

Chapter 1

Cyber SA: Situational Awareness for Cyber Defense

P. Barford, M. Dacier, T. G. Dietterich, M. Fredrikson, J. Giffin, S. Jajodia, S. Jha, J. Li, P. Liu, P. Ning, X. Ou, D. Song, L. Strater, V. Swarup, G. Tadda, C. Wang, and J. Yen

1.1 Scope of the Cyber SA Problem

Situation Awareness (SA) for cyber defense consists of at least seven aspects:

1. Be aware of the current situation. This aspect can also be called *situation perception*. Situation perception includes both *situation recognition* and *identification*. Situation identification can include identifying the type of attack (recognition is only recognizing that an attack is occurring), the source (who, what) of an attack, the target of an attack, etc. Situation perception is beyond intrusion detection. Intrusion detection is a very primitive element of this aspect. An IDS (intrusion detection system) is usually only a sensor, it neither identifies nor recognizes an attack but simply identifies an event that may be part of an attack once that event adds to a recognition or identification activity.
2. Be aware of the impact of the attack. This aspect can also be called *impact assessment*. There are two parts to impact assessment: 1) assessment of current impact (damage assessment) and 2) assessment of future impact (if the attacker continues on this path or more general if the activity of interest continues - what is the impact?). Vulnerability analysis is also largely an aspect of impact assessment (provides knowledge of us and enables projection of future impact). Assessment of future impact also involves threat assessment.

Paul Barford, University of Wisconsin · Marc Dacier, Symantec · Thomas G. Dietterich, Oregon State University · Matt Fredrikson, University of Wisconsin · Jon Giffin, Georgia Institute of Technology · Sushil Jajodia, George Mason University · Somesh Jha, University of Wisconsin · Jason Li, IAI Inc. · Peng Liu, Pennsylvania State University · Peng Ning, North Carolina State University · Xinming Ou, Kansas State University · Dawn Song, University of California, Berkeley · Laura Strater, SA Technologies, Inc. · Vipin Swarup, MITRE · George Tadda, Air Force Research Laboratory Rome NY · Cliff Wang, Army Research Office · John Yen, Pennsylvania State University

3. Be aware of how situations evolve. Situation tracking is a major component of this aspect.
4. Be aware of actor (adversary) behavior. A major component of this aspect is attack trend and intent analysis, which are more oriented towards the behaviors of an adversary or actor(s) within a situation than with the situation itself.
5. Be aware of why and how the current situation is caused. This aspect includes causality analysis (via back-tracking) and forensics.
6. Be aware of the *quality* (and trustworthiness) of the collected situation awareness information items and the knowledge-intelligence-decisions derived from these information items. The quality metrics include *truthfulness* (or soundness), *completeness*, and *freshness*. This aspect can also be viewed as part of situation perception or more specifically recognition.
7. Assess plausible futures of the current situation. This involves a multitude of technologies for projecting future possible actions/activities of an adversary, paths the adversary might take, and then constraining the possible futures into those that are plausible. This constraining requires an understanding of adversary intent, opportunity, and capability (knowledge of them) as well as an understanding of blue vulnerabilities, etc. (knowledge of “us”).

Without losing generality, cyber situation awareness can be viewed as a three-phase process [5]: situation recognition (including Aspects 1, 6, and 7), situation comprehension (including Aspects 2, 4, and 5), and situation projection (including Aspect 3).

Situation awareness is gained by a system, which is usually the (cyber-physical) *system* being threatened by random or organized cyber attacks. Although the ultimate “dream” system is one that can gain self-awareness (and do self-protection) without involving any humans in the loop, this vision is still very distant from the current reality, and there still does not exist a tangible roadmap to achieve this vision (in a practical way). In this paper, we view *human decision makers* as an indispensable “component” of the system gaining situation awareness. Practical cyber SA systems include not only hardware sensors (e.g., a network interface card) and “smart” computer programs (e.g., programs that can learn attack signatures), but also mental processes of human beings making advanced decisions [1, 3].

Finally, cyber situation awareness can be gained at multiple *abstraction levels*: (raw) data are collected at the lower levels and at higher levels, as data is converted to more abstract information. Otherwise, data collected at the lowest levels can easily overwhelm the cognitive capacity of human decision makers. Situation awareness based solely on low level data is clearly insufficient.

The following aspects are typically not included in cyber SA, but they and the aforementioned cyber SA aspects complement each other in achieving the overall goal of cyber defense.

- Identification of better response plans and actions. This aspect could be called *planning*. This aspect stays in the boundary between situation awareness and situation response, during which the planned course of action will be taken. Planning often involves estimating the effects of a response plan before the planned actions are taken. Planning, responses, and actions are all command and control functions (decide and act) and are not typically included in SA. However, without SA one can't effectively do response plans and actions.
- Made decisions on the course of action to take. Situation Awareness enables a decision maker's awareness of a situation and their understanding of the situation up to the point the decision is made. Once a decision is reached, planning and execution (of the response actions) occur.

1.2 Background

Regarding the state of the art of cyber situation awareness, our main observations are as follows:

- Cyber SA systems and Physical SA systems have fundamental differences. For instance, Physical SA systems rely on specific hardware sensors and sensor signal processing techniques, but neither the physical sensors nor the specific signal processing techniques play an essential role in Cyber SA systems (although there is research that has looked at applying signal processing techniques to analyze network traffic and trends). (Cyber SA systems rely on cyber sensors such as IDS', log file sensors, anti-virus systems, malware detectors, and firewalls; they all produce events at a higher level of abstraction than raw network packets.) For another instance, the cyber situation evolving speed is usually orders of magnitude quicker than in physical situation evolution. Finally, cyber attacks/situations have unique semantics.
- Existing approaches to gain cyber situation-awareness consist of vulnerability analysis (using attack graphs), intrusion detection and alert correlation, attack trend analysis, causality analysis and forensics (e.g., backtracking intrusions), taint and information flow analysis, damage assessment (using dependency graphs), and intrusion response. These approaches however only work at the lower (abstraction) levels. Higher level situation-awareness analyses are still done *manually* by a human analyst, which makes it labor-intensive, time-consuming, and error-prone.
- Although researchers have recently started to address the cognitive needs of decision makers, there is still a big gap between human analysts' mental model and the capability of existing cyber situation-awareness tools.
- Existing approaches need to handle *uncertainty* better.

- Uncertainty in perceived data could lead to distorted situation awareness. For example, attack graph analysis toolkits are designed to do deterministic attack consequence estimation. In real time cyber situation-awareness, such consequence estimates could be very misleading due to various uncertainties. Alert correlation techniques cannot handle the inherent uncertainties associated with inaccurate interpretations of intrusion detection sensor reports (such inaccurate interpretations lead to false positives/negatives in determining whether an IDS alert corresponds to an attack).
- Lack of data or complete knowledge may raise additional uncertainty management issues. For example, lack of data leads to incomplete knowledge of “us”. Such incompleteness may be caused by imperfect information about system configurations, incomplete sensor deployment, etc.
- Existing approaches lack the reasoning and learning capabilities required by gaining full situation-awareness for cyber defense.
- The seven aspects of cyber situation awareness (see Section 1.1) have been treated as separate problems, but full cyber situation awareness requires all these aspects to be integrated into one solution. Such a solution is in general still missing. Furthermore, looking beyond cyber SA and considering how cyber SA solutions complement the other cyber defense technologies, cyber SA activities need to be better integrated with effect achieving or environment influencing activities (e.g., intrusion response activities).

1.3 Research Goals

At a high level, the basic objectives of a comprehensive cyber SA research agenda may be the ones listed below.

- The objective is to develop new algorithms that will (a) greatly enhance machines’ intelligence in gaining self-situation-awareness so that one day machines could protect themselves, and (b) automate human decision maker’s cognitive situation-awareness processes.
- If successful, the systems being-protected will recognize and learn about evolving situations, generate and reason about situation response plans and actions, and automatically respond to intrusions.

1.4 Research Agenda

1.4.1 Principles and Rationales

The following principles and rationales can help us identify the important research concentration areas and issues.

Rational 1: Before machines have sufficient artificial intelligence, information systems need to be protected by humans who are able to gain full situation awareness.

Rational 2: Cyber situation-awareness has at least two purposes: (a) enhance machines' intelligence in gaining situation awareness so that one day machines can gain self-awareness and do self protection; (b) automate human decision maker's cognitive situation-awareness processes.

Principle 1: Full situation-awareness for cyber defense requires a holistic methodology to synthesize perception, understanding and projection.

Principle 2: Information systems with full situation-awareness must manage uncertainty (e.g., through hypotheses and reasoning).

Principle 3: Cyber situation awareness must be gained at multiple abstraction levels.

Principle 4: Cyber situation awareness has two largely orthogonal viewpoints: The *life-cycle view* contains the proper mechanisms for each phase of the cyber SA process, while the *human cognition view* contains the theories and techniques to integrate human analysts into the overall cyber SA framework (or solution). For automation to facilitate human situation awareness, the human has to model or identify activities of interest for which they wish to maintain awareness.

1.4.2 A Collection of Viewpoints on the Research Agenda

One of Cliff Wang's viewpoints on cyber SA is as follows.

“Over the past two decades, we have witnessed exponential increase in computing power and explosive applications of computing devices. During the same time, information system exploitation and compromises have grown from a novice hobby to the choice of targets by organized crime group and nation/state sponsored adversaries. Unfortunately, our current cyber defense capability is still at an infancy state. Information security practiced daily is art rather

than science. It is quite common for an enterprise to rely its information security on a few knowledgeable, but overwhelmed analysts and a collection of tools that may provide some useful defense against known or past attacks, but are ineffective against new exploits. It is hoped that new investment in cyber situation awareness research will substantially change this picture. New CSA technology will allow analysts to obtain a more complete comprehension of what is going on now, to predicate what might happen next, and to plan and response to ongoing and new cyber attacks effectively. Unlike traditionally machine learning applications which may only interact with physical systems, new CSA technology must deal with sophisticated adversaries with unpredictable behavior patterns. It is crucial that CSA research will take a multi-disciplinary approach and incorporate new advances in areas such as adversarial reasoning, machine learning, and uncertainty management to establish a new paradigm in cyber defense.”

One of John Yen’s viewpoints on cyber SA is as follows.

“The cyber situation awareness overlaps with the situation awareness in the physical battle space. Threats and attacks in the cyber space could affect missions in many different ways. Hence, it is important to integrate the cyber SA with the SA in the mission space. Combining the awareness of the situations in these two spaces enable war fighters to better detect, predict, prevent, and respond to attacks in each space by synthesizing information from both spaces. Rapidly, the cyber space has emerged as the fifth dimension of the battle space, in addition to land, sea, air, and space. A key to integrate the cyberspace SA with the physical SA is to introduce suitable “context” that describes situations across multiple layers. It is also important to allow analyst to maintain situation understanding about the dynamic evolution of multiple situations, so that they can maintain a holistic view in a bigger context, can connect different situations when their relationship emerges, and can predict the evolution of situations, and choose decisions and actions based on their predictive and holistic understanding of both the situation in the cyber space and the mission in the physical space.”

One of Peng Ning’s viewpoints on cyber SA is as follows.

“Several decades of research on intrusion detection and prevention has demonstrated that dealing with intelligent attackers is by no means an easy task. One particular difficulty comes from the uncertainty of information gathered and used for cyber situational awareness. How to reduce such uncertainty is thus critical to the success of this line of research. A promising direction is to take advantage

of the recent advances in trusted computing. For example, we may gain high confidence in the trustworthiness of data gathered for cyber situational awareness by protecting them using a Trusted Platform Module (TPM). Nevertheless, substantial research is necessary to guarantee the successful use of trusted computing technologies to support cyber situational awareness.”

One of Jason Li’s viewpoints on cyber SA is as follows.

“At current stage cyber SA can be extremely overwhelming for human analysts due to the inherent complexity, scalability, and uncertainty issues. To help ease this difficulty, extensive efforts are needed on transformation: from low-level data to meaningful information, from information to actionable knowledge, and from knowledge to trustworthy intelligence. Such bottom-up transformations can be achieved via enhancing the state-of-the-art alert correlation, vulnerability analysis, damage assessment, and machine learning techniques. These efforts can be very useful to help human analysts understand the current situation and project future situations.

On the other hand, human experts may exhibit unique analysis capabilities that surpass the most advanced security analysis software tools, especially with respect to insights and intuitions. While the knowledge possessed by human experts may vary or even conflict with each other, such expertise is extremely valuable and it is necessary to obtain and transfer such expertise into automated cyber SA software tools. This top-down transformation can be achieved via knowledge engineering techniques. Therefore, human-in-the-loop cyber SA means both helping human analysts to better understand as well as using human experts as the design and analysis guide. This will entail the design of some kind of novel human-machine interaction framework.

Lots of challenging problems need to be solved to meet the goal of true human-in-the-loop cyber SA. For example, how to connect the knowledge obtained via the top-down approach (from human) with that obtained via the bottom-up approach (from raw data and information) is an open problem, although the same term “knowledge” is used in both approaches. Without such a connection, it is not possible to realize a holistic (or even consistent) cyber SA solution. To solve this problem, systematic methodologies are needed.

One potential methodology is to treat the cyber enterprise as an organism and design decentralized solutions to handle the challenges related to complexity, scalability, and uncertainty. Essentially, local software agents can be designed to carry out low-level tasks such as monitoring, pattern recognition, reporting, and local reaction. Regional managers (another kind of software agents) can work on a higher-level to coordinate local agents as well as providing a

broader view. Finally, some top-level control center can obtain the global picture and coordinate overall action planning and responses. On a nutshell, the overall cyber SA solution using this methodology can be scalable (local events will be handled locally), effective (views are broadened as needed), and amenable to implementation (using distributed computing paradigm which is the nature of cyber enterprise). Uncertainty management can also be naturally incorporated in this distributed framework using mainstream approaches such as Bayesian networks. Finally, such a framework can also leverage the Trusted Computing approach for uncertainty management.”

One of Xinming Ou’s viewpoints on cyber SA is as follows.

“One thing the technical community can benefit from is talking to security practitioners who have to handle cyber situation awareness manually due to the lack of automated tools. Even though human reasoning is not always accurate or effective, human brains work in a much more flexible manner than what machines can do now and studying how humans react to uncertain and vague information in making quick decisions will be an important first step to automate the process. This can foster a bottom-up theory development process, where we first simulate what a human does in algorithmic ways, then extract from the process the mathematical foundation behind them, and eventually lead to even more accurate and effective automatic situation awareness technologies.

Another closely related question is quantitative analysis. Can we give a metric on the confidence of assertions coming from an SA system? Can we say that with 80% confidence this machine is compromised? Such quantitative metric is not only useful in deciding upon the optimal countermeasure, but also crucial in risk mitigation before an incident happens. According to an article published by IEEE Security & Privacy, 2003, “most organizations call security a top management priority, and chief executive officers rank security as 7.5 out of 10 in importance, but funding levels don’t match the rhetoric. On average, companies spend only 0.047 percent of their revenue on security. Why the disconnect? Simple questions, readily answered in any other business context, are met by information security experts with embarrassed silence. These questions include: Is my security better this year? What am I getting for my security dollars? How do I compare with my peers? Answering such questions requires rigorous security metrics and a risk-management framework in which to compare them.” [2].

One of George Tadda's viewpoints on cyber SA is as follows.

"In order for a human to trust automated decision making, the decision process had to be deterministic. Right now, most human decision makers don't trust a machine to decide if they don't know how the decision was reached or if they wouldn't reach the same decision."

One of Somesh Jha and Matt Fredrikson's viewpoints on cyber SA is as follows.

"Previous work in the area of formal methods and automated reasoning can be brought to bear in simplifying and enhancing many tasks in cyber-SA. For instance, we have previously studied the effectiveness of applying data-driven logic programming to the problem of system and network-level intrusion awareness, and found that most of the work involved in producing representative yet manageable views of intrusion scenarios can be automated. However, applying this technique to yield useful results required substantial research effort; it is clear that before further progress towards a more sophisticated reasoning engine can be made, a groundwork must be laid. The entities and principles essential to tasks in cyber-SA need to be established and formalized, and the reasoning techniques themselves must be modified to suit the particular needs of cyber-SA consumers. For example, provenance is of special concern in cyber-SA, as users may need to "dig deeper" past the results of a reasoning engine to learn more or verify conclusions, but this issue has received little attention in the formal methods literature. Once these issues have been sorted out more, we may benefit from general-purpose reasoning and decision engines for cyber-SA."

One of Sushil Jajodia's viewpoints on cyber SA is as follows.

"The Need for Vulnerability Context in Network Defense: Network defense is inherently difficult. Many internet protocols are insecure, software applications are often buggy, and security measures such as firewalls and intrusion detection systems are complex and error prone. There are large volumes of relevant data, such as detected software vulnerabilities, firewall rules, and intrusion alarms. These data are interrelated, but security tools generally lack the ability to provide the context necessary for effective network defense. What is needed is a capability for "connecting the dots" that shows patterns of attack and corresponding paths of network vulnerability. Such a capability would provide a powerful framework for situational awareness in network defense.

Network defense is labor-intensive, requires specialized knowledge, and is error prone because of the complexity and frequent changes in network configurations and threat behaviors. Furthermore, the correct priorities need to be set for concentrating efforts

to secure a network. Security concerns in a network are highly interdependent, so that susceptibility to an attack often depends on multiple vulnerabilities across the network. Attackers can combine such vulnerabilities to incrementally penetrate a network and compromise critical systems.

However, traditional security tools are generally only point solutions that provide only a small part of the picture. They give few clues about how attackers might exploit combinations of vulnerabilities to advance a network attack. Even for experienced analysts, it can be difficult to combine results from multiple sources to understand vulnerability against sophisticated multistep attacks. In other words, what is lacking is an understanding of the roles of vulnerabilities within the context of overall network defense.”

One of Thomas G. Dietterich’s viewpoints on cyber SA is as follows.

“Existing machine learning approaches to intrusion detection and anomaly detection tend to produce many false alarms. The fundamental reason is that the learning systems have a very narrow view (e.g., sequence of system calls; sequence of packets) that is missing key information (e.g., which vulnerabilities have been patched, changes in local network configuration). An important challenge is to develop learning methods that can integrate and fuse a much broader array of contextual information. Traditional statistical methods break down, because the broader the array of information, the more training examples are required to achieve good performance. We need to develop methods for breaking the learning problem up into modules that can be learned separately and then combined (e.g., [4]).

A second challenge for machine learning in cyber situation awareness is that over time, the relevant features and relationships change as the threats change. Currently, this requires re-engineering the learning system which is costly and requires machine learning expertise. We need machine learning algorithms and user environments that support end users (i.e., system administrators) to that they can diagnose and repair machine learning systems in the field.

A third challenge is to learn from adversarial noise. Machine learning systems typically assume that the input data has random, non-adversarial, measurement noise. An important challenge is to consider cases where malware has the partial ability to delete or modify a subset of the log entries. Can we develop learning methods for learning from adversarial data? One possibility is to first have an abductive “data interpretation” level that maps from the raw logs to the most likely low-level interpretation of events. These, more reliable interpretations then provide a basis for learning.”

One of Peng Liu's viewpoints on cyber SA is as follows.

"Some of the main research issues in the area of cyber SA are as follows. (a) Uncertainty and risk mitigation in cyber situation awareness via such techniques as hypothesis-based (probabilistic) reasoning. (b) Situation (knowledge and semantics) representation and modeling: transforming plain (situation-describing) English to machine-readable models; digitizing human (situation-awareness) intelligence. (c) Automating human analysts' cognitive situation-awareness processes. (d) Situation awareness across multiple abstraction levels. (e) Hypotheses and reasoning against incomplete and imperfect situation information. (f) Gaining better cyber situation awareness through machine learning. (g) Integration of situation perception, comprehension, and projection. (h) Identifying cyber SA measures of performance and effectiveness. (i) Information fusion for cyber situation-awareness. (j) Achieving machine self-awareness (k) Attacker behavior and intent analysis."

1.5 Conclusion

The goal of this article is to clarify the cyber situational awareness problem and to propose a tentative research agenda for solving the cyber SA problem. A set of research issues viewed as important by the authors are also briefly discussed.

Acknowledgements

We would like to thank the Army Research Office for sponsoring the workshop on Cyber Situation Awareness held at George Mason University in March 3-4, 2009.

References

- [1] H. Gardner, *The Mind's New Science: A History of the Cognitive Revolution*, Basic Books, 1987.
- [2] D. Geer Jr., K. S. Hoo, A. Jaquith, "Information security: Why the future belongs to the quants," *IEEE Security & Privacy*, 2003.
- [3] P. Johnson-Laird, *How We Reason*, Oxford University Press, 2006.
- [4] C. Sutton, A. McCallum, "Piecewise Training for Structured Prediction," *Machine Learning*, To appear.
- [5] G. Tadda and et al., "Realizing situation awareness within a cyber environment," In *Multisensor, Multisource Information Fusion: Architectures, Algorithms, and Applications*, B. V.

Dasarathy, eds., *Proceedings of SPIE Vol. 624* (SPIE, Bellingham, WA, 2006) 624204, Kissimmee FL, April 2006.