support, privacy and cyber threat intelligence. Teaching wise, he tutors and lectures a variety of computer graphics and security related-subjects at both Royal Holloway and Oxford.

THOMAS BASHFORD-ROGERS is a Senior Lecturer in Games Technology in the Department of Computer Science and Creative Technologies at the University of the West of England. He has a degree in Computer Science from the University of Bristol in 2003 and a doctorate in Computer Graphics from the University of Warwick in 2011. He worked as a Research Fellow in graphics at the University of Warwick from 2011 to 2016, and then in Cyber Security at the University of Oxford in 2017. His research interests include global illumination, raytracing, Monte Carlo methods, machine learning, deep learning, high dynamic range imaging and cyber security.

IOANNIS AGRAFIOTIS is an Officer at the European Union Agency for Cybersecurity (ENISA) working in the areas of Data Protection and Certification. Ioannis has an affiliation with the University of Oxford, where he acts as a Senior Researcher in cybersecurity with the Department of Computer Science and as a James Martin Fellow at the Global Cyber Security Capacity Centre (GCSCC). In his academic role, he has participated in various projects conducting research in areas such as capacity building in cybersecurity, risk analysis and resilience in the cyber domain, cyber insurance, and anomaly detection techniques for internal and external threats. Ioannis completed his doctoral studies in Engineering at the University of Warwick (2012, EPSRC-funded). He also holds an MSc in Analysis, Design and Management of Information Systems from the London School of Economics and Political Science (2008) and a BSc in Applied Informatics from the University of Macedonia in Greece (2006).

MICHAEL GOLDSMITH is a Director of the Global Cyber Security Capacity Centre at the Oxford Martin School and a Senior Research Fellow in the Department of Computer Science and at Worcester College, University of Oxford. His research spans a wide range of topics within security, from the mathematical to the social. He received his DPhil in Computation from Oxford University three decades ago for work on support for specification logics, and has also worked in concurrency theory and formal verification through exhaustive state-exploration.

SADIE CREESE is a Professor of Cyber Security in the Department of Computer Science at the University of Oxford. She teaches threat detection, risk assessment and operational aspects of cyber security. Her current research portfolio includes developing mathematical models for calculating Cyber-Value-at-Risk for an organisation, threat modelling and intrusion detection including insiders, visual analytics for understanding and communicating cyber security postures, logics for predicting risk propagation, resilience strategies, privacy vullnerability, threats to distributed ledgers, and understanding the nature of cyber-harm. She also researches what constitutes cyber security capacity for a nation, and how to deliver it, working with international organisations, nations, practitioners and academics around the world. She is Principal Investigator on the AXIS Insurance Company-sponsored project "Analysing Cyber Value-at-Risk" focused on developing a method for predicting the range of potential losses arising from cyber-attacks taking account of risk control practices. She was the founding Director of the Global Cyber Security Capacity Centre (GCSCC) at the Oxford Martin School and continues to serve as a Director. She was the founding Director of Oxford's Cybersecurity network launched in 2008 and now called CyberSecurityOxford, and is a member of the Advisory Board for the World Economic Forum's Centre for Cybersecurity. Sadie is a Fellow of Worcester College, Oxford.