



Principles for Human-Centered Interaction Design, Part 2: Can Humans and Machines Think Together?

Larry Bunch, Jeffrey M. Bradshaw, Robert R. Hoffman, and Matthew Johnson, Florida Institute for Human and Machine Cognition

Although the word “compute” does not immediately bring a social context to mind, its original meaning is grounded in the interaction of two or more parties. The word comes from the Latin *computare*, from the roots *com* (together) and *putare* (to think, reckon, clear up, or settle). With apologies to Alan Turing, this column will discuss an example of joint human-machine sensemaking for cyber events under the rubric of, “Can humans and machines think together?”

In Turing’s exploration of the question, “Can machines think?” he laid out an experiment in the form of a game.¹ The challenger in the game is given the task of comparing the separate answers of a human and a machine in order to determine which is which. By way of contrast to Turing’s game, the challenge in our example is not in distinguishing between humans and machines. Instead, it explores some initial efforts to blur the line between human and machine thinking—to understand what it might be like someday for humans and machines to be coactively engaged in a form of joint sensemaking so closely and continuously linked and fitted to the humans and machines that it seems as if the parties are a thinking system.

How Is Cyber Sensemaking Different Than Aircraft Flight?

The system we describe in this article is based on similar interaction-design principles as OZ, the human-centered flight display described in Part 1 of the essay.² In both the OZ flight display and the cyber sensemaking display, the goal is to improve the work system’s performance by amplifying and extending human understanding of the current

situation and facilitating anticipation of unfolding events. In both of these examples of interaction design, the goal is to enable appropriate action by portraying events of potentially overwhelming scale and complexity—in essence, fitting a visualization lens to a semantically rich interpretive model of the original data. However, unlike the OZ flight display, the responsibility of our cyber sensemaking display would be not only to support perception and action relating to a fixed interpretive model of phenomena, but also to support the understanding of *changes* to the interpretive model of a potentially fundamental nature as the sensemaking process posits new hypotheses. Simply put, the plasticity of the interpretive model underlying the design of interaction for sensemaking would attempt to match the plasticity of the sensemaking process itself.

Consider this motivating example from the aviation domain. Many cockpits share a similar set of gauges, including an attitude indicator, altimeter, airspeed indicator, compass, and vertical speed indicator. These gauges provide the necessary information for interpreting the flight, but they do not change as the aircraft transitions between maneuvers. Pilots must be trained to alter their scan patterns to effectively interpret the gauges, because the important characteristics change.

One might imagine a better flight display that could adapt on the fly to change its rules of interpretation whenever its flight mode changes. Or, even better, the same display could have features tailored to address all common flight regimes without requiring a scan pattern. An example of this is the OZ flight display, which supports level flight, constant turns, and constant descents through a single bent-wing primitive. Note that in cyberwork,

the need for such adaptation is even greater than in flight, because we are not talking about toggling a display among a few fixed modes or changing a scan pattern, but rather fundamentally reconfiguring the display to reveal the essential properties of a virtually infinite number of potential network attack and defense strategies. Moreover, although it is possible with the OZ flight display for trained pilots to learn to discriminate reliably between the major modes of flight operations, cyberwork is characterized by uncertainty, ambiguity, and deliberate deception by the adversary.

Another way of describing the differences between the OZ flight display and a cyber sensemaking display is in the difficulty in finding the equivalent of a normative flight performance model for computer network analysis. Whereas the pilot's primary task is to fly effectively within the known parameters of a fixed aerodynamic model, the cyber analyst's job is to accurately understand emerging threats against the moving target of a network that is constantly changing in many dimensions. For this reason, what cyber sensemaking requires is not a control device, nor merely an informative picture of the world, but rather a tool for formulation, exploration, and testing of hypotheses about a dynamic situation—essentially the framing and reframing aspects of sensemaking that allow humans and software agents to think together. It follows, then, that the utility of a cyber sensemaking display should be evaluated in terms of its effectiveness in asking and answering a serviceable range of relevant questions for the analyst. Examples might include:

1. Are attacks happening?
2. What might be their origin?
3. What might the attackers be trying to do?

4. What might the attackers do next?
5. Is deception and/or counter-deception involved?
6. How might the attacks affect my mission now and in the future?
7. What options do I have to defend against these attacks?
8. How effective will a given option be against these attacks, what effect will exercising it have on my mission, and how is it likely to affect the future actions of allies and adversaries?
9. Might a defensive action “give me away”?
10. How do I prevent or mitigate the impact of such attacks in the future?

Questions 6 through 10 involve mental projection to the future and deciding. Questions 7 and 9 involve flexible execution.³ The answers to questions 5 through 9 depend on being able to answer questions 1 through 3. These first three questions involve both sensemaking and situation assessment.⁴

With this perspective as background, we introduce the Network Observatory, a human-centered, agent-supported cyber sensemaking system, and highlight how it addresses several of the kinds of questions that analysts must answer. We then give an overview of the visual design principles employed in the Network Observatory. Finally, we address the sensemaking design principles we have explored in the Network Observatory. Specifically, we describe our initial intuitions about what we call “coactive emergence,” the iterative process where useful interpretations of data are continuously developed through the interplay of joint human-machine sensemaking and decision-making activities.

Toward a Human-Centered Cyber Sensemaking System

Despite the attention being given to critical cyberoperations problems,

the ability to keep up with the increasing volume and sophistication of network attacks is lagging. Throwing more computing horsepower at fundamentally limited visualization and analytic approaches will not get us anywhere. Instead, we must rethink the way cyberoperations tools and approaches have been conceived, developed, and deployed. We need a human-centered system for cyber sensemaking. The Network Observatory is an important part of our first attempt at such a system.

Network Observatory Overview

The Network Observatory (or Observatory, for short) is a highly configurable, interactive 4D visualization of network traffic. The Observatory—and the Sol framework of which it is a part⁵—were designed to support several individual and group sensemaking functions, including continuous knowledge discovery across individuals, groups, and software agents.⁶ In addition, Sol was designed to support continuous knowledge preservation by collaborative logging of cases and workflows by analysts and software agents.^{5,7} The framework manages a large logical pool of event data that is shared by many analysts and software agents. All actors can collaboratively explore, filter, and annotate the data within the constraints established by KAOs, a semantic policy framework that governs data use.⁸ The Observatory is a primary way analysts view and manipulate this event data to make sense of it. The Observatory also supports creating and directing the population of software agents that help analyze the massive volumes of high-tempo event data.

Figure 1 shows a snapshot of the Observatory. As with the OZ cockpit display, the Network Observatory relies on colored lines and dots on a black background. These perceptual

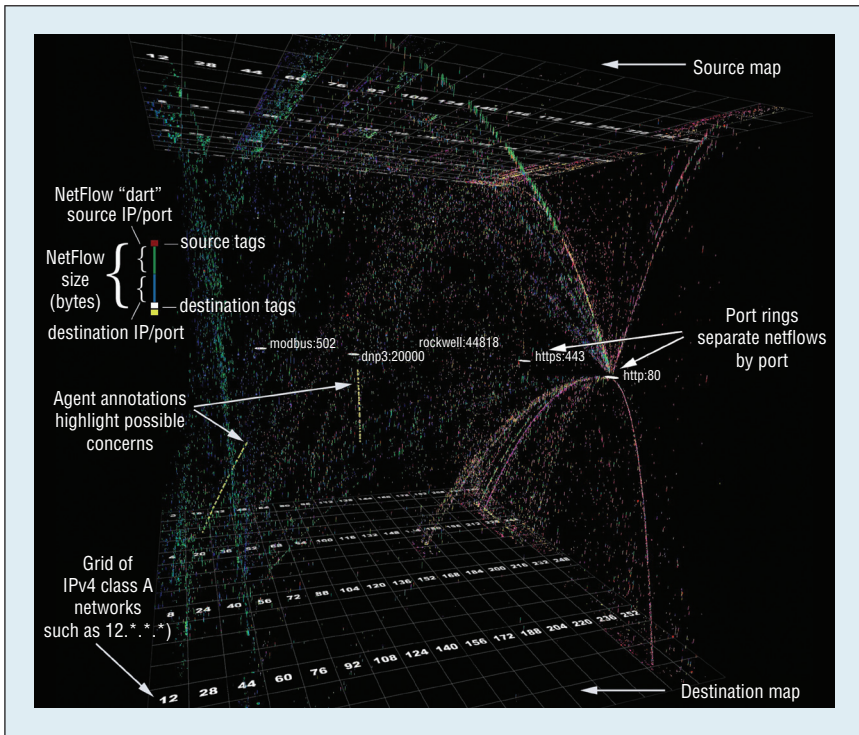


Figure 1. A screenshot of the Network Observatory annotated to illustrate the background context comprised of source and destination IP address maps.

primitives allow for resilience in the face of optical and neurological demodulation and exploit the properties of ambient vision on the basis of the principles discussed previously.²

The input to the Observatory visualization is NetFlow or virtually any other type of record that concerns cyber or physical events happening in time. For example, NetFlow or IPFIX records contain information about a network event time, source and destination addresses, protocols and ports used, and size and rate of the data exchanged.

Planes

The two planes at the top and bottom of the display provide a spatial context for the graphical layout of events. For instance, in Figure 1, the top plane shows a source IP map, and the bottom plane shows a destination IP map. Each of these two planes represents the full IPv4 address space, where each

point on a plane is a unique IP address. Analysts can drill down at any time to see a more detailed projection of the traffic on a plane, displaying, for example, current event records to or from all addresses within a given network.

As alternatives to the IPv4 maps shown, different plane types can be defined and used. For instance, the framework can geolocate the IP addresses and project the source and destination locations as latitude and longitude on a map of the world (see Figure 2). Conceptually based planes—for instance, categorizing events from certain types of groups (such as criminals or nation-state attacks) or economic sectors (such as financial or energy)—can also be defined. The number of planes need not be limited to two; a number of them can be stacked and arranged to suit the topology of the networks involved and the questions of interest.

Darts

Individual event records are depicted in the Observatory both as colored dots on the planes forming source and destination heat maps, and also as individual “darts” that emanate from the top plane and move downward over time. Thus, the darts represent the history of events, with the oldest events at the bottom and the newest events at the top. In the configuration of the display shown in Figures 1 and 2, each dart’s length is proportional to the number of bytes that are being transferred between the source and destination during that network event. Although the figures are not sufficiently zoomed in to reveal detail on the darts, each dart is individually configured with color and various graphical annotations. For example, the top half of each dart typically reflects selected properties of the source plane, whereas the bottom half usually reflects selected properties of the destination plane. The event properties on which dart visual characteristics are based can be easily and dynamically redefined and remapped to represent other properties such as—in the case of computer network data—protocol, duration, and TCP flags.

Rings

The white rings labeled with protocols and port numbers (for example, http:80 and https:443) attract NetFlows that have a matching source or destination port value. This lets them be visually grouped by the ring as they travel downward. The rings are initially placed in sorted order but can be manipulated with a pointing device. For example, an analyst can move the ring to a less-congested area to more easily separate and monitor certain kinds of traffic. Any event property can be used to define rings.

Controls

The Observatory is interactive and dynamically configurable, enabling analysts to manipulate the presentation to answer various questions about the event data. The visualization is a 3D model that analysts can rotate, zoom, and pan to handle data occlusion, observe patterns that may be apparent only from certain perspectives, and reveal patterns of structure from motion. The fourth dimension the Observatory displays is animation over time. Interface controls let the analyst specify a timeframe of interest and the playback rate. Analysts can pause, rewind, and fast forward the display for instant replay in slow or fast motion, enabling them to engage in different kinds of attentive and pre-attentive visual information processing. The time period represented between the top and bottom planes can also be adjusted to show an event history ranging from weeks or days to milliseconds.

Analysts select and filter the events to temporarily isolate phenomena of interest. Selection allows textual metadata for individual events or groups of events to be viewed. Filters determine what events and aspects of events are shown, including any combination of the event properties (for example, source and destination addresses, ports, protocol, countries, regions, cities, and domain names). Analysts can direct existing software agents to associate additional metadata with event records based on the current filters or selection. New classes of software agents that may implement significant changes to the interpretive model can also be constructed programmatically and dropped into the mix at any time. Tagged events will visually pop out when the filters are removed. This ability to tag events and direct the analysts' attention is key to analysts and software agents

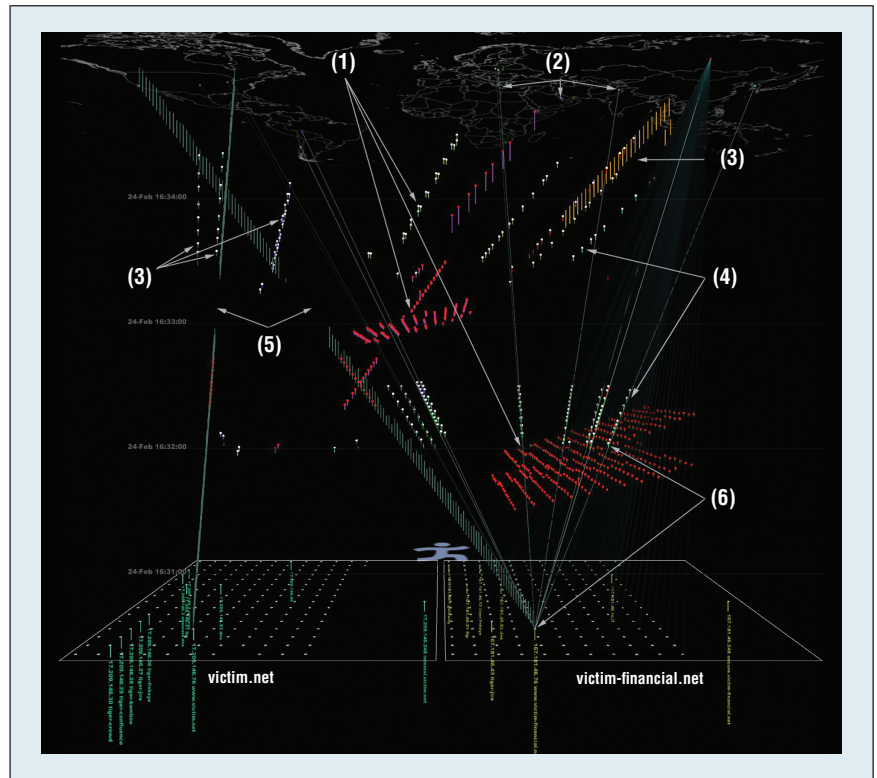


Figure 2. A Network Observatory illustrating answers to common cyber sensemaking questions. (1) Network scanning and subsequent attacks are clearly happening. (2) These attacks originate from several locations worldwide. (3) Attackers are launching distributed denial-of-service (DDoS) attacks on Web services (left), possibly to deceptively mask other attacks (right). (4) Similar attacks are being repeated on both networks. (5) Attacks have disrupted Web services. (6) Defensive options include blocking the set of attacking addresses and relocating the Web services.

working together to make sense of the events. Selections of interest can also be shared among analysts and groups and viewed in other cyber displays in the Sol framework.

Figure 2 illustrates how analysts and software agents use visualization to collaboratively answer the analysts' key questions about the situation. In this instance, one software agent identified the sources of network scanning behavior and tagged all events originating from one of these sources, such that the tails of the corresponding darts are shown in red. Another software agent identified the sources of distributed denial-of-service (DDoS) attacks and tagged all traffic from these IP addresses, such that the tails of all darts are shown in

white. Analysts use the agent-provided highlights to focus on pertinent events and context that can answer important cyber sensemaking questions:

- *Are attacks happening?* Attacks have been launched against both the “victim” and “victim-financial” networks.
- *What is their origin?* The attacks originate from multiple locations worldwide, as shown on the world maps at the top of Figure 2. By way of contrast, the display shows that command-and-control elements of the attack are located in only a few geographic regions.
- *What are the attackers trying to do?* The grid-like patterns shown in red indicate network scanning

to identify which addresses in each network have machines and services that may be vulnerable. The simultaneous sets of frequent requests from multiple sources indicate a DDoS attack that is intended to disrupt the provision of services from the target machine. This type of flooding behavior might also serve as a ploy to dazzle a defender and mask other types of attacks. Indeed, this may be the case with the set of event glyphs shown in orange.

- *What might the attackers do next?* The earlier attack on the victim-financial network that appears at the lower right uses the same sources and similar methods as the subsequent attack, so the analysts hypothesize that events will unfold in much the same way during the attack on the victim network at the lower left.
- *How do the attacks affect my mission now and how might they affect it in the future?* The sets of green glyphs forming nearly solid crossing lines on the left represent the expected continuous interaction between Web services in the two networks. Thus, the gap indicating a lack of activity can mean the attacks have temporarily disrupted the service interactions and negatively impacted the organizational mission.
- *What options do I have to defend against these attacks?* Because the attacks originate from a few sources and there do not appear to be other legitimate requests from the attacking addresses, one option is to block requests from the source IP addresses. Another option is to change the IP addresses of the hosts being attacked. This would require the attacker to recognize the change and retarget the attacks.

In our work, we apply two sorts of human-centered display design

principles: *visual design principles* to support primary perception and actions related to network events, and *sensemaking design principles* to support the process of coactive emergence. We will now outline how these principles are used in the Observatory.

Visual Design Principles Employed in the Observatory

We have adapted the principles of visual design used in the creation of the OZ flight display to our cyber display, with some variation.

First, we note that visualizations, counter to received wisdom, do not have to be immediately intuitive or “natural.” This is particularly true when the primary tasks have no natural analogue, such as when we must recognize and interpret significant events in network data. After all, many tasks are difficult, have a long and steep learning curve, and require significant skill to be acquired through perceptual learning. The measure of the cognitive work should be that once the work system is understood, making sense of the underlying phenomena becomes as effortless as possible. The models that drive interaction can only be designed and created in context of the functional dynamics of the work that they are intended to support.⁹

To harness both focal and ambient vision channels, the dart glyphs in the Observatory, like the graphic elements in the OZ display, use visual-perceptual primitives that are resilient to optical and neurological demodulation. In addition to the design principles for visual primitives discussed in the OZ article, we offer the following, which are well known from other visual perception work.^{10,11}

Use Proportionately Scaled Symbology

This widely applicable design principle can be used in any type of interface

design. Symbols are proportionately scaled when their communicative aspects are not overshadowed by the size, shape, or change in other neighboring symbols. For example, having small glyphs that change color slowly arranged next to large glyphs that change color rapidly will reduce the operator’s ability to detect and respond to the slow color changes. Sizing these appropriately can ensure that the information to be communicated by the symbols is neither muted nor excessively exaggerated. By iterative refinement, we have modulated changes in the size and shape of glyphs representing different network properties to assure that important information is made salient without obfuscating neighboring glyphs.

Set a Holistic Foreground against a Contextual Background

Visualizations that are designed to be processed by the ambient visual channel can exploit movement sensitivity and a large field of view when the display’s visual elements are constructed. This principle is evident in the choice to have darts in the foreground move against the background of static planes.

Create Structure from Motion

This is the phenomenon in which people perceive meaningful objects on the basis of the movement of several elements. One example of this principle is evident in the movement of what appear to be grid-like patterns of scan attacks that are derived from the movement of individual darts (see Figure 2).

Pop Out

Pop out occurs when features of a search target differ significantly from their surroundings, and the target becomes the most salient element.¹² This principle is evident in the way the

important events are tagged by agents and made visually salient on the display.

Chunking

Visual displays of complex data can benefit from chunking conceptually interrelated stimulus units.¹³ One example of this principle is evident in the use of visual rings to group events with similar properties (see Figure 2).

Sensemaking Design Principles Employed in the Network Observatory

A significant design challenge in adapting principles from the OZ flight system to our cyber sensemaking system is the lack of a normative interpretation model across all network situations. We compensated for this lack by using sensemaking design principles that let one or more possible interpretive models be generated and refined by analysts and software agents working together. Through the iterative interplay of joint activity, analysts and agents converge on useful interpretations.

The key features that support coactive emergence in the cyber sensemaking process are design for coactivity, for coevolution of tasks and artifacts, and for second-order emergence.

Design for Coactivity

Our use of the word “coactive” is meant to emphasize the joint, simultaneous, and interdependent nature of collaboration among analysts and automated agents.¹⁴ This is in contrast to more common attempts at implementing human-machine work systems that rely on schemes whereby tasks or subtasks are allocated wholesale to a person or a machine. Simple task-allocation approaches not only introduce a single point of failure for a given task but also hinder others from contributing collaboratively to a teammate’s work.

Interaction design is not simply a matter of putting a human “in the loop,” and it certainly is not a matter of relegating the human to be “on the loop,” as some have recently advocated. It requires understanding where people and machines can each best contribute and knowing how to design a work system to support resilient performance and the kind of interdependence that enables humans and machines to work effectively as teammates.

For example, the Observatory is not a traditional, dedicated single-machine-to-single-person display punctuated sporadically by visual updates and user commands, but rather a common surface—a “mediating representation”—on which any number or combination of analysts and software agents can engage as a team in continuous interaction. In spirit, the Observatory becomes a visual equivalent of the AI blackboard systems of the 1980s: continuously running software agents post interesting results in the context of the Observatory’s single frame of reference at any time while analysts make their contributions on the identical surface asynchronously. The coactive nature of our “visual blackboard” design lets team members make their contributions while maintaining a common understanding of the evolving situation.

That said, the value of coactivity is realized only to the extent that the work system design supports significant coevolution of tasks and artifacts in response to these contributions.

Design for Coevolution of Tasks and Artifacts

The Observatory’s design addresses one of the most important problems of cyber work: namely, how cyber systems must be designed for coevolution of tasks and technology artifacts in order to be adaptive and stay ahead of adversaries that continuously adapt

and escalate their attacks. This never-ending cycle of coevolution between mutually dependent tasks and artifacts is an inevitable challenge to software developers.¹⁵ It means that the capabilities of software will always lag behind current needs, particularly in domains such as cyberwork where the nature of threats is constantly changing. What is new in our approach to cyberwork tools is the idea that the work system’s coevolution can occur rapidly and continuously as tasks and threats develop, rather than requiring potentially long delays as updated versions of software are released to meet new requirements. However, the extent of coevolution is constrained by the degree to which the work system design supports second-order emergence.

Design for Second-Order Emergence

In systems theory, emergence denotes the phenomenon whereby unexpected phenomena or behaviors arise from interactions among the system’s functional components. Emergence in complex systems is often studied through agent-based models. Such models combine the interaction of individual agents possessing individual strategies with deliberately imposed constraints particular to a given environment of interest (see, for example, *Emergence: From Chaos to Order*,¹⁶ p. 117). Although the individual agents are governed by fixed rules, new patterns can arise from their interaction that reveal previously hidden relationships among what once seemed to be disparate particulars. In this fashion, we can view the production of new knowledge as a form of emergence.¹⁷

In the case of the Observatory, the process of emergence is meant to operate at two levels. *First-order emergence* occurs when software agents and human analysts apprehend meaningful patterns in the results obtained by the interpretive models (that is,

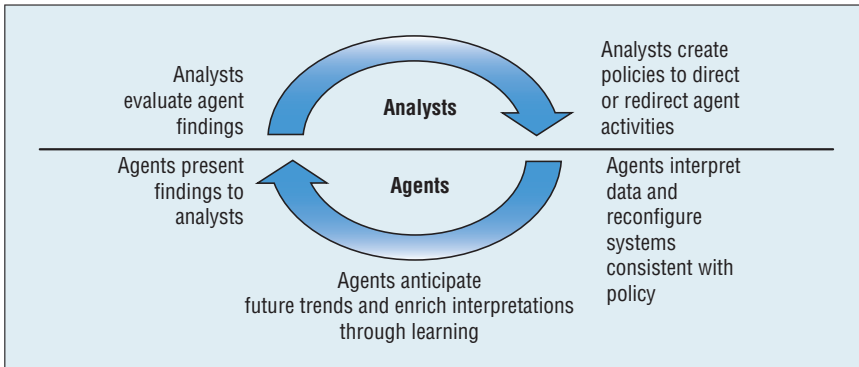


Figure 3. The coactive emergence cycle.

agent policies and software configurations currently in force) for a given dataset (see, for example, *Emergence: From Chaos to Order*,¹⁶ pp. 117–118). This sort of emergence resembles the process of frame elaboration in the data/frame model of sensemaking, where the algorithms specific to particular software agent classes, coupled with policies currently in place, drive the ongoing interpretation of incoming data.

Second-order emergence arises from changes made by software agents and analysts to the interpretive models themselves. This sort of emergence resembles the process of reframing in the data/frame model of sensemaking. Changes to the interpretive model affect the entire system’s behavior, much in the way genetic evolution operates over a population of genes.¹⁸ However, in contrast to those natural systems that affect second-order changes in response to environmental influences that are indifferent to the objectives of the system itself, the analysts and software agents collaborating through the Observatory mutually seek to influence the direction of adaptations so that they converge on shared sensemaking objectives. In other words, analysts seek to change the behavior of software agents in helpful ways—and vice versa. For instance, analysts can add, delete, or change software agent policies on the fly to modify the interpretations or threat responses of agents. In their

turn, software agents present their interpretations to analysts in ways that can influence analyst interpretations or responses.

The ability to create or modify software agents on the fly to support second-order emergence is the capstone feature that enables the kind of plasticity needed to support cyber sensemaking. Figure 3 illustrates how this process can be seen as an experience of mutual teaching and learning among humans and software agents:

- Analysts gather evidence relating to their hypotheses through high-level declarative policies that direct or redirect the ongoing activities of existing or new software agents.
- Within the constraints of policy, software agents interpret and enrich event data. Because of their built-in abilities to reason about and explain their actions using ontologies and to work together at multiple levels of organization, software agent interpretations can be more easily made to match the kinds of abstractions found in human interpretations.
- Software agents aggregate and present their findings to analysts within the context of integrated graphical displays, and analysts interact with those displays to explore and evaluate how software agents’ findings bear on their hypotheses.
- Based on refinements of their hypotheses and questions from these explorations and evaluations, analysts can

redirect software agent activity as appropriate.

The Network Observatory has been running without interruption (except for maintenance updates) on the Florida Institute for Human and Machine Cognition (IHMC) network since mid-2013 to allow visualization and monitoring of Internet traffic interacting with our servers. It has also been delivered to government sponsors for use in various other settings. For instance, a specially configured version of the software was delivered for routine use in ongoing exercises for cyber analysts in training. Highlighting its usefulness for sensemaking of physical rather than cyber events, the National Center for Food Protection and Defense (NCFPD) is partnering with IHMC on tools and strategies for monitoring and protecting critical supply chains.

The Observatory enables network analysts to see and understand Internet traffic in effective new ways. One experienced analyst from a prominent government lab who has used the display to analyze gateway traffic at his network operations center wrote, “The files with DDOS really do ‘pop’ in the display. The images dramatically clarify the abnormality of the DNS during that timeframe; you cannot miss the widths standing out from the rest of the DNS transactions.”

Our experience with human-centered interaction design highlights the benefits of what Don Norman has called “cognitive artifacts.”¹⁹ Such designed objects anticipate their own proper usage, reifying domain knowledge in the intrinsic structure and affordances of the artifact itself, and making this knowledge manifest through constraints in how the artifact may or may not be used. They distrib-

ute actions across time and multiple actors (whether human or machine) and make good use of scientific principles of perception and cognition. What OZ exemplifies, in addition to these general properties, is the path toward increasing the reliability and effectiveness of pilots through exploiting a fixed performance model of flight. The Observatory takes this feature one step further, illuminating how one might eventually harness the power of human-machine teamwork to iteratively generate and refine a series of interpretive models to converge on the most useful understandings of cyber events. Though the current version of the Observatory takes only a few steps toward the ultimate goal of matching the plasticity of the sense-making display to that of the sense-making process, we are certain that the future holds many exciting possibilities for humans and machines to think together. ■

Acknowledgments

The work described in this article is relevant to the following IHMC Intellectual Property: US Patent 8,803,884 for the Event Data Visualization Tool (12 August 2014); US Provisional Patent for the Policy-Governed Software Agent System and Method of Operation.

References


1. A.M. Turing, "Computing Machinery and Intelligence," *Mind*, vol. 59, no. 236, 1950, pp. 433–460.
2. T.C. Eskridge, D. Still, and R.R. Hoffman, "Principles for Human-Centered Interaction Design, Part 1: Performative Systems," *IEEE Intelligent Systems*, July/Aug. 2014, pp. 88–94.
3. G. Klein, "Flexexecution, Part 2: Understanding and Supporting Flexible Execution," *IEEE Intelligent Systems*, Nov./Dec. 2007, pp. 108–112.
4. G. Klein, B. Moon, and R.R. Hoffman, "Making Sense of Sensemaking 2: A Macrocognitive Model," *IEEE Intelligent Systems*, Nov./Dec. 2006, pp. 88–92.
5. J.M. Bradshaw et al., "Sol: An Agent-Based Framework for Cyber Situation Awareness," *Künstliche Intelligenz*, vol. 26, no. 2, 2012, pp. 127–140.
6. L. Bunch et al., "Human-Agent Teamwork in Cyber Operations: Supporting Co-Evolution of Tasks and Artifacts with Luna," *Proc. 10th German Conf. Multiagent System Technologies*, 2012, pp. 53–67.
7. J.M. Bradshaw et al., "Coactive Emergence as a Sensemaking Strategy for Cyber Security Work," *Psychosocial Dynamics of Cyber Security Work*, S.J. Zaccaro et al., eds., Routledge, 2015.
8. J.M. Bradshaw et al., "The KAoS Policy Services Framework," *Proc. 8th Cyber Security and Information Intelligence Research Workshop (CSIIRW 13)*, 2013; www.jeffreybradshaw.net/publications/CSIIRW%20KAoS%20paper-s.pdf.
9. R.R. Hoffman et al., "The Practitioner's Cycles, Part 2: Solving Envisioned World Problems," *IEEE Intelligent Systems*, May/June 2010, pp. 6–11.
10. J.R. Bergen, "Theories of Visual Texture Perception," *Spatial Vision, Vol. 10: Vision and Visual Dysfunction*, D. Regan, ed., CRC Press, 1991, pp. 71–92.
11. B. Julesz, "Figure and Ground Perception in Briefly Presented Isodipole Textures," *Perceptual Organization*, M. Kubovy and J. Pomerantz, eds., 1981, pp. 27–54.
12. S. Kastner, H.C. Nothdurft, and I.N. Pigarev, "Neuronal Correlates of Pop-Out in Cat Striate Cortex," *Vision Research*, vol. 37, no. 4, 1997, pp. 371–376.
13. C.D. Wickens and H.G. Hollands, *Engineering Psychology and Human Performance*, 3rd ed., Prentice-Hall, 2000.
14. M. Johnson et al., "Coactive Design: Designing Support for Interdependence in Joint Activity," *J. Human-Robot Interaction*, vol. 3, no. 1, 2014, pp. 43–69.
15. J.M. Carroll, W.A. Kellogg, and M.B. Rosson, "The Task-Artifact Cycle," *Designing Interaction: Psychology at the Human-Computer Interface*, J.M. Carroll, ed., Cambridge Univ. Press, 1991.
16. J.H. Holland, *Emergence: From Chaos to Order*, Addison-Wesley, 1998.
17. B.R. Gaines, "The Emergence of Knowledge through Modeling and Management Processes in Societies of Adaptive Agents," *Proc. 10th Knowledge Acquisition Workshop*, 1996, pp. 9–14.
18. C. Langton, *Artificial Life: Proceedings of an Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, Addison-Wesley, 1989, p. 90.
19. D.A. Norman, "Cognitive Artifacts," *Designing Interaction: Psychology at the Human-Computer Interface*, J.M. Carroll, ed., Cambridge Univ. Press, 1992, pp. 17–38.

Larry Bunch is a research scientist at the Florida Institute for Human and Machine Cognition. Contact him at lbunch@ihmc.us.

Jeffrey M. Bradshaw is a senior research scientist at the Florida Institute for Human and Machine Cognition. Contact him at jbradshaw@ihmc.us

Robert R. Hoffman is senior research scientist at the Florida Institute for Human and Machine Cognition. Contact him at rhoffman@ihmc.us.

Matthew Johnson is a research scientist at the Florida Institute for Human and Machine Cognition. Contact him at mjohnson@ihmc.us.

 Selected CS articles and columns are also available for free at <http://ComputingNow.computer.org>.