Determining the Effectiveness of Representing Intrusion Detection System Log Files with Visualization Techniques

Gina Thomas and Kristen Liggett
Collaborative Interfaces and Teaming Branch, Interfaces & Applications Section Air
Force Research Laboratory, 2255 H St, Wright Patterson AFB, OH 45433
USA

gina.thomas.2@us.af.mil; kristen.liggett@us.af.mil

ABSTRACT

After several knowledge elicitation activities with two intrusion detection system (IDS) analysts, two different visualization techniques were selected that best portrayed the most important information relevant to a cyber security task. These two visualizations (customized parallel coordinates and scatter plots) were then investigated for their effectiveness in portraying potential attack patterns to novices. Both visualizations were also subjected to a distortion factor and a complexity factor to determine the robustness of the visualization techniques. A total of 44 volunteers participated in the study. Results revealed a number of significant interactions but overall, performance data showed that scatter plots consistently provided faster response times while parallel coordinates showed more consistent accuracy performance. Lessons learned and future research ideas are discussed.

1.0 INTRODUCTION

According to a report published by a security firm last year, the estimated cost of cybercrime is $113 billion a year and has victimized approximately 378 million people (White Hats, 2014). Not only are cyber crimes costly, they are becoming more common. According to the Ponemon Institute's annual report on the cost of cyber crime (2013), companies experience approximately two successful cyber attacks per week (infiltration of the company's core network) and the most costly are denial of service attacks, malicious insiders, and web-based attacks.

On the military side, the same trend of increasing attacks is occurring (Leopold, 2013). Superiority in the cyberspace domain is essential to maintaining Air Force mission assurance and enabling new mission capabilities. As such, this domain must be mastered in order to ensure success. The budget more than doubled in 2014 from 2013 and is estimated at $5.1 billion in 2015, a portion to be invested in science and technology programs to support cyber defense (McCaney, 2014). A common response has been to develop new software to automate processes, but no one has been successful at creating computer systems that are immune to the attacks by adversaries who are capable of adapting their weapons and tactics as new defense technologies come on-line (Jackson, 2012). The cyber domain has the unique characteristic in that, the resources necessary for superiority in the domain are not aircraft and airfields as in the air domain, but highly trained and competent personnel (Bryant, 2013). The Department of Defense expects to add 4000 cyber operators over the next four years with the majority slated to defend the networks that support worldwide operations (Costlow, 2013). The success or failure of cyber superiority rests on the human factor - in the hands of the cyber operator. Therefore, it is imperative that the systems with which these operators interact are designed for effective human use. Understanding adversaries and their intent is still primarily a human skill that cannot be automated with even the most sophisticated computer software (Aucsmith, 2011). “Cyber situational awareness” has become a popular concept (Barford et. al., 2010); however, there is much debate over whom or what should maintain this awareness – a computer or a human operator. Instead of taking the human further and further out the loop (from in-the-loop to on-the-loop to observing-the-loop), perhaps keeping the operators in-the-loop and designing their systems to better support them is the most efficient way to enhance cyber operations in the future.

Putting the human in charge of defending computer networks requires a change in the way in which cyber traffic data are collected, processed, and presented. Currently, cyber analysts sit diligently monitoring vast amounts of cyber traffic trying to determine if said traffic is suspicious or anomalous (D’Amico, Tesone, Whitley, O’Brien, and Roth, 2008). Specifically, we found during our interactions with experts that analysts are responsible for triaging about 3000 intrusion detection system (IDS) events per hour (represented as plain text rows in a table.) Optimizing the presentation of information that analysts need to perform the job based on human cognitive and perceptual skills and abilities, as well as limitations, should make analysts more effective in this task (Goodall, 2005).

Whatever the attack, the keys to combating it are recognizing it and responding to it appropriately. Recognizing an attack involves understanding the signatures and sources of the malicious traffic (Jackson, 2013). Signatures can mean visual patterns, and visual pattern matching is a skill at which humans excel. Using visualization techniques to portray cyber security data has been studied for a number of years (Conti, 2007; Marty, 2009). Common visualizations for cyber network data include scatter plots, node-link diagrams, histograms, tree maps, and parallel coordinates (see Table 3-1 in Marty, 2009). Each of these visualization techniques may be capable of showing patterns or signatures of attacks, some better than others. A typical IDS log file may contain upwards of 10,000 rows of data for a day’s worth of network traffic. This is especially true if the rule set for the IDS is configured to favor false positives over false negatives, which is often the case due to the potential high cost of missing an attack. If the information in thousands of rows of defensive log files can be presented to analysts in a single visualization containing patterns that might imply an attack, this visualization should help analysts detect and respond to cyber attacks more effectively. Visualizations have been proposed to help with this task; however, the visualizations typically encode all of the information contained in the log file without focusing on only the important elements from the analysts’ perspective to detect anomalies (Harrison and Lu, 2012; Goodall, 2009; Thompson, Rantanen, Yurcik, and Bailey, 2007; Livnat, Agutter, Moon, Erbacher, and Foresti, 2005; Koike and Ohno, 2004).

In this paper, we will discuss how we performed knowledge acquisition activities with cyber defenders to determine the most important dimensions of IDS log files and applied this knowledge to the design of two visualizations for testing. The study reported here focuses on presenting cyber defensive data as from IDS log files graphically in scatter plots and parallel coordinates plots to determine which of these visualizations best aids an analyst in detecting attacks. Scatter plots are most effective at encoding two dimensions of data, as these dimensions can be given a direct spatial encoding. Additional dimensions can be encoded using color, shape, size, texture, etc., but these are less effective than spatial encodings and should be reserved for less important dimensions. Scatter plots are typically used to examine data relationships or to detect clusters and trends in the data. Parallel coordinates (Inselberg & Dimsdale, 1991) depict n-1 pair-wise relationships out of a possible n^2 number of pair wise relationships, where n is the number of dimensions. In contrast to scatter plots where data points are depicted as points, in parallel coordinates data points are depicted as line segments that connect the parallel pairs of axes. These two visualization techniques were compared using the same data set. Determining the elements to be portrayed with these two visualization techniques and then designing them to highlight these element may help analysts more quickly and accurately determine the security of their network. The objective of this research was twofold: 1) to determine if IDS log data could be transformed into visualizations in which network attacks appeared as patterns (visual signatures), and 2) test user human effectiveness (time and accuracy) of the developed visualization candidates for portraying attack information.

1.1 The Human Factors Design Process

Our research is grounded in the human factors design process. Operator challenges in the cyber domain are well-suited for the human factors engineering design/cognitive engineering design process because cyber systems provide operators with a large amount of data and designers want to automate more processes (Roth, Patterson, and Mumaw, 2002). This approach has been successfully implemented by many including Mahoney, Roth, Steinke, Pfautz, Wu, and Farry (2010) to design, develop and evaluate a system supporting cyber situation awareness. Also particular to the cyber security visualization domain, Shiravi, Shiravi, and Ghorbani (2012) acknowledge that "a design process centered on the needs, behaviors, and expectations of security analysts can greatly influence and impact the usability and practicality of such systems." (p. 1313). A diagram of this process is shown in Figure 1 (Kilgore, Godwin, Hogan, Davis, Pfautz, Woods, Branlat, and Kaufman, 2012). It begins with an analysis phase; effective execution of this phase is ***crucial*** for the success of the design process as it provides the foundation for the design activities. Operators and/or subject matter experts (SMEs) provide the design team valuable understanding of domain information, such as the tasks and processes by which they accomplish their work. Members of the design team can then determine the cognitive processes that need to be supported when the operators are performing these tasks and incorporate this knowledge into the design to support them in a way that facilitates their work flow.

Once the information from the analysis phase is converted into useful design guidelines, the design team can begin designing, refining, and investigating candidate solutions. The iterative design process leverages users at various stages of the design process to ensure the design concepts are supporting them when conducting their tasks.

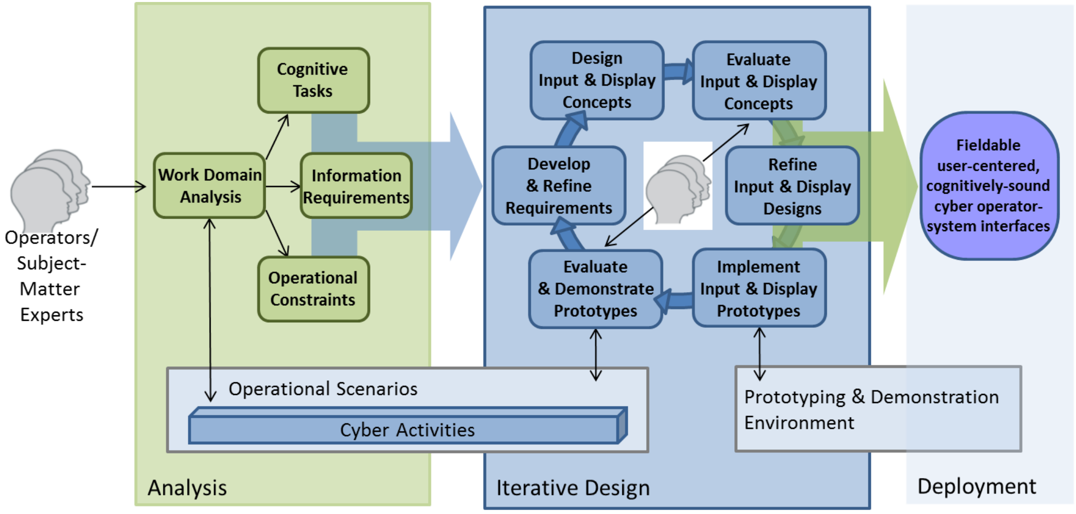
Operators/
Subject-
Matter
Experts
Work Domain Analysis
Cognitive Tasks
Information Requirements
Operational Constraints
Design Input & Display Concepts
Evaluate Input & Display Concepts
Develop & Refine Requirements
Refine Input & Display Designs
Evaluate & Demonstrate Prototypes
Implement Input & Display Prototypes
Fieldable user-centered, cognitively-sound cyber operator-system interfaces
Operational Scenarios
Cyber Activities
Prototyping & Demonstration Environment
Analysis
Iterative Design
Deployment

Figure 1: Human Factors Design Process

1.2 Analysis Phase

Cyber analysts have, at their disposal, a large amount of information about the network they are trying to defend. First and foremost, they have an understanding of the machines on the network, each machine's role, and how everything is connected. This information may be in the form of a physical network diagram/document or the analysts' mental model of the network. They also have access to a variety of log files that document all of the traffic on the network. These log files include packet data log files, firewall log files, and IDS log files. Based on the literature and from discussions with IDS analysts for commercial companies, third-party security vendors, and a commercial financial institution, cyber analysts typically monitor events generated by the IDS in near real-time to maintain network security. If data generated by the IDS is deemed anomalous or suspicious by the analyst, they perform further investigation using firewall and packet data log files.

There are three approaches to cognitive task analyses; one of which is using empirical techniques (informal interviews, retrospective analyses of actual incidents, observation of practitioners as they attempt to perform domain tasks, and formal behavioral studies) to analyze how people actually go about performing the task either in the actual task environment or in a simulated task environment (Roth et al., 2002). For this study, the design team conducted a series of knowledge elicitation sessions with two IDS analysts using observations of practitioners performing the task in a simulated environment. Although CTA activities for various defensive cyber analysts have been documented (Mahoney et al., 2010; D'Amico, Whitley, Tesone, O'Brien, and Roth, 2005; D'Amico and Whitley, 2007), none were specifically focused on IDS analysis and contained the level of cognitive requirements necessary to design effective visualizations for their particular tasks. We essentially followed a think-aloud protocol (Rubin, 1994) in which we had the analysts talk through their activities as they reviewed an IDS log file looking for anomalous activity. The IDS log file from the VAST 2011 Mini-Challenge 2 Day 2 (Grinstein, Cook, Havig, Liggett, Nebesh, Whiting, Whitley, and Koneci, 2011) was used for these sessions. Snort was the IDS used for the challenge, and each log entry contains many details. If opened in a plaintext editor such as WordPad, as shown in Figure 2, each log entry takes up to 6 lines and is followed by a blank line before the next entry. The important parts of a log entry are identified in Table 1. Here, the line number refers to the line number of the entry shown in the WordPad example. We parsed the dataset and transformed it into a table representation such that each row represented an IDS event and each column represented one field (Snort rule, text of rule violated, classification of rule, priority of alert, date, time, source/destination IP addresses, etc.). The file contained 22,425 rows corresponding to 7 unique rules firing during Day 2. We provided these data to the analyst in Microsoft Excel. In general, the analysts did a lot of sorting, scrolling, and filtering. Each analyst spent about 45 minutes talking aloud as they tried to determine what, if anything was concerning in the data. At the end of the session, they indicated places where they would investigate further (look at the packet capture data).

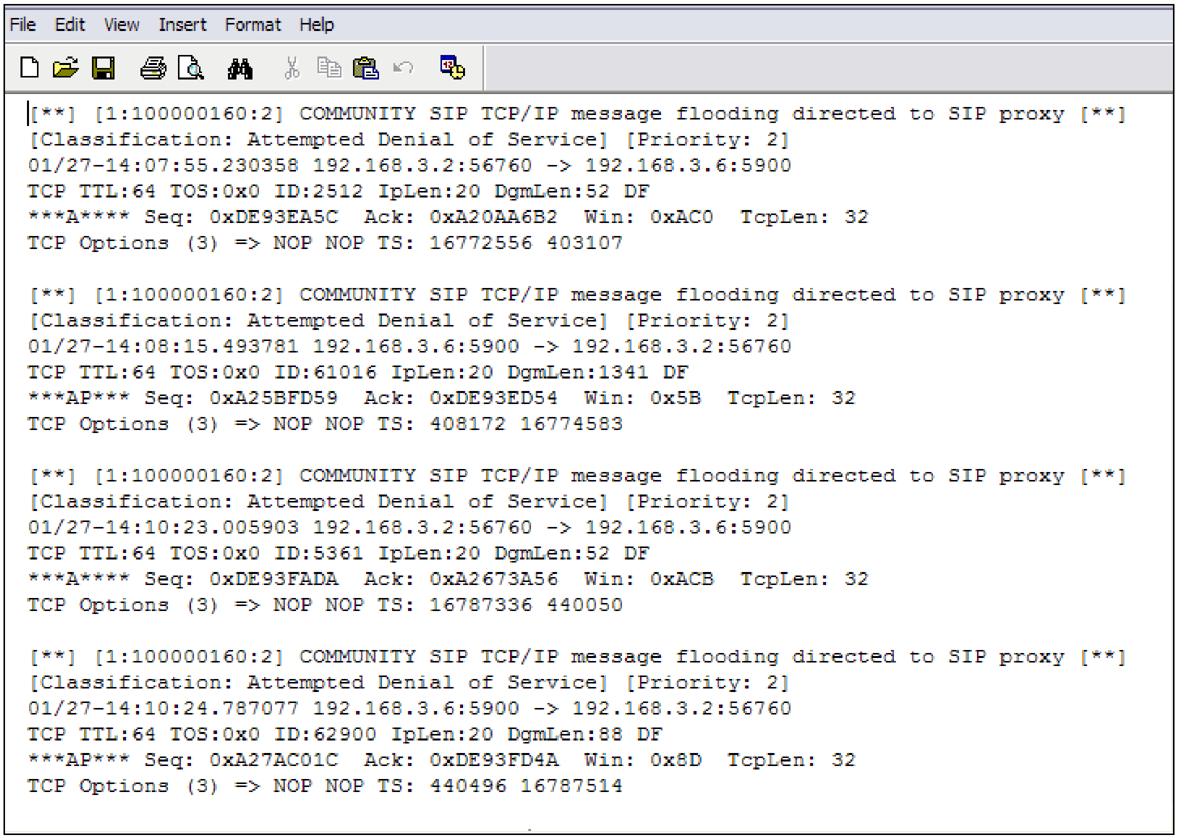

```
[**] [1:100000160:2] COMMUNITY SIP TCP/IP message flooding directed to SIP proxy [**]
[Classification: Attempted Denial of Service] [Priority: 2]
01/27-14:07:55.230358 192.168.3.2:56760 -> 192.168.3.6:5900
TCP TTL:64 TOS:0x0 ID:2512 IpLen:20 DgmLen:52 DF
***A**** Seq: 0xDE93EA5C  Ack: 0xA20AA6B2  Win: 0xAC0  TcpLen: 32
TCP Options (3) => NOP NOP TS: 16772556 403107

[**] [1:100000160:2] COMMUNITY SIP TCP/IP message flooding directed to SIP proxy [**]
[Classification: Attempted Denial of Service] [Priority: 2]
01/27-14:08:15.493781 192.168.3.6:5900 -> 192.168.3.2:56760
TCP TTL:64 TOS:0x0 ID:61016 IpLen:20 DgmLen:1341 DF
***AP*** Seq: 0xA25BFD59  Ack: 0xDE93ED54  Win: 0x5B  TcpLen: 32
TCP Options (3) => NOP NOP TS: 408172 16774583

[**] [1:100000160:2] COMMUNITY SIP TCP/IP message flooding directed to SIP proxy [**]
[Classification: Attempted Denial of Service] [Priority: 2]
01/27-14:10:23.005903 192.168.3.2:56760 -> 192.168.3.6:5900
TCP TTL:64 TOS:0x0 ID:5361 IpLen:20 DgmLen:52 DF
***A**** Seq: 0xDE93FADA  Ack: 0xA2673A56  Win: 0xACB  TcpLen: 32
TCP Options (3) => NOP NOP TS: 16787336 440050

[**] [1:100000160:2] COMMUNITY SIP TCP/IP message flooding directed to SIP proxy [**]
[Classification: Attempted Denial of Service] [Priority: 2]
01/27-14:10:24.787077 192.168.3.6:5900 -> 192.168.3.2:56760
TCP TTL:64 TOS:0x0 ID:62900 IpLen:20 DgmLen:88 DF
***AP*** Seq: 0xA27AC01C  Ack: 0xDE93FD4A  Win: 0x8D  TcpLen: 32
TCP Options (3) => NOP NOP TS: 440496 16787514
```

Figure 2: IDS excerpt from Snort opened in WordPad

.

Table 1: IDS File Content

Line #	**Content / Format**	**Useful Content for this Challenge**
1	[**] [*Snort Rule File*] *Text of specific rule violated* [**]	Text of specific rule violated
2	[Classification: an optional generalized description of the alert] [Priority of the alert]	Classification (if present) Priority
3	Date/Time MM/DD-HH24:MI:SS.milliseconds Source IP/Port Destination IP/Port	Date/Time Source IP/Port Destination IP/Port
Remaining lines	*No content used in the challenge*	None

A discussion followed these sessions to ensure that we understood the analysts' procedures and tactics taken during the session. Although there were a total of 7 columns of information available in the Excel file, we discovered that the analysts searching focused primarily on 4 variables – time, source internet protocol (IP) address, destination IP address, and the text of the specific rule violated. They performed sorts on these variables and talked about looking for patterns in the rules being fired or patterns of timing or rules firing, or patterns of source IPs communicating with destination IPs. With that information, we began the visualization design process.

1.3 Visualization Design

The idea that a large amount of IDS alerts could potentially be summarized in a simple visualization (Shiravi et al., 2012) was attractive. Langton and Newey (2010) point out that current cyber security visualization typically focus on static visualizations of specific applications, which is exactly what we propose below, however, they also point out that operators focusing on network monitoring tasks need an overview of activity with the capability to drill down to investigate suspicious activity. Therefore, once our static overview concept is tested and proven, the static representation can be adapted to an interactive visualization allowing which would allow filtering and a drill down capability to raw data. This methodology supports the visual information-seeking mantra; overview first, zoom and filter, details on demand (Shneiderman, 1996). Since we defined 4 variable of interest, we needed to find visualizations that could support the display of these 4 variables in ways that effectively showed patterns in the data. We started by considering the standard visualization techniques of node-link diagrams, parallel coordinates, scatter plots, and tree maps. We quickly ruled out node-link diagrams because there was no obvious way of representing time, which was a critical variable. We mocked up parallel coordinates with 4 axes (one for each variable) and scatter plots using x position, y position, color of marker, and shape of marker to represent the 4 variables. We were also able to create a variety of nested visualizations embedded in tree maps by partitioning the dataset according to unique rules, sources, or destinations, but because this approach afforded a large number of options (too many to adequately test) we decided to pursue this approach in a follow-up study.

For the parallel coordinates, we went through several iterations of changing axis order, selecting colors, grouping variables into categories versus plotting the actual variables themselves, etc. We considered grouping variables because studies have shown that parallel coordinates suffer issues of occlusion when applied to large data sets (Van Wijk and Van Selow, 1999; Fua, Ward, and Rundensteiner, 1999). The continuous time variable was grouped into 3 categories (morning, work day, and evening); the source IP variable as well as the destination IP variable were organized into 6 categories representing the types sub-network (data center, DC/DNS, office workstations, external web, internet, and unknown); the rule variable was organized into 4 sections based on the rule category (not the individual rule itself). We showed the different versions to our SMEs to get feedback about the variations of visualizations. We settled on a version for which known attacks from the ground truth created discernible patterns. In arranging the rules in this hierarchical way, we discovered that patterns were specific to the rule category. Figure 3 shows a representation of a parallel coordinates visualization and highlights activity indicative of a port scan.

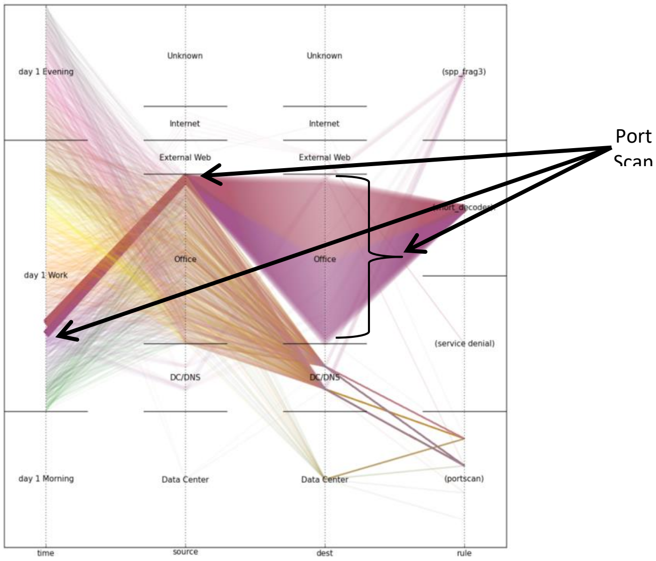

Figure 3: Parallel coordinates representation of IDS log file for Day 1 of VAST Challenge 2011 data. Focusing on the first three axes, a port scan pattern looks like a sideways broom with the handle to the left.

To help evaluate our hypothesis that specific attacks might have unique visual signatures when presented in parallel coordinates, we worked with the SMEs to list the most common attacks, then thought about how time, source, and destination would "look" on a parallel coordinates plot for each attack. For instance, a denial of service attack involves a large number of sources communicating with one destination in a very short period of time. If you plot that type of attack on the first three axes of a parallel coordinates visualization shown in Figure 3, it would look like a diamond (a single point on the time axis, many points on the source axis, and a single point on the destination axis). The fourth axis (rule category) always fell into one category because the various attacks only triggered a single IDS rule category. Generalizing these rules to a "one" or "many" indication for each of the 3 primary axes (time, source, destination) allowed us to develop 8 distinct visualization patterns to use as test stimuli.

Although some of the 8 patterns created are not indicative of a known attack, the patterns allow the user to readily sort and group by rule and to recognize things that are more likely to be an attack versus things that are more likely to be noise. Additionally, since each axis has meaning, defenders can derive meaning from the pattern that allows them to make an assessment about which rule violations are more likely to require follow up. For example, a one-many-one relationship of violations of a particular rule tells the user that at a single time (very short time period), multiple IP addresses are pinging a single computer. If, for instance, that computer is a web server, such a pattern might indicate that someone is attempting a denial of service attack. Table 2 shows the 8 parallel coordinates stimulus patterns.

Table 2: Eight Attack Stimuli in Parallel Coordinates

1, 1, 1, 1 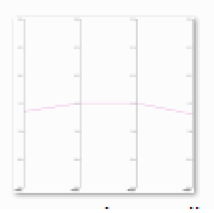	m,1, m, 1 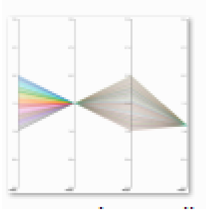	1, m, 1, 1, 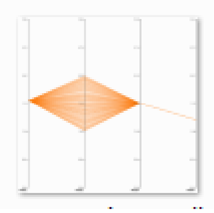	1, 1, m, 1 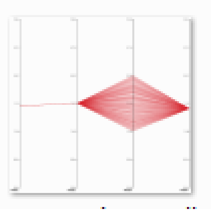
m, m, 1, 1 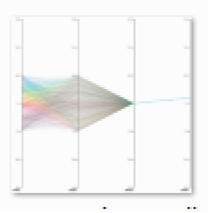	m, 1, 1, 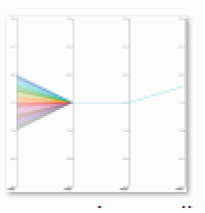 1	1, m, m, 1 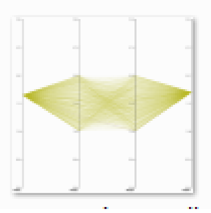	m, m, m, 1 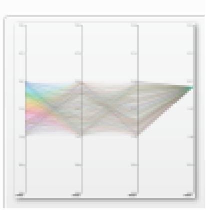

Note: Since the last axis always only contains one point, the patterns may be fully described by first three axes and ignoring the rightmost portion. The patterns are described using commonplace language such as brooms, bowties, houses, etc. For instance, 1,1,1,(1) is a straight line, m,1,m,(1) is a bowtie, 1,m,1,(1) is a diamond, 1,1,m,(1) is a broom with the handle to the left, m,m,1,(1) is a side-ways house with the roof to the right, m,1,1,(1) is a broom with the handle to the right, 1,m,m,(1) is a house with the roof to the left, and m,m,m,(1) is a rectangle.

For the scatter plots, we again mocked up many variations of variable assignments among the 4 representational dimensions (x, y, color, and shape). Since the rule category for each type of attack was always the same for a single attack, we set rule category to shape, which is the most difficult of the dimensions to discriminate (each of the 8 patterns were always drawn using the same shaped marker). Time is represented on the x axis, source IP is represented on the y axis, and destination IP is represented by color of marker. Table 3 shows the 8 scatter plot stimuli. The pattern in the 1,m,1,1 block, as in the prior example, might represent a denial of service attack - at 1 time (1 x value), many sources (many y values) sent traffic to one destination (1 color), and only 1 category of rules fired (1 shape of point markers).

Table 3: Eight Attack Stimuli in Scatter Plots

1, 1, 1, 1 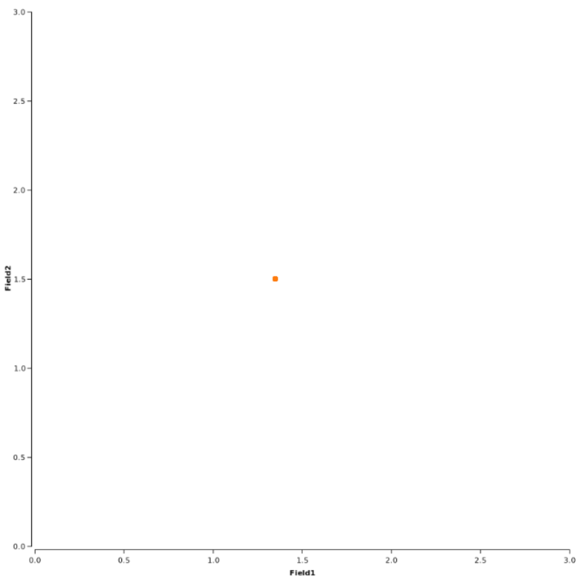 	m,1, m, 1 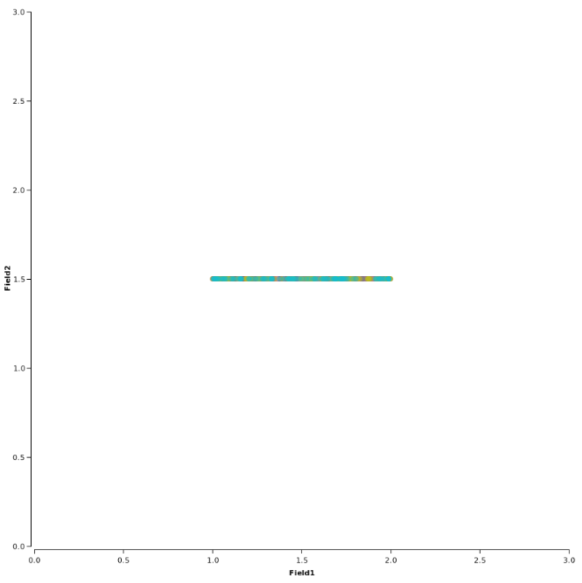

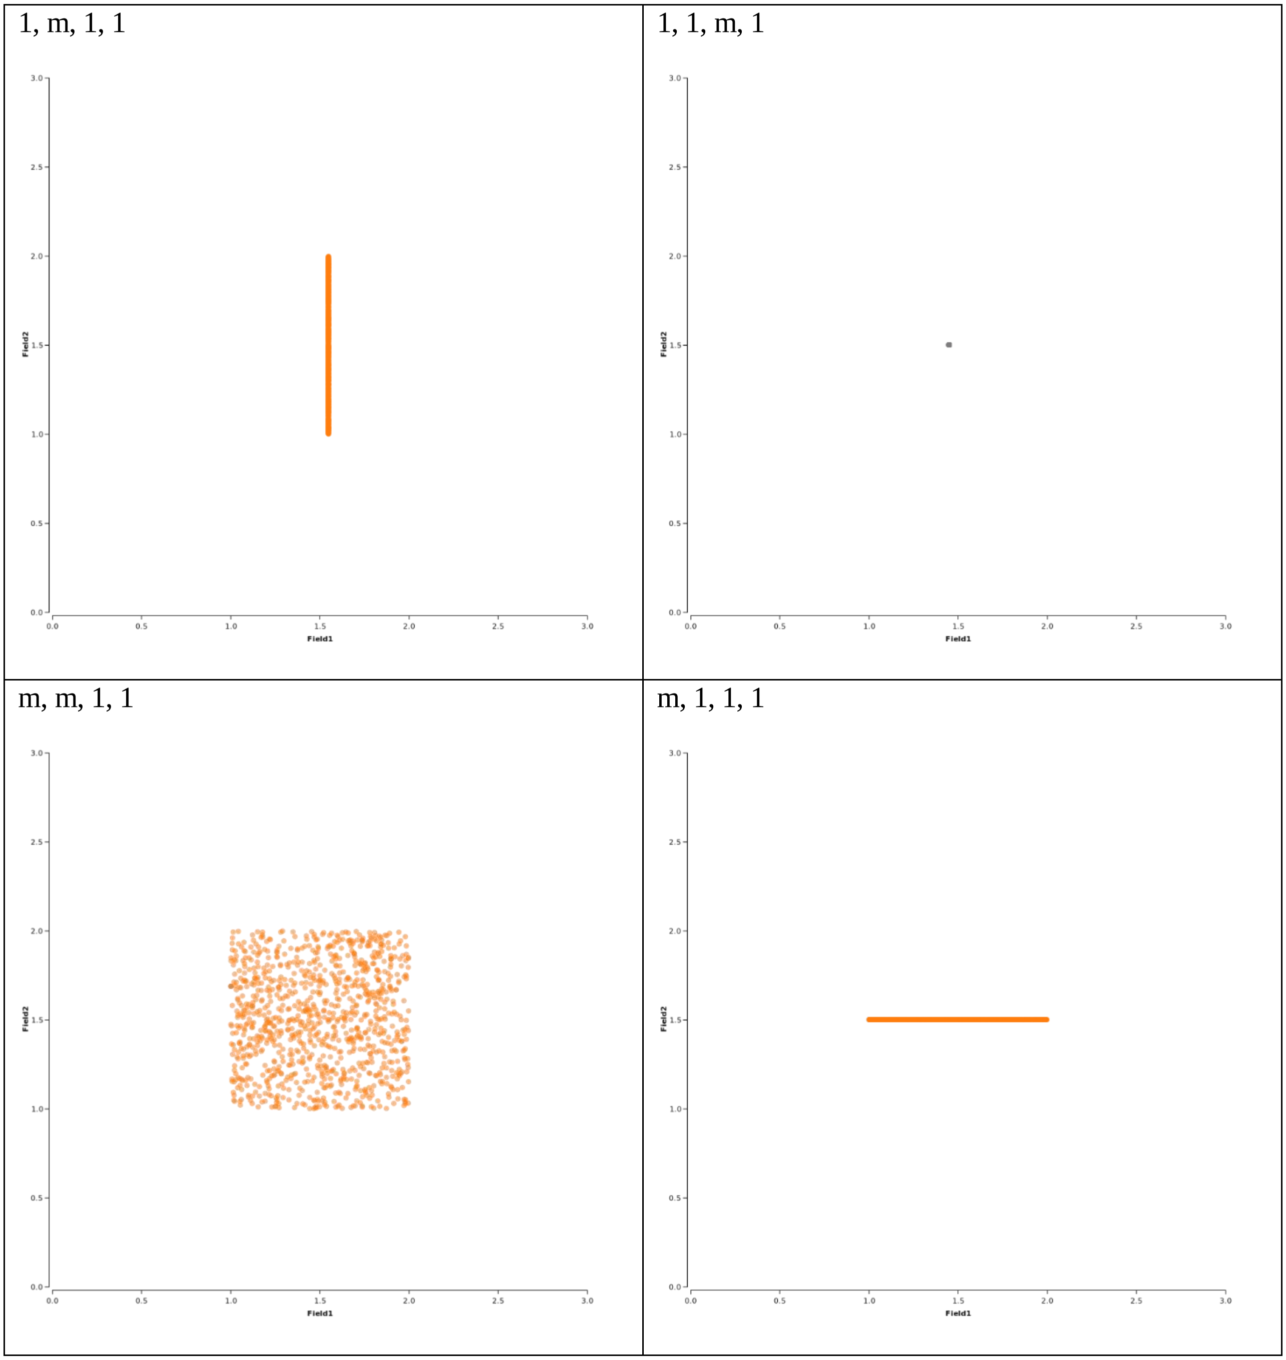
1, m, 1, 1
1, 1, m, 1
m, m, 1, 1
m, 1, 1, 1
Field1
Field2
0.0
0.5
1.0
1.5
2.0
2.5
3.0

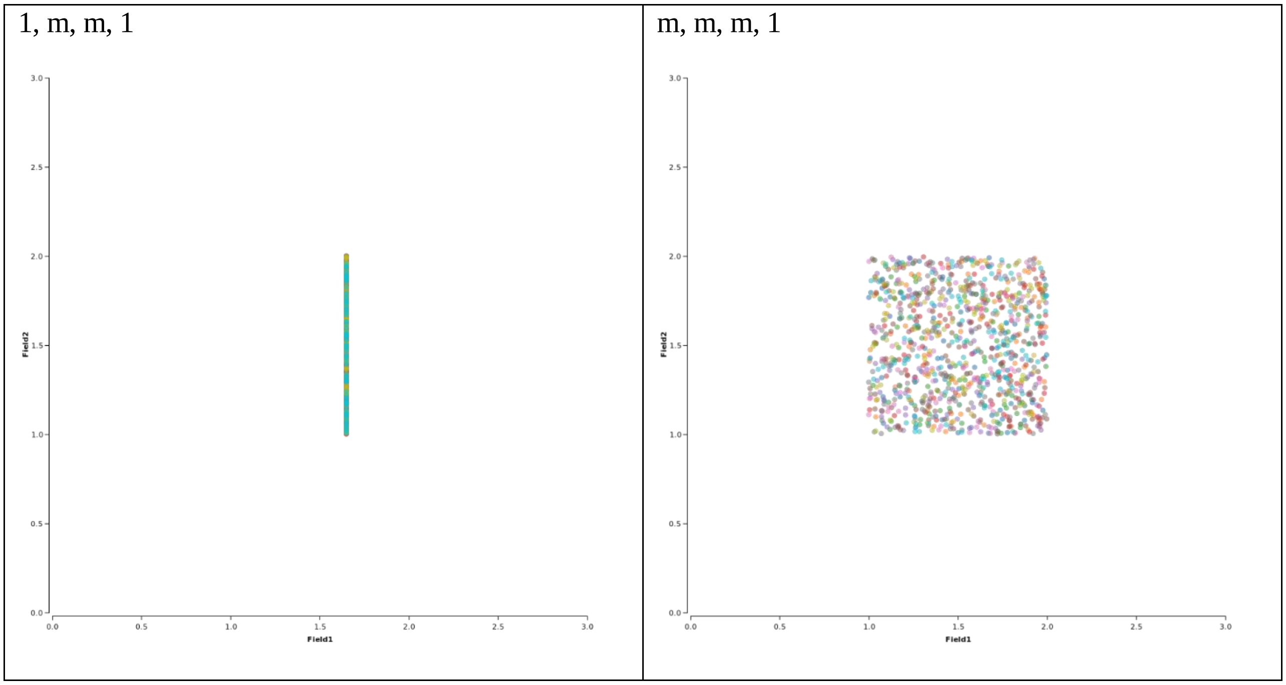

1, m, m, 1

m, m, m, 1

Note: Patterns here are 1,1,1,(1) is a single-colored marker in single location, m,1,m,(1) is multi-colored markers in a horizontal line, 1,m,1,(1) is single-colored markers in a vertical line, 1,1,m,(1) is multi-colored markers in single location [appears grey], m,m,1,(1) is single-colored markers in a square, m,1,1,(1) is single-colored markers in horizontal line, 1,m,m,(1) is multi-colored markers in a vertical line, and m,m,m,(1) is multi-colored markers in a square.

2.0 THE RESEARCH STUDY

With 8 distinct attack patterns and 2 visualization techniques, we had the basis for a study. The primary objective of the study was to compare the effectiveness of parallel coordinates and scatter plots as a means of portraying the data in IDS log files. More specifically, we were interested in whether these techniques would be effective in portraying attack patterns when the attacks were embedded in noise due to normal activity occurring at the same time as an attack, and when the patterns were distorted due to particular axis values representing different variable values. In other words, would defenders be able to pick out the attack patterns even if there were numerous attacks and/or among rule violations that were not indicative of an attack, as well as when the patterns were distorted by the position of the values along each axis? We, therefore, considered 2 factors in addition to the visualization type to be varied in this study. These factors included distortion of the pattern, and the amount of visual density each pattern contained both individually and when additional patterns were added to them (noise). Figure 4 shows both scatter plots and parallel coordinates with distortion and noise.

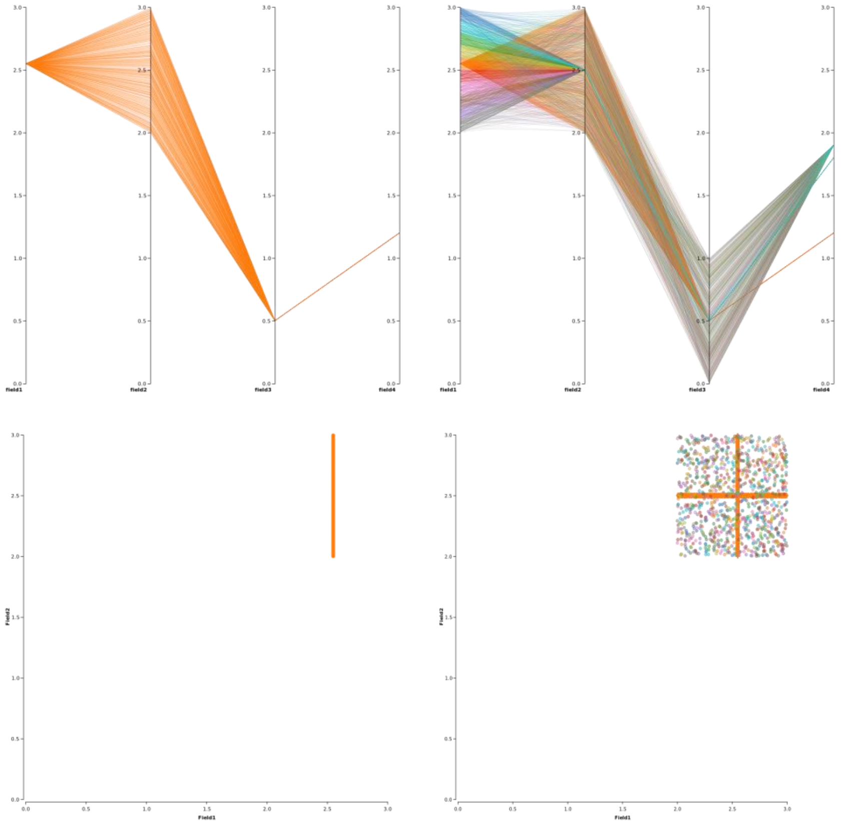

Figure 4: Parallel coordinates and scatter plots with distortion and noise. Top visualizations are parallel coordinates; bottom visualizations are scatter plots. Top left is the 1,m,1,1 (diamond) attack pattern with distortion level 2; top right is the 1,m,1,1 (diamond) attack pattern with noise patterns m,1,1,1 (broom with handle to right) and m,m,m,1 (rectangle); bottom left is 1,m,1,1 (single-colored vertical line) attack pattern with distortion level 2; bottom right is the 1,m,1,1(single-colored vertical line) attack pattern with additional noise patterns m,1,1,1 (single-colored horizontal line) and m,m,m,1 (multi-colored square).

2.1 Method

2.1.1 Participants

A total of 44 participants were recruited from Wright-Patterson Air Force Base and Wright State University. Wright-Patterson AFB participants were recruited via e-mail and flyers. Wright State Introductory Psychology students were recruited by way of the university's research participation system (a web-based system developed by Sona Systems and used by the Department of Psychology for students to sign up to participate in approved experiments). These students are required to participate in a minimum of 8 hours of research as part of the Introductory Psychology course and receive extra credit for further participation.

2.1.2 Stimuli

Stimuli for the study consisted of various representations of data in both parallel coordinates format and scatter plot format. To make the stimuli less cluttered and eliminate the instinct of people reading the legends, we removed all textual information from the stimuli (as shown in Tables 2 and 3 and Figure 4). The objective of the task was to have participants determine whether there was an attack pattern in each of the visual stimuli presented. The study employed 5 independent variables; visualization type (parallel coordinates vs. scatter plots), level of distortion (5 levels), level of complexity (a discrete ordinal variable determined by the sum of the visual density of each attack pattern and the pattern or patterns of noise displayed with it; range in this experiment is from 5 to 18), presentation order or "session" (parallel coordinates set first vs. scatter set first), and attack pattern set (which 4 of the 8 patterns were designated as attacks; 2 categories). In further explanation of this last variable there were 8 basic patterns; 4 were defined as "attack" and 4 were defined as "noise" for each of 2 sets. In order to ensure that results are not restricted to particular patterns of attack, we reversed the definitions (patterns that were considered noise in one case were defined as attacks in another and vice versa). The attack pattern sets were selected such that in one set, the overall complexity of the noise was greater than that of the attacks (simple attacks in complex noise) and in the other, complexity was greater for the attacks than the noise (complex attacks in simple noise). For instance, simple attacks were (refer to Tables 2 and 3) 1,1,1,1; m,1,m,1; 1,m,1,1; and 1,1,m,1; and complex noise for those attacks were m,m,1,1; m,1,1,1; 1,m,m,1; and m,m,m,1. This situation was reversed for the other attack pattern set (complex attack patterns were m,m,1,1; m,1,1,1; 1,m,m,1; m,m,m,1 and simple noise patterns were 1,1,1,1; m,1,m,1; 1,m,1,1; 1,1,m,1;). Since it would be difficult for participants to consider a particular pattern as an attack in one situation and as noise in another situation, we treated this variable as a between-groups variable (half of the subjects were told patterns 1-4 were attacks and patterns 5-8 were noise; the other half of the subjects were told that patterns 1-4 are noise and patterns 5-8 were attacks).

2.1.3 Experimental Design

The study utilized a mixed design; attack pattern set was varied between groups and visualization type, level of distortion, level of complexity, and session were varied within-subjects. Each block of trials contained 2 sets of 6 noise only pattern combinations by 5 levels of distortion plus 40 combinations of attack plus noise patterns by 5 levels of distortion, yielding 60 noise only and 200 attack plus noise trials for a total of 260 trials per participant. The dependent measures collected included time to complete each trial and the accuracy of the response. Each presentation was displayed for a maximum of 10 seconds, and participants were instructed to respond as quickly and accurately as possible. Each block took approximately 45 minutes on average. Assignment of the session and attack pattern sets were systematically varied (the 4 possible combinations were alternated) between participants, and trials were randomly presented within blocks.

2.1.4 Hypotheses

Overall, we were uncertain as to which visualization type would provide the best performance. The parallel coordinates had more visual content due to the inherent layout of the 4 axes, and every variable was presented spatially. However, the scatter plot patterns were more simplistic (single location markers, lines, and squares). Also, the scatter plots seemed to be more robust to distortion as changing axis values in the scatter plots simply relocated the stimuli in the scatter plot space (refer to Figure 4). Therefore, we anticipated an interaction between visualization type and distortion with scatter plots being less affected by distortion than parallel coordinates. In terms of complexity, the parallel coordinates were effected in a uniform way as increasing complexity resulted in overlaying patterns on each other, whereas for the scatter plots, since some of the patterns were single markers, increasing complexity might "hide" the single location patterns. Therefore, we hypothesized that the scatter plots would be more negatively affected as complexity increased. We also anticipated an interaction between complexity and attack pattern set as Attack Pattern Set A contained simpler attack patterns and Attack Pattern Set B contained more complex attack patterns, which meant that higher levels of complexity in B did not decrease the "signal to noise ratio" nearly as much as with A. Therefore, it was hypothesized that Attack Pattern Set B would be unaffected by higher complexity levels but Attack Pattern Set A would be affected by high complexity levels. We expected to see an effect for session (essentially a training effect), so we counterbalanced presentation order to control for systematic learning effects.

2.1.5 Procedure

Upon arrival, the research objectives and procedures were explained, and participants were asked to read and sign an informed consent form and were assigned a participant number. Participants were then administered a color vision test (Ishihara plates). Which visualization type (either parallel coordinates representation or scatter plot representation) was presented first and which set of attack patterns that they were assigned was dependent on their participant number. They were shown pictures that contained simple patterns representative of four cyber attacks as well as examples of distorted versions of those attacks. Once they were comfortable that they knew what the attacks looked like, they began a training block in which sixteen attack and noise patterns were presented at various levels of distortion, and they had to determine whether each was an attack or noise pattern. Feedback was provided. Training blocks were repeated until they were able to identify all attacks and noise with 100% accuracy, at which time data collection proceeded.

During each block of data collection, participants were simply asked whether an attack was present in the stimuli. Participants were informed that there was not a 50% chance that an attack was present so they should not make an effort to balance their responses. Upon completion of the first block of data collection, participants were offered a break. Then, the process of training and data collection was repeated for the other visualization type. Another break was offered, and then participants were asked to complete a brief questionnaire to record any challenges, observations, or comments they might have.

2.2 Results

2.2.1 Main Effects

A total of 44 participants were recruited from Wright-Patterson Air Force Base and Wright State University. Wright-Patterson AFB participants

2.2.1.1 Visualization type

With regard to accuracy, there was a small but statistically significant main effect of visualization type, $F(1,519) = 5.70$, $p < 0.05$. Overall, participants were slightly more accurate when determining whether attacks were present in scatter plot (M = 75.4%) than in parallel coordinates (M = 74.8%). There was also an effect of visualization type on response time, $F(1,519) = 2163.21$, $p < 0.05$. On average it took 2.84 seconds for participants to decide whether an attack was present in parallel coordinates visualizations and about 1.71 seconds to make the determination when viewing scatter plots. Since time varied negatively with accuracy, there was not a speed accuracy trade-off that needs to be considered. In general, participants who were more accurate tended to have shorter response times. Visualization type also interacted with several other variables; those interactions will be discussed below.

2.2.1.2 Other variables

There were main effects of distortion, both for response time, $F(4,519) = 19.90$, $p < 0.05$ and accuracy, $F(4,519) = 10.64$, $p < 0.05$ as well as on complexity for both response time, $F(13,519) = 21.18$, $p < 0.05$, and accuracy, $F(13,519) = 76.76$, $p < 0.05$. There also were main effects of attack pattern set on response time, $F(1,519) = 184.36$, $p < 0.05$, as well as on accuracy, $F(1,519) = 522.16$, $p < 0.05$. Finally, there were significant main effects of session for response time, $F(1,519) = 233.87$, $p < 0.05$, and accuracy, $F(1,519) = 4.57$, $p < 0.05$.

2.2.2 Two-Way Interactions

2.2.2.1 Visualization type by distortion

The effects of distortion significantly interacted with visualization type for both response time, $F(4,519) = 29.08$, $p < 0.05$ and accuracy, $F(4,519) = 15.20$, $p < 0.05$. Response times were consistently better with scatter plots than with parallel coordinates; this effect increased at higher levels of distortion. Also, parallel coordinates resulted in higher accuracies than scatter plots for low levels of distortion, but parallel coordinates accuracies fell below those of scatter plots at the highest level of distortion (Figure 5).

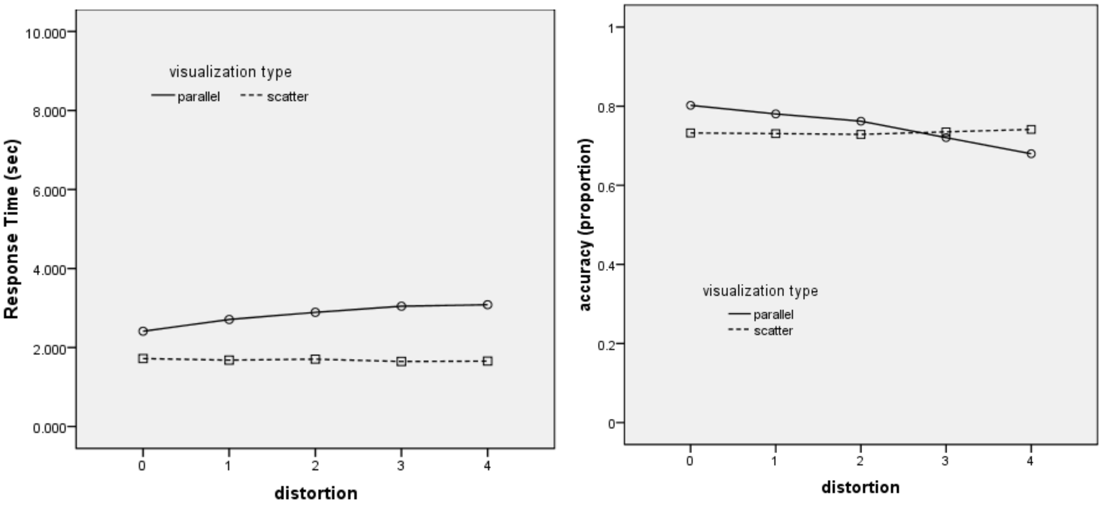

Figure 5: Interaction between visualization type and distortion for response time and accuracy.

2.2.2.2 Visualization type by complexity

Complexity interacted with visualization type for both response time, $F(13,519) = 12.77$, $p < 0.05$ and accuracy, $F(13,519) = 42.35$, $p < 0.05$. In general, complexity had a greater effect on response times of parallel coordinates than those of scatter plots. However, scatter plots showed much more variation in accuracy across levels of complexity (Figure 6).

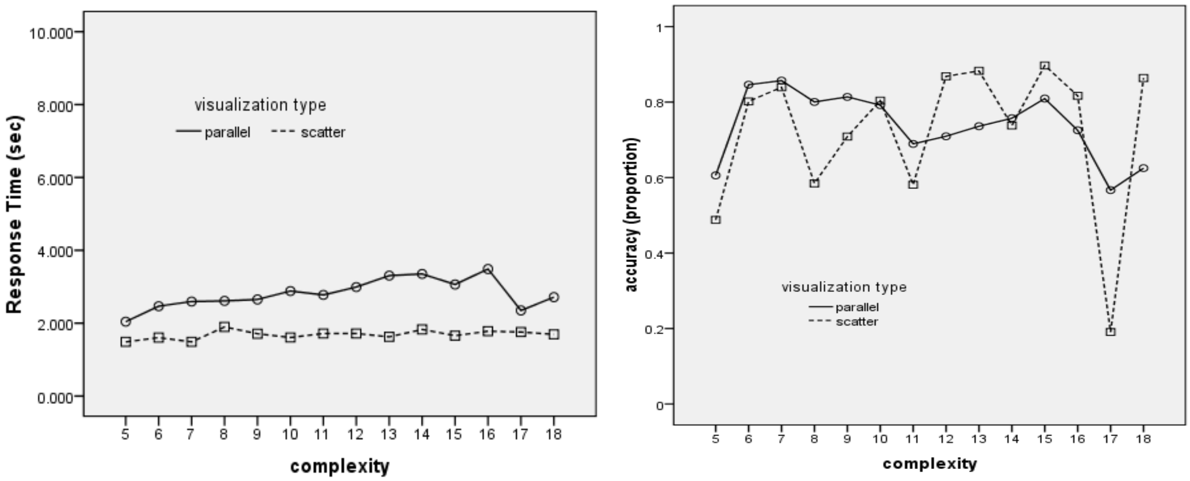

Figure 6: Interaction between visualization type and complexity for response time and accuracy.

2.2.2.3 Visualization type by attack pattern set

There was a significant two-way interaction between visualization type and attack pattern set for response time, $F(1,519) = 311.32$, $p < 0.05$, and accuracy $F(1,519) = 8.49$, $p < 0.05$. In terms of response time, the attack pattern set had a greater effect on parallel coordinates than on scatter plot performance. Specifically, it took longer for participants to detect complex attacks in simple noise (Condition B) than detecting simple attacks in complex noise (Condition A) when using the parallel coordinates but attack pattern set had no effect on response time performance when using the scatter plots. However, for the condition where a simple attack was embedded in more complex noise (Condition A), performance accuracy was higher with parallel coordinates, which was generally not true for Condition B (Figure 7).

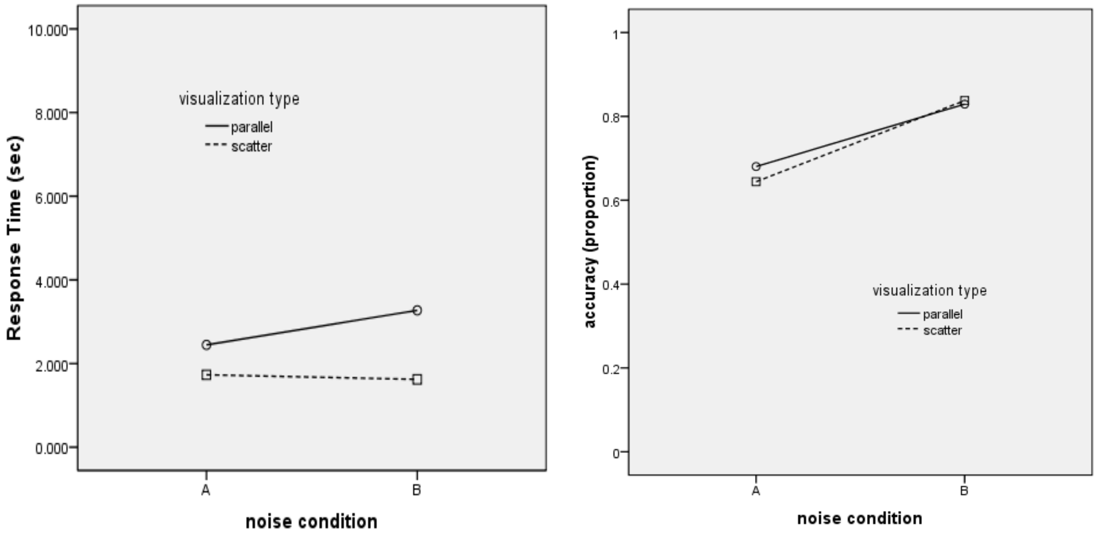

Figure 7: Interaction between visualization type with attack pattern set (noise condition) for response

time and accuracy. Condition A is simple attacks in complex noise and condition B is complex attacks in simple noise.

2.2.2.4 Visualization type by attack session

There was a significant interaction between visualization type and session in terms of both response time, *F*(1,519) = 301.38, *p* < 0.05, and accuracy, *F*(1,519) = 27.68, *p* < 0.05. Response time differences were less for Session 2 than for Session 1, in general, but the difference was chiefly from lower response times with parallel coordinates for Session 2; session did not make a difference in reaction times with scatter plots (Figure 8). In terms of accuracy, Session 1 participants were more accurate with parallel coordinates, but were slightly more accurate with scatter plots when comparing Session 2 accuracies.

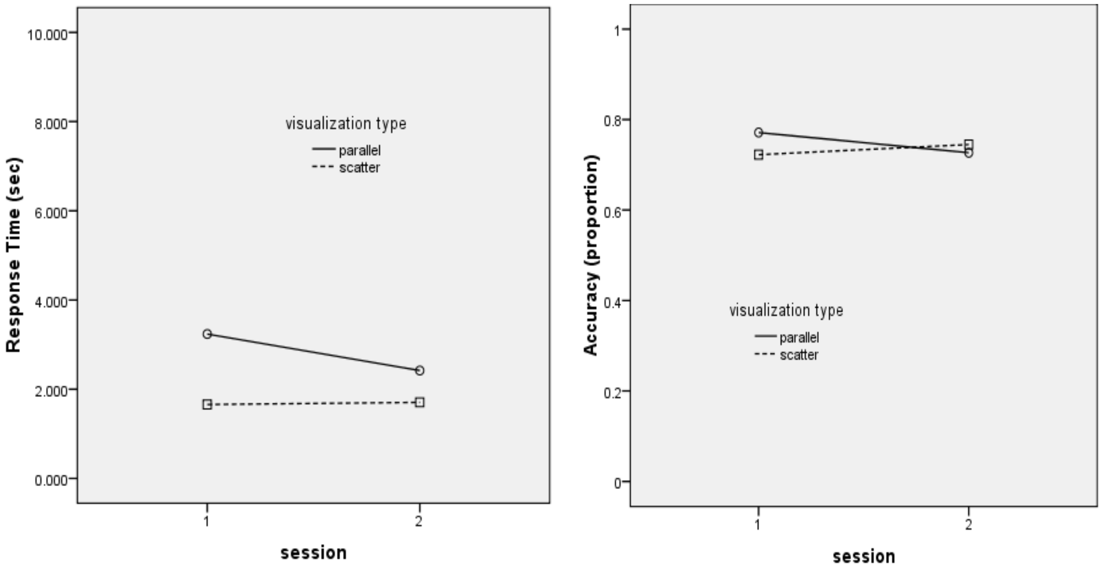

Figure 8: Interaction between visualization type and session (presentation order) for response time and accuracy.

2.2.2 Higher Order Interactions

2.2.2.1 Visualization type by complexity by attack pattern set

The There were significant three-way interaction between visualization type, complexity, and attack pattern set for both response time, *F*(11,519) = 12.34, *p* < 0.05, and accuracy, *F*(11,519) = 19.96, *p* < 0.05. The effects of complexity on response times for parallel coordinates were less for simple attacks in complex noise than vice versa (Figure 9). What this effect generally means is that as more noise was added, complex attack patterns

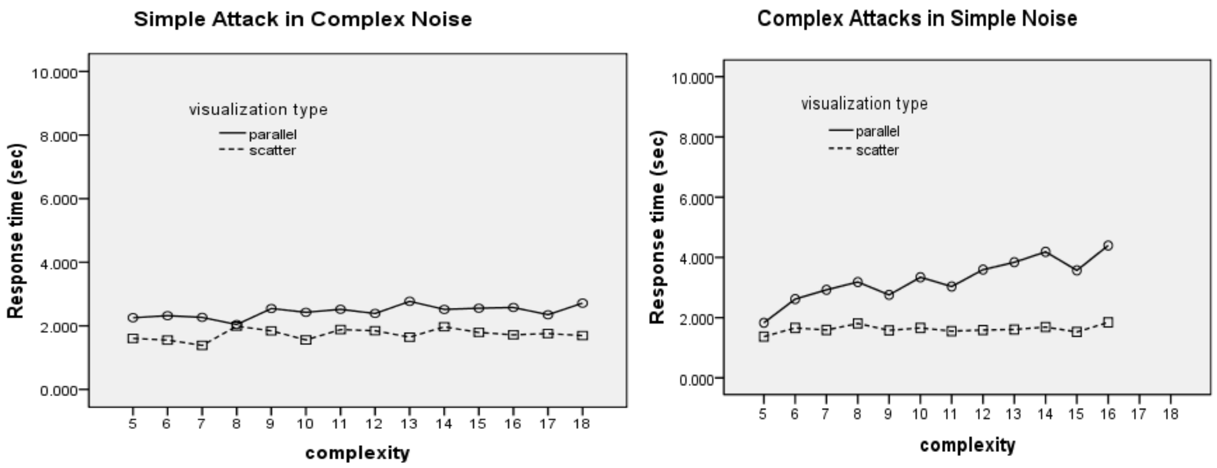

Figure 9: Differential interactions of visualization type and complexity for the two attack pattern sets with regard to response time.

were more difficult to find but simple attacks were not. In terms of accuracy, there was a high level of variability for the scatter plots when simple attacks were presented in complex noise versus when complex attacks were presented in simple noise (Figure 10).

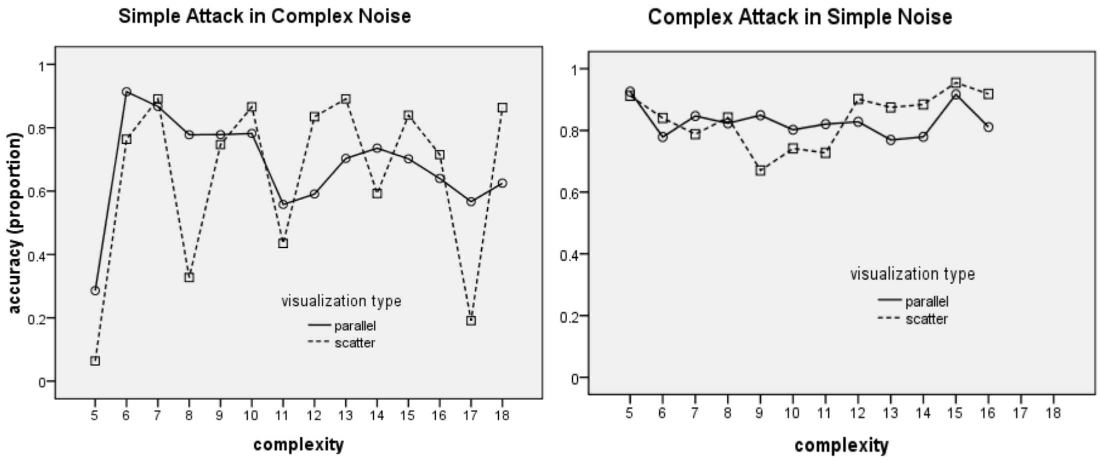

Figure 10: Differential interactions of visualization type and complexity for the two attack pattern sets with regard to accuracy.

2.2.2.2 Visualization type by attack pattern set by session

Analyses revealed a three-way interaction among visualization type, attack pattern set, and session on response time, $F(1,519) = 15.65$, $p < 0.05$. The differences in response time for the parallel coordinates visualization by session were smaller for simple attacks in complex noise than for complex attacks in simple noise, but did not vary for scatter plots (Figure 11).

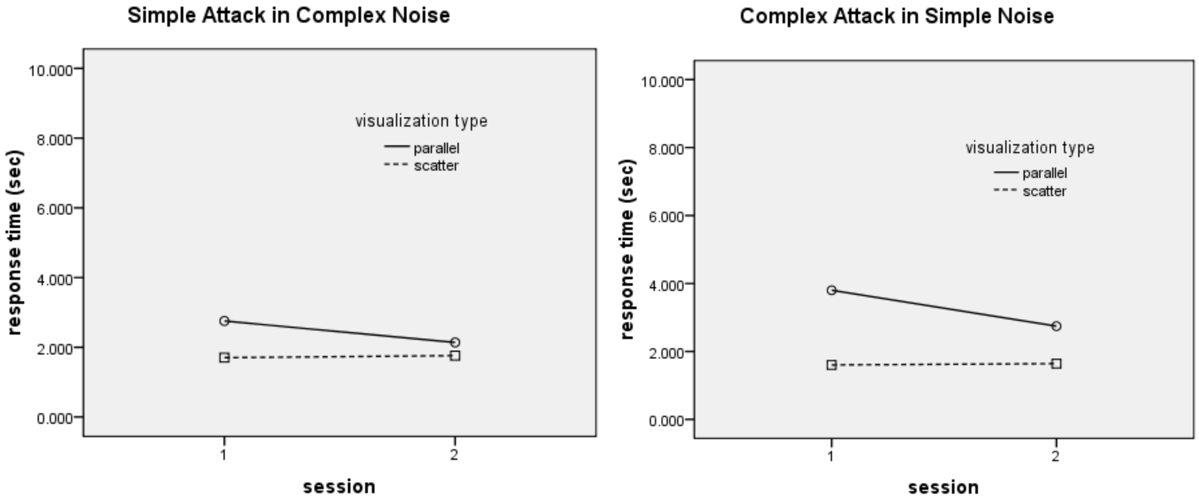

Figure 11: Differential interactions of visualization type and session for the two attack pattern sets with regard to response time.

2.3 Discussion

In terms of the main effects for visualization type, participants performed more quickly and accurately with the scatter plots than with parallel coordinates. However, these effects, although statistically significant, were fairly small from a practical standpoint (scatter plots were 0.6% more accurate and 1.14 seconds faster than parallel coordinates). Questionnaire data revealed that 28 of the 44 participants (64%) thought they performed better with the scatter plots; 13 or the 44 (29%) thought they performed better with the parallel coordinates; and the remaining 3 participants (7%) didn't think they performed differently regardless of the visualization type. Interestingly, during the course of the experiment, it became fairly clear that participants understood the one-to-many relationships better as they were depicted in parallel coordinates than in scatter plots. However, this artifact did not manifest itself in a significant performance advantage for the parallel coordinates.

In terms of the effects of distortion, the significant interaction between visualization type and distortion showed that scatter plots allowed for better performance in terms of response time for all levels of distortion and better performance in terms of accuracy for the higher levels of distortion, which is in line with our hypothesis that scatter plots would be less affected by distortion. Again, distortion manifested itself differently in the two visualization types (refer to Figure 4) resulting in a lesser effect on performance (faster response times and higher accuracies for higher levels of distortion) for the scatter plots (refer to Figure 5).

Since the variables involved in the remaining two-way interactions were also present in the three-way interactions, we will proceed directly to the discussion of the three-way interactions. Regarding the three-way interaction among visualization type, complexity, and attack pattern set, parallel coordinates response times were slightly more affected by complexity than the scatter plots; however this effect differed by attack pattern set. Oddly, for simple attacks in complex noise, response times did not vary much by complexity for either visualization type (refer to Figure 9). When one considers some of the simple attack patterns in the scatter plot condition and how the addition of multi-colored lines and squares to single-location attack patterns may lead to confusion, these results are surprising and contrary to our hypotheses. For instance [Refer to Table 3 for the following example] , one of the simple noise patterns is the 1,1,1,1 (single-colored marker in a single location for the scatter plot condition). To increase overall complexity of the stimuli, complex noise had to be added to this simple attack pattern. This addition could take the form of the m,m,m,1 (multi-colored markers in a square) noise pattern. It seems that it would be challenging to locate a single-colored single location marker in a field of multi-colored markers in a square, but perhaps the fact that the shapes of the markers were different for the two patterns, made that task reasonable. Finding a singular circular marker in a field of diamonds was not as challenging as it sounds. Also, consider adding the complex noise pattern m,m,m,1 (multi-colored markers in square shape) to the simple attack pattern m,1,1,1 (single-colored marker in a horizontal line). As in the previous example, the shapes of the markers for the two patterns would be different, but if the line was drawn first, then the square, there might be some overlapping of multi-colored markers from the square on the single-colored markers of the horizontal lime, making the line look multi-colored. Again, the different marker shapes might have been a strong cue reducing the uncertainty of this situation. While this might explain the consistent time variable, this theory is challenged when looking at the accuracy data of the same 3-way interaction.

Figure 10 shows a high level of variability in accuracy with regard to the scatter plot data when looking at simple attacks in complex noise. Specifically, complexity levels 5, 8, 11, and 18 produce lower accuracies than other complexity levels. When looking at the specific stimuli that generate the levels of complexity stated above for the scatter plots, the attack patterns that participants were to identify in these cases were 1,1,1,1 and 1,1,m,1 (the single location marker, both single-colored or multi-colored). Note, the other two simple attack stimuli are m,1,m,1 and 1,m,1,1 (lines – which contained more visual content than the single location markers). Therefore, this effect may be more related to the specific patterns that participants were looking for versus the level of complexity variable we calculated for each. So to summarize the meaning of this interaction in terms of the scatter plots, participants responded fairly consistently regardless of attack pattern set, but their accuracy suffered when they were looking for 2 particular simple attack patterns.

With the parallel coordinates, response times for complex attacks in simple noise increased with complexity such that they were generally higher than with simple attacks in complex noise (refer to Figure 9). Accuracies, although not always greater, were more stable for parallel coordinates, particularly for simple attacks in complex noise (refer to Figure 10). The fact that, in the parallel coordinates format, the visual stimuli of the attack pattern with additional noise patterns that are required to increase the complexity score resulted in a total overlapping of the patterns on each other, made the task of finding complex patterns in simple noise more difficult (refer to Figure 4). For instance [Refer to Table 2 for the following example], think about the situation where the attack pattern is the m,m,m,1 (rectangle) combined with any of the simpler noise patterns, such as m, 1, 1, 1 (broom with handle to right). Participants many have been spending too much time trying to identify the individual patterns present in the stimuli before making a choice. For instance, in the above example, they may have spent more time trying to determine if the many lines on the first axis belonged to the many colored lines of the rectangle or the many lines of the broom before responding. The complete overlapping of these patterns may have caused participants a bigger challenge then deciphering the complexity of the scatter plots. Also consider a visual stimuli that might include both brooms (1,1,m,1 and m,1,1,1). Although the 1,1,m,1 pattern would be a single color and the other m,1,1,1 pattern would be multi-colored, at a quick glance, one might mistake the two brooms for the bowtie. The color coding for the parallel coordinates in this experiment was redundant with the first axis (time). It occurred to us during the course of the experiment that the patterns would be more separable if the color had instead been redundant with the fourth axis (rule category) thereby making each pattern one color similar to the ones in Table 2 that have a 1 in the first axis (1,1,1,1 [line]; 1,m,1,1 [diamond]; 1,1,m,1 [broom with handle to the left]; and 1,m,m,1 [house with roof to left]). So in summary, the three-way interaction between visualization type, complexity, and attack pattern set challenged our hypotheses and led us to some additional designs to test in future research.

In terms of the three-way interaction among visualization type, attack pattern set, and session, participants tended to perform better in the second session than in the first (refer to Figure 11). Although we did not predict an interaction, it is not particularly surprising that the effects of session were greater for parallel coordinates. It appears that participant tried to go faster in the second session when the second session was with parallel coordinates. Perhaps, this tendency is due to the fact that they had already experienced faster response times in the first session with scatter plots and this influenced the pace they set for themselves in the second session. Response times were generally closer and differed slightly less by session for simple attacks in complex noise than for complex attacks in simple noise. The fact that when parallel coordinates were presented in the first session resulted in longer response times is supported by the training data. Participants were required to go through the training session for each visualization type until they reached 100% accuracy on all 16 trials. Training data showed that the 44 participants ran through 112 training trails before reaching 100% accuracy for the parallel coordinates versus 88 training trials before reaching 100% accuracy for the scatter plots. Therefore, it took participants 2.5 practice runs on average versus 2 practice runs to reach 100% accuracy for the parallel coordinates as compared to the scatter plots.

3.0 CONCLUSIONS

This research shows that, based on knowledge elicitation activities with experts, it is possible to represent IDS log data in static visual representations which can portray visual signatures of attack patterns. The subsequent experiment showed that this is a reasonable approach to represent IDS attack data; that participants can quickly and accurately identify these visual signatures. While overall performance showed that scatter plots consistently provided the faster performance times, parallel coordinates showed more consistent accuracy performance. Based on the complex interactions revealed in this study, ideas for optimizing both visualization techniques should be pursued. Also, future research should investigate additional visualization techniques to determine how best to portray complex, abstract cyber information to operators in a meaningful way.

4.0 REFERENCES

Aucsmith, D. W. (2011). Rethinking cyber defense. *High Frontier, 7*(3), 35-37.

Barford, P., Dacier, M., Dietterich, T. G., Fredrikson, M., Griffin, J., Jajodia, S., Jha, S., Li, J., Liu, P., Ning, P., Ou, X., Song, D., Strater, L., Swarup, V., Tadda, G., Wang, C., & Yen, J. (2010). Cyber SA: Situational awareness for cyber defense. In S. Jajodia, P. Liu, V. Swarup, & C. Wang (Eds.) *Cyber Situational Awareness* (pp. 3-13). US: Springer.

Bryant, W. D. (2013). Cyberspace Superiority A Conceptual Model. *Air & Space Power Journal, 27*(6), 25-44.

Conti, G. (2007). *Security Data Visualizations: Graphical Techniques for Network Analysis.* San Francisco: No Starch Press.

Costlow, T. (2013). Commercial tools playing expanded role in cyberops. *Defense Systems*, *7*(6), 12-14.

D’Amico, A., Tesone, D., Whitley, K., O’Brien, B, & Roth, E. (2008). *Understanding the Cyber Defender: A Cognitive Task Analysis of Information Assurance Analysts* (Report No. CSA-CTA-1). Northport, NY: Secure Decisions.

D’Amico, A., & Whitley, K. (2007). The real work of computer network defense analysts. *Proceedings of the 4th International Workshop on Visualization for Cyber Security* (pp. 19-37).

D’Amico, A., Whitley, K., Tesone, D., O’Brien, B., & Roth, E. (2005). Achieving cyber defense situational awareness: A cognitive task analysis of information assurance analysts. *Proceedings of the Human*

Factors and Ergonomics Society 49th Annual Meeting (pp. 229-233).

Fua, Y, H., Ward, M. O., & Rundensteiner, E. A. (1999). Hierarchical parallel coordinates for exploration of large datasets. *Proceedings of the IEEE 1999 Conference on Visualization* (pp. 43-50).

Goodall, J. R. (2005). User requirements and design of a visualization for intrusion detection analysis. *Proceedings of the 2005 IEEE Workshop on Information Assurance and Security* (pp. 394-401).

Goodall, J. R. (2009). Visualization is better! A comparative evaluation. *Proceedings of the 6th International Workshop on Visualization for Cyber Security*. (pp. 57-68).

Grinstein, G., Cook, K., Havig, P., Liggett, K., Nebesh, B., Whiting, M., Whitley, K., & Koneci, S. (2011). VAST 2011 challenge: Cyber security and epidemic. *IEEE VAST 2011*. (pp. 299-301)

Harrison, L., & Lu, A. (2012, November). The future of security visualization: Lessons from network visualization. *IEEE Network, 26*(6), 6-11.

Inselberg, A., & Dimsdale, B. (1991). Parallel coordinates. In A. Klinger (Ed.) *Human-Machine Interactive Systems* (pp. 199-233). US: Springer.

Jackson, W. (2012). How to design a network that fights for itself. *Government Computer News, 31*(12), 18-21.

Jackson, W. (2013). Surviving your next denial-of-service disaster. *Government Computer News, 32*(2), 24-27.

Kilgore, R., Godwin, A., Hogan, C., Davis, A., Pfautz, J., Woods, D. D., Branlat, M., Kaufman, H. (2012). Tangible trustworthiness for mixed-initiative network defense (Final Report from SBIR Phase I). Cambridge, MA: Charles River Analytics.

Koike, H. & Ohno, K. (2004). SnortView: Visualization system of snort logs. *Proceedings of the 2004 ACM Workshop on Visualization and Data Mining for Computer Security* (pp. 143-147).

Langton, J. T., & Newey, B. (2010). Evaluation of current visualization tools for cyber security. *SPIE Defense, Security, and Sensing* (pp. 10-20).

Leopold, G. (2013). Dempsey says cyberattacks are the new normal. *Defense Systems, 7*(5), 28.

Livnat, Y., Agutter, J., Moon, S., Erbacher, R. F., & Foresti. S. (2005). A visualization paradigm for network intrusion detection. *Proceedings of the 2005 IEEE Workshop on Information Assurance and Security* (pp. 92-99).

Mahoney, S., Roth, E., Steinke, K., Pfautz, J., Wu, C., & Farry, M. (2010). A cognitive task analysis for cyber situational awareness. *Proceedings of the Human Factors and Ergonomics Society 54th Annual Meeting* (pp. 279-283).

Marty, R. (2009). *Applied Security Visualization*. Boston: Addison-Wesley.

McCaney, K. (2014). 2015 budget proposal reflects the impact of cyber operations. *Defense Systems, 8*(2), 28.

Ponemon Institute. (2013). *2013 Cost of Cyber Crime Study: United States,*

Sponsored by HP Enterprise Security. Retrieved from http://media.scmagazine.com/documents/54/2013_us_ccc_report_final_6-1_13455.pdf

Roth, E. M., Patterson, E. S., & Mumaw, R. J. (2002). Cognitive engineering. In J. J. Marciniak (Ed.) *Encyclopedia of Software Engineering* (2nd ed., Vol. 1, pp. 163-179). New York: Wiley.

Rubin, J. (1994) . *Handbook of Usability Testing: How to Plan, Design, and Conduct Effective Tests*. New York: Wiley.

Shiravi, H., Shiravi, A., & Ghorbani, A. A. (2012). A survey of visualization systems for network security. *IEEE Transactions on Visualization and Computer Graphics, 18*(8), 1313-1329.

Shneiderman, B. (1996). The eyes have it: A task by data type taxonomy for information visualization.

Proceedings of the IEEE Symposium on Visual Languages (pp. 336-343).

Thompson, R. S., Rantanen, E. M., Yurcik, W., & Bailey, B. P. (2007). Command line or pretty lines?: Comparing textual and visual interfaces for intrusion detection. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 1205).

Van Wijk, J. J., & Van Selow, E. R. (1999). Cluster and calendar based visualization of time series data. *IEEE Information Visualization* (pp. 4-9).

White hats to the rescue: Law-abiding hackers are helping businesses to fight off the bad guys. (2014). Economist. Retrieved from http://www.economist.com/news/business/21596984-law-abiding-hackers-are-helping-businesses-fight-bad-guys-white-hats-rescue