Modeling the Effects of Base-rates on Cyber Threat Detection Performance

Varun Dutt (varundutt@cmu.edu)
Young-Suk Ahn (ysahn@altenia.com)
Noam Ben-Asher (noamba@andrew.cmu.edu)
Cleotilde Gonzalez (coty@cmu.edu)

Dynamic Decision Making Laboratory, Department of Social and Decision Sciences
Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA

Abstract

Cyber attacks cause major disruptions of online operations, and might lead to data and revenue loss. Thus, appropriately training security analysts, human decision makers who are in charge of protecting the infrastructure of a corporate network from cyber attacks, on different frequencies of cyber threats (base-rates) is indispensable to improving their on-job performance. However, little is currently known about how training analysts on different cyber attacks, that differ in the base-rate of cyber-threats, affects their on-job performance in a highly dynamic environment, while confronting novel transfer conditions. We report a laboratory experiment where human participants are trained on two different cyber-threat base-rates, high and low, and are transferred to an intermediate base-rate level of threats. The experiment helps us to develop an understanding of the situational attributes that participants attend to during their detection of cyber-threats. A linear model that is based upon participants' attended attributes and calibrated to the two base-rates during training does well to capture the performance during transfer. We use the calibrated model to generate predictions in novel real-world transfer conditions that contain a low cyber-threat base-rate and a shorter training period.

Keywords: cyber-threat; linear model; security analyst; training; transfer; base-rate.

Introduction

Cyber attacks, i.e., the disruption of computers' normal functioning and the loss of sensitive information in a network through malicious network events (cyber-threats), are becoming widespread. With "Anonymous" and other threats to corporate and national security, guarding against cyber attacks is becoming a significant part of IT governance, especially because most government agencies and private companies have moved to online systems (Sideman, 2011). Recently, President Barack Obama declared that the "cyber-threat is one of the most serious economic and national security challenges we face as a nation" (2011). According to his office, the nation's cyber-security strategy is twofold: (1) to improve our resilience to cyber incidents; and (2) to reduce the cyber-threat. To meet these goals, the role of a security analyst (called the "analyst" hereafter), a human decision maker who is in charge of protecting the online infrastructure of a corporate network from random or organized cyber attacks, is indispensable (Jajodia, Liu, Swarup, & Wang, 2010).

Given that the threat of cyber attacks is growing, there is an urgent need to emphasize training programs for analysts that will acquaint them with different kinds of attacks. For example, the U.S. Department of Homeland Security (DHS) has recently started offering a weeklong training program to help analysts learn how to deal with intrusions into their computer networks (Zakaria, 2011). In this training program, the DHS uses a scenario that contains industrial espionage: a fictitious company ACME has built a new chemical product and another company Barney Advanced Domestic (BAD) Chemicals tries to steal its "secret sauce" and disrupt operations to put ACME out of business (Zakaria, 2011).

Although training programs like the one by DHS are an important step towards improving cyber-threat detection, there is little literature on how training analysts on different kinds of cyber attacks will influence their on-job performance (Dutt, Ahn, & Gonzalez, 2011). One aspect of analyst training is base-rate: In the real world, cyber threats are likely to occur with a low base-rate and given this rare occurrence, analysts might underweight them by relying on their personal experience (Hertwig & Erev, 2009). Another aspect of analysts' training is length. Lengthy training is likely to benefit analysts, but also becomes costly in resources and time (Kanellis, 2006). In this regard, little is known about how the training period might influence analysts' performance. The current documentation about analysts' decisions is sensitive information that is classified, as it could be exploited by attackers (D'Amico et al., 2005). Thus, in the absence of real human data, literature has proposed a simulation approach towards evaluating the effects of training manipulations on performance at transfer (Dutt, Ahn, & Gonzalez, 2011). For example, Dutt, Ahn, and Gonzalez (2011) have proposed a computational model based upon the Instance-Based Learning Theory (IBLT; and, "IBL model" hereafter; Gonzalez & Dutt, 2011), to generate predictions about the effects of training simulated analysts with different base-rates on their transfer performance. They pre-populated the model's memory with experiences of a threat-prone base-rate (90% threats and 10% non-threats) and a nonthreat-prone base-rate (10% threats and 90% non-threats).

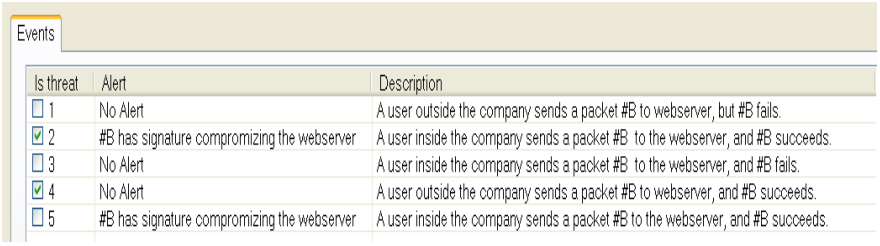

Figure 1: A typical trial showing events' description and accompanying alert information. Participants marked an event as a threat by checking the "Is threat" box against the event's description.

The model with a threat-prone memory possessed a greater hit rate and a smaller false-alarm rate (i.e., greater accuracy) compared with the model with a nonthreat-prone memory. These results have been replicated by Dutt and Gonzalez (in press). Thus, when human analysts are exposed to higher base-rates during their training, they are more likely to perform better at transfer. However, these model predictions are limited by the lack of empirical validity from real human observations.

In this paper, we build upon existing literature and report a laboratory experiment that aims to empirically evaluate analysts' cyber-threat detection accuracy when they are trained on different base-rates before transferring to novel conditions. Furthermore, we evaluate the attributes to which participants attend during their training with self-reported strategies. Using this information, we propose a linear model that classifies cyber events as threats and non-threats. The proposed model fits the existing data well. Finally, we use this model to generate predictions in novel real-world transfer conditions that contain a low cyber-threat base-rate and a shorter training period. The low base-rate is likely to be representative of situations encountered in the real world, where cyber-threats rarely occur in a certain period of time (Jajodia et al., 2010). Shorter training periods might be plausible, given the resource and time costs imposed by lengthy training sessions.

Experiment: Effects of Cyber-Threat Base-rates on Threat Detection Performance

We report an experiment where we train participants in two different conditions, high and low, that differ in the proportion of cyber-threats present during training. We evaluate the effects of their training with base-rates on their transfer performance. Our main motivation is to find the attributes to which participants attend most. According to prediction from the IBL model proposed by Dutt, Ahn, and Gonzalez (2011), we expect better performance (i.e., greater hit rates and smaller false-alarm rates) from participants who are trained in conditions involving higher base-rates. That is because higher base-rates provide participants with more frequent opportunities to formalize and test different hypotheses regarding what defines a threat.

The experiment involved two between-subjects conditions, low (N=20) and high (N=18), that differed in the proportion of cyber-threats presented in a trial. Both conditions contained 10 training trials followed immediately by a transfer trial. Each trial contained 25 network events that were presented to participants sequentially one at a time. An event included a description and might be accompanied with an alert. The alerts were generated by an intrusion detection system (IDS) that might indicate whether these events were threats or not. The IDS systems generated both false-positives and false-negatives. Figure 1 shows a snapshot of a trial with 5 events (some with alerts and some without). The participants' main goal in the task was to correctly classify each event as a threat or a non-threat by checking or unchecking the corresponding "Is threat" box for each event. A new event appeared in the window after every four seconds. Participants could go back and check/uncheck any previously presented event during the duration of the trial, but not after a trial had ended.

During training in the low base-rate condition, 12% of events (=3) were actual threats in each training trial. For the high base-rate condition, 52% of events (=13) were threats in each training trials. A threat was defined as any event that by its description was: (1) initiated by a user outside the company; and (2) against which there was an alert generated. However, participants were not told this definition and they were expected to discover it with practice. In both conditions, a single transfer trial was presented and 32% of events were threats (=8). Thus, the base-rate in the transfer trial was in between the low and high conditions (participants were not told about any base-rates). Participants' performance was evaluated in terms of hit and false-alarm rates at the end of each trial.

Instructions informed participants about how they would be paid based upon their performance. After participants read the instructions, they started the experiment with the first training trial. At the end of each training trial, they were asked to write down the strategy that guided their events' classification. After participants submitted their written explanations, they were informed of the number of hits, misses, false-alarms, and correct-rejections they made in the last trial; along with their current and total earnings based upon performance. However, they were not shown which exact events in the last trial were actual threats and non-threats. Participants were compensated with $5 as base payment. In addition, participants earned 1 cent for each threat and non-threat correctly classified and lost 1 cent for each threat and non-threat incorrectly classified during training. During transfer, participants earned 3 cents for each threat and non-threat correctly classified and lost 3 cents for each threat and non-threat incorrectly classified. After participants had completed their experiment, they were paid and thanked for their time.

Results

We expected superior performance (i.e., a greater hit-rate and a smaller false-alarm rate) in the high condition compared with the low condition. Figure 2 shows the aggregated hit and false-alarm

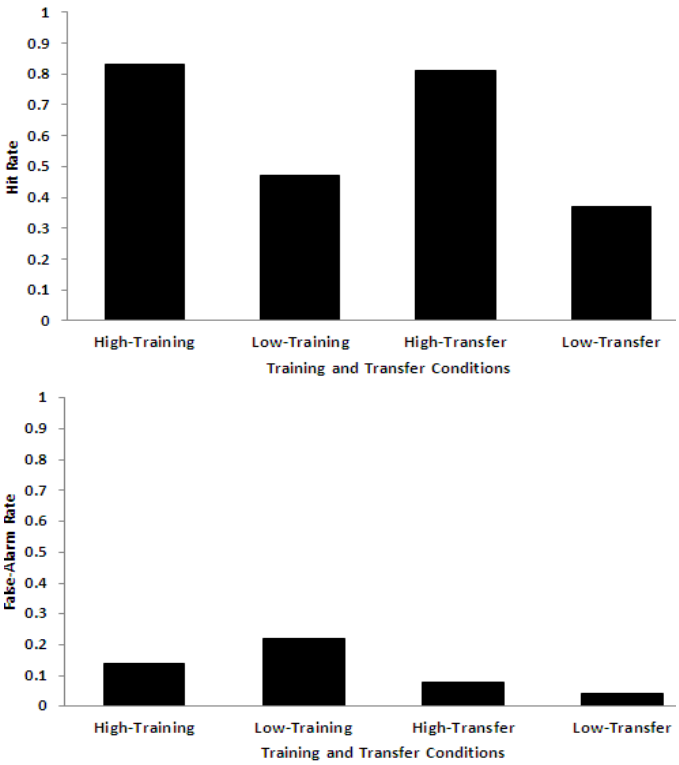

Figure 2: The human hit rate and false-alarm rate in the low and high conditions during training and transfer.

rates in both conditions during training and at transfer. As can be seen in Figure 2, during training, the average hit rate in the high condition (83%) was significantly greater than that in the low condition (47%), $t(36) = -3.89$, $p < .001$. The same relationship existed for the hit rate at transfer between the low condition (37%) and the high condition (81%), $t(36) = -4.95$, $p < .001$. Although false-alarm rates generally appeared to be greater in the low condition compared with the high condition, these differences were insignificant, both during training ($t(36) = 1.74$, *ns*) and at transfer ($t(36) = -1.23$, *ns*). Overall, these results are in agreement with our expectation of superior performance in the high condition.

Attention to Situational Attributes

In order to gain deeper understanding of the overall performance in the detection task, we analyzed the explanations that participants provided about their classifications at the end of each trial. Although there was diversity present in their explanations across trials, five main categories emerged in each condition for which a majority of explanations could be categorized into (45% and 35% in the low and high conditions). Table 1 provides a description of these categories with statistical differences between the two conditions. We believe participants' explanations under different categories were based upon the descriptions of events and corresponding alerts that were presented to them during training and transfer trials (see Figure 1 for the format of presentation). Four out of the five categories, "word *malicious* (present/absent)," "fileserver (attacked/not attacked)," "operation (successful/unsuccessful)," and "user (inside/outside)," overlapped between the two conditions. Two categories, "alerts (present/absent)" and "files (manipulated/not-manipulated)," were uncommon across conditions. For the overlapping categories, the proportions of usage were somewhat different for the two conditions. However, a significant difference was found only for the "word *malicious*" category, its proportions were greater in the low condition compared with the high condition. As the "word *malicious*" category did not always classify threats correctly in this experiment, the difference in attention between the two conditions might partly explain participants' superior performance in the high condition.

As seen in Table 1, the two leading categories of rules that participants used in the low condition were user location and the presence of the word *malicious*. Similarly, in the high condition, the two categories that participants attended to most were user location and the presence of an alert. The user location and alerts represent categories that could correctly classify an event in this task. In the high condition, participants paid attention to both of these categories among the top five; in the low condition, however, they only paid attention to the user category and ignored the alerts category. This difference in attention likely accounts for the superior performance in the high condition.

A Linear Model: Validating Attention to Situational Attributes

The literature on heuristics and biases shows that linear models have been highly successful in explaining human behavior and leads to approximate correct responses that are more accurate than even expert judgments (Dawes, 1979; Goldberg, 1970). For example, when Dawes and Corrigan (1974) applied different linear models to five different datasets to predict a criterion, an equal weighting linear model (the simplest assumption of linearity) outperformed all other competing models. If the categories and weights (i.e., the relative proportion) reported in Table 1 are representative of our behavioral findings, then they should produce a close fit to the hit and false-alarm rates observed in human data when these categories are simulated in a linear model. To explore this idea further, we developed a stochastic linear model consisting of derived categories and attention weights.

Stochastic Linear Model

We represent the binary decision process of classifying events as threats and non-threats in a trial according to the following rule:

If the model's Outcome > *threshold*, then classify the event as threat; otherwise, classify the event as a non-threat. (1)

The *threshold* is a free parameter that is calibrated in the model. The Outcome is defined according to a linear model:

For the low condition: Outcome = 0.23 * Word *malicious* + 0.15 * Fileserver + 0.11 * Operation + 0.34 * User + 0.17 * Files

Table 1: The top five attention categories and their descriptions in the low and high conditions.

Categories	Low Condition	High Condition	Statistic
	Relative Proportion (%)[1]	Relative Proportion (%)[1]	Low versus High Condition
Word malicious (present/absent) If the description of the event contains the word malicious, then treat the event as a threat; otherwise treat it as a non-threat.	22.51	7.69	$t(36) = -2.34$, $p < .05$, $r = 0.36$
Fileserver (attacked/not attacked) If the description of the event contains an attack on fileserver, then treat the event as a threat; otherwise treat it as a non-threat.	15.37	8.13	$t(36) = -1.45$, *ns*, $r = 0.23$
Operation (successful/unsuccessful) If the description of the event contains execution of an operation successfully, then treat the event as a threat; otherwise treat it as a non-threat.	10.52	11.13	$t(36) = -0.09$, *ns*, $r = 0.01$
User (inside/outside) If the description of the event contains a user generating the event from outside the company, then treat the event as a threat; otherwise treat it as a non-threat.	34.01	53.43	$t(36) = -1.57$, *ns*, $r = 0.25$
Alerts (present/absent) If the description of the event is accompanied by an alert, then treat the event as a threat; otherwise treat it as a non-threat.	-	19.59	-
Files (manipulated/not-manipulated) If the description of the event details files being manipulated, then treat the event as a threat; otherwise treat it as a non-threat.	17.56	-	-
Total	**100%**	**100%**	

Note: The *r* is the effect size.

For the high condition: Outcome = 0.08 * Word *malicious* + 0.08 * Fileserver + 0.11 * Operation + 0.53 * User + 0.20 * Alerts (2)

Where the "Word *malicious*," "Fileserver," "Operation," "User," "Files," and "Alerts" are dummy variables (taking a values of 0 or 1) corresponding to the five categories in Table 1. The weights (i.e., coefficients) that multiply the dummy variables are the relative proportions of the respective categories reported in Table 1. The model is "stochastic" because the exact value of a dummy variable (0 or 1) for each event in a trial depends upon comparing $U(0, 1)$ with the dummy variable's *attention probability* parameter. The rule for paying attention to a dummy variable in the model is the following:

If the category is applicable (i.e., present in the event's a description) to the network event and $U(0, 1) \leq$ dummy variable's *attention probability*, then the dummy variable equals 1; otherwise, the dummy variable equals 0 (3)

Each dummy variable's *attention probability* is a free parameter that is calibrated in the model. This parameter represents whether a model participant pays attention to a category when it is present in the event's description or in the accompanying alert. Furthermore, if the category is attended to, then the Outcome takes a weighted contribution of the category into the binary decision. Also, more than one category could be attended to (dummy variable = 1) for an encountered event. Therefore, the model captures the property that human participants might stochastically pay attention to multiple categories for an encountered event.

Parameter Calibration

The model's free parameters, i.e., each dummy variable's *attention probability* and *threshold*, were calibrated to human data in the low and high conditions, separately. Calibrating the model to human data means running it in the same training conditions experienced by human participants to find the parameters values which minimize the sum of mean squared deviations (Sum of MSDs) between the model's hit and false-alarm rates and human hit and false-alarm rates, respectively. The smaller the sum of MSDs, the better the model's ability to capture human behavior is. The model was calibrated separately to training trials in the low and high conditions using a genetic algorithm program. To calibrate the model, we varied the *threshold* and *attention probability* parameters between 0.0 and 1.0 (their minimum and maximum values). The model was run using the same number of simulated participants as the number of human participants that participated in the two conditions.

Table 2 presents a summary of the calibrated parameters and the models' performance (MSD) in the two conditions. The model performed reasonably well to capture the hit and false-alarm rates in both conditions; however, it was slightly better in capturing the false-alarm rate than the hit rate. The calibrated value of the *threshold* parameter was found to be similar in both conditions (close to 30%). Moreover, based upon the *attention probability* parameters, the model seemed to frequently attend to the correct categories, "user" and "alerts," in the high condition to make decisions. In fact, the model's attention to the user category was greater in the high condition (.76) than in the low condition (.50). Figure 3 shows the model fits to human data in both conditions for training and transfer.

The MSDs between the model and human hit rates at transfer in the low and high conditions were 0.0003 and 0.0002, respectively. Similarly, the MSDs between the model and human false-alarm rates at transfer in the low and high conditions were 0.0001 and 0.0052, respectively. These MSDs are very small and therefore the model provides a good approximation to the human transfer performance.

Table 2: Summary of calibration of the linear model during training.

Condition	Parameters	MSD (Hit Rate)	MSD (False-alarm Rate)
High	*threshold*= .31; *attention probability* (Word *malicious*) = .10; *attention probability* (Fileserver) = .99; *attention probability* (Operation) = .44; *attention probability* (User) = .76; *attention probability* (Alerts) = .53	0.0112	0.0016
Low	*threshold*= .33; *attention probability* (Word *malicious*) = .53; *attention probability* (Fileserver) = .60; *attention probability* (Operation) = .83; *attention probability* (User) = .50; *attention probability* (Files) = .30	0.0130	0.0071

Predictions in Novel Transfer Conditions

Typically, one could think of cyber-threats as rare events in the real world (Jajodia et al., 2010). If these threats occur rarely at transfer, then participants trained in the low condition might benefit more from their training. That is because, the experiences gained in the low base-rate training condition are likely to be more suited to the rare transfer condition compared with those gained in the high condition. One way to test this expectation is by creating a rarer transfer trial (i.e., whose base-rate is less than that in the low condition's training trials and that in the original transfer trial). One such rare transfer trial could have a threat base-rate of 4% (i.e., only 1 event out of 25 events is an actual threat in the trial).

In the transfer trial of our experiment, the human hit rate in the high condition was 81% while it was 37% in the low condition (i.e., a gap of 44%; see Figure 2). In the transfer trial, the human false-alarm rate in the high condition was 8% while it was 4% in the low condition (i.e., a gap of 4%). Thus, we expect the accuracy to be greater in the high condition than in the low condition; however, based upon the discussion above, it is also possible that people trained in the low condition will perform better in the rare transfer trial (with a 4% threat base-rate). Therefore, we expect the model's hit rate predictions in the low condition to increase and in the high condition to decrease at transfer, closing the overall gap. However, for the same rare transfer trial, we expect the model's false-alarm rate predictions in the low condition to decrease and in the high condition to increase at transfer, widening the overall gap.

Predictions generated from our calibrated model were in agreement with these expectations. The model's performance in the rare transfer trial showed a hit rate in the high and low conditions to be 76% and 46%, respectively. Therefore, its overall performance in terms of hit rate was superior in the high condition compared with the low condition; however, the gap between the hit rates in the two conditions was reduced to 29%. Similarly, the model's performance showed a false-alarm rate in the high and low conditions of the rare transfer trial to be 16% and 4%, respectively (i.e., an increased gap of 11% as expected). Overall, the gap between hit and false-alarm rates in the low and high conditions moved in the direction as expected above.

Another aspect to consider is the length of analysts' training sessions. In the real world, lengthy training might become costly because of resources and time (Kanellis, 2006). In such situations, one method to save costs is to reduce the training length and evaluate cyber-threat detection accuracy at transfer. Here, we derive predictions

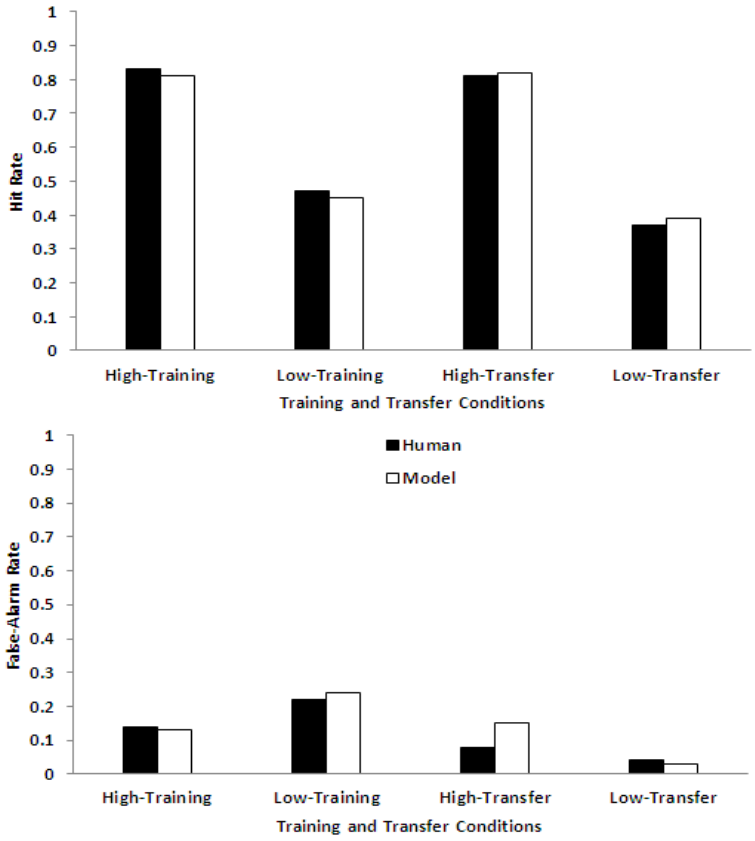

Figure 3: The model and human hit rate and false-alarm rate in the low and high conditions during training and transfer.

Table 3: Summary of half calibration length of the linear model during training.

Condition	Parameters	MSD (Hit Rate)	MSD (False-alarm Rate)
High	*threshold*= .33; *attention probability* (Word *malicious*) = .20; *attention probability* (Fileserver) = .71; *attention probability* (Operation) = .93; *attention probability* (User) = .96; *attention probability* (Alerts) = .53	0.0149	0.0004
Low	*threshold*= .29; *attention probability* (Word *malicious*) = .44; *attention probability* (Fileserver) = .71; *attention probability* (Operation) = .94; *attention probability* (User) = .46; *attention probability* (Files) = .88	0.0208	0.0073

from the linear model in a situation where the training in both the low and high conditions is reduced by halving its original length (i.e., only first 5 trials long), and the model is then transferred to the rare transfer trial with a 4% base-rate. Table 3 provides the summary of the calibrated parameters and the MSDs to the halved training length.

Again in the high condition, the free parameters have greater values in the "user" and "alert" categories compared to those in the low condition. Furthermore, the MSDs in Table 3 are slightly higher compared with those reported for full-length training in Table 2. At transfer, the calibrated model's hit and false-alarm rates were 67% and 29% in the high condition, respectively; whereas, these rates were 37% and 14% in the low condition, respectively. When these proportions are compared with those reported above for the rare transfer trial with full-length training, we find a drop in hit rate as well as an increase in false-alarm rate in both conditions. Thus, our model predictions suggest that reducing training conditions by half their original length is likely to save time and costs, but also likely to decrease the analysts' detection accuracy at transfer.

Discussion

In this paper, we evaluated the effects of training security analysts in conditions with different threat base-rates (low and high) and transferring them to novel conditions (that were either in-between those encountered during training or possessing a very low base-rate). We found that their transfer performance is superior when their training environments provide them with lengthy training and higher threat base-rates. That is likely because higher base-rates allow participants to form improved hypotheses about threats that they could test during their training and transfer performance (Dutt & Gonzalez, in press). This reasoning is clearly reflected in the greater proportions of calibrated attention probability for the "user" and "alert" attributes (i.e., the attributes that reveal the ground truth) in the high condition compared with in the low condition.

Our results suggest that any training interventions for analysts should pay close attention to how the base-rate of threats compare to their actual work conditions. Also, the length of training (e.g., weeklong or half a week) is likely to influence analysts' learning and performance at transfer. Thus, the training length is likely to affect their performance in actual work conditions. Generally, it would be advisable to keep the training extended in length, as well as train analysts on scenarios that makes them experience a high threat base-rate. In fact, the linear model could be used to derive the optimal length of training session for a desired level of accuracy. Although we can only speculate, but our results are also likely to be valid for other emergency situations like training miners for a low-probability mine emergency, or training air-traffic controllers for low probability air accidents.

In this paper, we contribute to the growing literature on cyber security by evaluating the benefits and costs of training analysts in scenarios that differ in threat base-rates. Although base-rates were different, other aspects of the scenario (i.e., the sequence of attack, computers compromised, etc.) were identical. Thus, future research is likely to benefit from our results by manipulating other aspects of attack scenarios and evaluating the influence on training and transfer interventions. Also, as the linear model might be more mathematical than cognitive in its formulation, future research is also likely to benefit by comparing how other cognitive models, which use memory and activation (including some data-mining algorithms), perform compared to the linear model. Finally, we also contribute a method of going from an experiment about detecting cyber threats to developing a cognitive model based upon participants' self-reported strategies. This "model discovery" approach that uses human data to construct cognitive models might provide useful insights for alternative modeling approaches to model development.

Acknowledgements

This research was a part of a MURI Award on Cyber Situation Awareness (MURI; #W911NF-09-1-0525) from ARO to Cleotilde Gonzalez.

References

Obama, B (2011, May). Remarks by the President on securing our nation's cyber infrastructure. Retrieved from http://www.whitehouse.gov/the_press_office/Remarks-by-the-President-on-Securing-Our-Nations-Cyber-Infrastructure/

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 Annual Meeting*, 49(3), 229-233.

Dawes, R. M. (1979). The robust beauty of improper linear models. *American Psychologist*, 34, 571-582.

Dawes, R.M., & Corrigan, B. (1974). Linear models in decision making. *Psychological Bulletin*, *81*, 95–106.

Dutt, V., Ahn, Y. S.., & Gonzalez, C. (2011). Cyber situation awareness: Modeling the security analyst in a cyber-attack scenario through Instance-Based Learning. *Lecture Notes in Computer Science*, *6818*, 280-292. doi: 10.1007/978-3-642-22348-8_24.

Dutt, V., & Gonzalez, C. (in press). Cyber situation awareness: Modeling the security analyst in a cyber attack scenario through Instance-based Learning. In C. Onwubiko & T. Owens (Eds.), *Situational Awareness in Computer Network Defense: Principles, Methods and Applications*. Hershey, PA: IGI Global.

Goldberg, L. R. (1970). Man vs. model of man: A rationale, plus some evidence for a method of improving clinical inferences. *Psychological Bulletin*, *73*(6), 422–432.

Gonzalez, C., & Dutt, V. (2011). Instance-based learning: Integrating decisions from experience in sampling and repeated choice paradigms. *Psychological Review*, *118*(4), 523-551.

Hertwig, R., & Erev, I. (2009). The description-experience gap in risky choice. *Trends in Cognitive Sciences*, *13*(12), 517-523.

Jajodia, S., Liu, P., Swarup, V., & Wang, C. (2010). *Cyber situational awareness*. New York: Springer.

Kanellis, P. (2006). *Digital crime and forensic science in cyberspace*. New York: Idea Group Publishing.

Sideman, A. (2011, February 23). Agencies must determine computer security teams in face of potential federal shutdown. *Federal Computer Week*. Retrieved from http://fcw.com/articles/2011/02/23/agencies-must-determine-computer-security-teams-in-face-of-shutdown.aspx

Zakaria, T. (2011, October 1). Government simulates cyber attack for training. *Reuters*. Retrieved from http://www.reuters.com/article/2011/10/01/us-usa-cyber-idaho-idUSTRE78T08B20111001