Interacting with Synthetic Teammates in Cyberspace

Scott D. Lathrop[✉]

Soar Technology Inc., 3600 Green Court, Ann Arbor, MI 48105, USA
scott.lathrop@soartech.com

Abstract. This paper explores the interaction of humans and autonomous, intelligent agents working together as teammates in cyberspace operations. Though much research has investigated human-machine teams in domains such as robotics, there is a dearth of research into human-agent dynamics in cyberspace operations Some challenges are similar, such as trust between human and agent. Other challenges, such as representation and interface, are unique to cyberspace given that topological, logical, and temporal relationships are first class constructs with different semantic interpretations from their counterpart visual and spatial representations that are prevalent in physical domains. These challenges arise as the software behaves less like a tool and increasingly becomes more like a synthetic teammate.

Keywords: Human factors · Cyberspace operations · Cybersecurity · Human-agent teaming · Knowledge representations

1 Introduction

There has been a plethora of human factors research on human-machines interactions addressing issues such as trust, communication, and user interfaces. For human-machine interaction, the research typically addresses machines that operate in the physical world, such as robotics platforms or training systems, which emulate a physical world. On the other hand, there is a dearth of research regarding how humans interact with a synthetic teammate for cyberspace operations. In fact, there has been very little research in general regarding teaming in cyberspace [1].

Cyberspace is a relatively new domain where concepts such as teaming are being developed as many of the current capabilities used in the domain are built by and for expert cybersecurity professionals for individual purposes rather than for collections of individuals. For cyberspace operations, where military concepts such as fire and maneuver apply, teaming is an inherent requirement.

Consideration must also be given to the velocity and volume of data that must processed and comprehended in cyberspace operations to drive decision-making. Just the shear amount of data one has to understand to make sense of underlying actions demands more automation. Also at play is a well-documented shortfall of a workforce that can scale to can make sense of this data. These shortcomings point towards more autonomy, transferring some of the tactical decision-making to synthetic teammates

D. Nicholson (ed.), *Advances in Human Factors in Cybersecurity*,
Advances in Intelligent Systems and Computing 593, DOI 10.1007/978-3-319-60585-2_14

that can augment humans by supporting them with data analysis, hypothesis generation, and confirming or denying key attributes or indicators of compromise.

When considering such a human-machine teaming construct, representational issues arise as data in the domain describes topological and logical representations that do not always correlate to visual-spatial representations that are first-class constructs in physical domains. A question also surfaces as to the degree to which such a teaming arrangement requires the personification of the synthetic teammate. This is a fundamental question one must answer because it drives the need for whether natural interaction is required or not (e.g. Siri or some form of augmented reality). For cyberspace operations where there are aspects that are similar in physical domains, such as command and control, maneuver, fires, etc., personification may be an important aspect to the design as it helps support explain-ability and ultimately trustworthiness.

We begin by reviewing what is meant by cyberspace operations. We then apply human-machine teaming concepts to cyberspace operations. Following this discussion, we present some of the representational and interface considerations inherent to building trust for human-machine teaming in cyberspace operations. We then conclude with aspects of our future work.

2 Teaming in Cyberspace Operations

Cyberspace operations are actions conducted in cyberspace—the information environment created when we connect computational nodes together through some physical and logical transmission medium such as Ethernet, fiber, or RF [2]. It includes both cyber-pure or cyber-physical systems, which are systems where these compute nodes receive input from sensors in the physical world or compute solutions that cause an effect on an electro-mechanical actuator in the physical world. Examples of cyber-physical systems include automobiles, electrical power plants, robotics platforms, and military weapon systems.

The cyberspace environment also includes a human element—the cognitive and social factors that enable human interaction through this environment. Direct communication is part of this interaction, but the environment supports a much broader array of behaviors between humans. Examples include social meeting places where the exchange of ideas occur, economic activity and transactions, monitoring and controlling of physical systems, and malicious activity such as stealing information or money, or perhaps worse, physical damage to systems [3].

It follows that from a military perspective, cyberspace operations are pro-active actions to defend these cyber-pure and cyber-physical systems from an active adversary in order to retain freedom of maneuver (defensive) while projecting power to achieve military objectives (offensive). The use of the traditional military functions of intelligence, maneuver, fire support, protection, sustainment, and command and control are important in achieving these objectives as well as the integration of cyberspace actions into physical domains (i.e. land, sea, air, space). The integration of these functions and domains demands teaming at tactical, operational, and strategic levels. This research focuses on tactical-level teaming, specifically between agents and humans.

The types of teams in consideration are the Cyber Mission Forces (CMF), which the U.S. Department of Defense established after the standup of U.S. Cyber Command [4]. The CMF is composed of the teams that are the maneuver elements executing cyberspace actions such as reconnaissance, defense, and attack to achieve both defensive and offensive oriented goals.

Large companies increasingly are applying more military-style processes and techniques to drive their cybersecurity operations, so these observations will apply there also. Security operations centers (SOCs), share some similarities with certain CMF teams, where individuals work collectively to maintain persistent observation of the information flowing in and out of that organization while actively searching for potential compromises. This effectively changes these cybsersecurity teams from a reactive security posture to a proactive defensive posture.

An example of this proactive defense are the procedures, techniques, and tools that these teams employ to support cyber threat hunting [5, 6]. The ability to hunt for adversarial threats across networks of enterprise-scale is becoming an increasingly important part of the CMF and a SOC's tactics, techniques, and procedures (TTPs). The goal of hunting is to identify malicious behavior in an organization's network through indicators of compromise (IOC). IOCs include hash values of malicious software; Internet or domain name addresses; host-based (e.g. logs) or network-based evidence (e.g. netflow data); harvested malware binaries or source code (e.g. implants, command and control malware); and, at a more abstract level, adversary TTPs. Identifying adversary tools and TTPs are the most valuable evidence as they are the costliest for an adversary to change.

Threat hunting uses open-source or classified threat intelligence, that when combined with the organization's asset inventory and known vulnerabilities, facilitates generation of hypotheses as to where potential adversaries may, or already have, compromised systems. These hypotheses focus the team's attention on specific aspects of the data to determine if a compromise has occurred and how it might have happened. This information is then feed into an overall representation of the situation generating new hypotheses and repeating the cycle (Fig. 1a).

As this activity is very much conducive to task-decomposition and requires a somewhat persistent presence, cyber-threat hunting is typically carried out by a small team of individuals composed of different skill sets. Figure 1b lists example work roles (operator, analyst, planner, leader) that might make up such a team. As current state of the art for hunting is resource intensive, especially when considering a network on the order of magnitude of 10–100K nodes, there is a need for automated, or autonomous, tools that offload cognitive tasks performed by operators and analysts, enabling them to hunt more efficiently so that measures such as the number of breaches, dwell time (i.e. how long an adversary in the organization's network), and response time can improve. Some of the activity is conducive to automation such as some of the operators and analysts' functions and thus favorable for human-machine teaming.

The ultimate goal of our research is to reduce the workload requirements for cyberspace operators, analysts, and planners so that they can spend more time comprehending and responding to the broader situation. We have demonstrated progress in building autonomous cognitive agent models to support training [8]. These agents work independently of an overall team, avoiding issues such as trust, communication, and human-machine interfaces, although we have incorporated some of the representations described in Sect. 4.

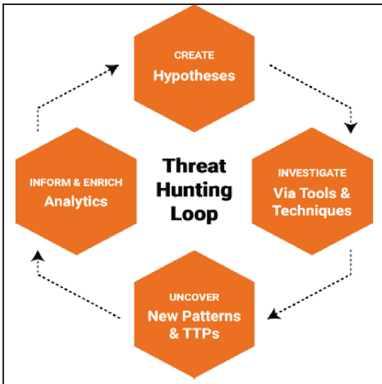 CREATE Hypotheses INFORM & ENRICH Analytics Threat Hunting Loop INVESTIGATE Via Tools & Techniques UNCOVER New Patterns & TTPs	• **Operator** – Collects and processes information from systems to exploit, locate, and/or track targets of interest. Navigates network and conducts immediate forensics analysis. • **Analyst** – Uses collected information (e.g. threat intelligence, IDS alerts, firewall logs, network traffic logs) to identify, analyze, and report threats that are occurring or might occur. Works closely with operators to help determine their next steps. Works closely with planners to identify and validate requirements for analysis. • **Planner** – Develops detailed plans to satisfy operational requirements. Collaborates with other planners and analysts to identify, validate, and levy requirements for collection and analysis. • **Leader** – Synchronizes activity across the team, executing decision-making authority and establishing vision and direction for an organization's operations. Keeps external organizations and leaders informed of planned and ongoing activities.
(a) Cyber threat hunting loop ([6])	(b) Work roles (adapted from [7])

Fig. 1. Cyber threat hunting loop and workroles

3 Human-Machine Teaming for Cyberspace Operations

There are several aspects of human-machine teaming that have been well studied, many revolving around the issue of trust [9]. The factors associated with trust are also important for human-machine teaming in cyberspace/cybersecurity operations where the state of the practice is transitioning from five-year-old soccer, where teammates bunch around the moving ball, to fourteen-year-old soccer where the teammates play their positions. As organizational structures and processes for cyberspace operations mature, the ability to include autonomous agents as part of the teaming structure to facilitate and reduce human workload becomes more feasible and practical.

For example, Abbass et al. [9] illustrate key components for human-machine teaming for autonomous systems (Fig. 2), connecting desired supporting behaviors with what others [10] argue are the baseline functions for teaming: *information exchange*, *communication*, *shared understanding*, and *communication of human intent* (depicted by the four boxes in the bottom left-hand corner of Fig. 2).

Sycara and Lewis [10] point out that the exchange of information, supported by communication, requires bringing to bear all relevant sources of knowledge given the current situational context (e.g. perceptions, past experiences, current internal state). This communication requires internal semantic representations and interfaces that are both general across many functions but also specific to the domain of interest while

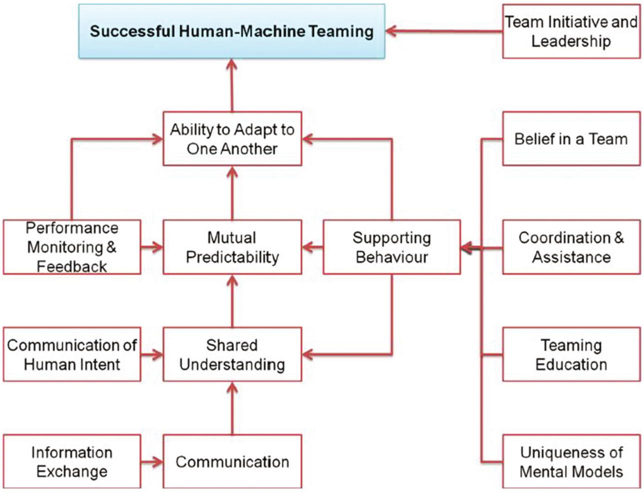
Successful Human-Machine Teaming
Team Initiative and Leadership
Ability to Adapt to One Another
Belief in a Team
Performance Monitoring & Feedback
Mutual Predictability
Supporting Behaviour
Coordination & Assistance
Communication of Human Intent
Shared Understanding
Teaming Education
Information Exchange
Communication
Uniqueness of Mental Models

Fig. 2. Model for successful human-machine teaming [9]

ensuring that noisy data or visual clutter is filtered as much as possible. To support such teaming, they argue architectures should support cognitive processes (e.g. decision-making, planning, reasoning) along with behaviors to support multi-entity relationships (situation assessment and monitoring).

For example, in the cyber threat hunting activity described in the previous section, we have found that an understanding of network topology is important to human operators and analysts [1]. However, the 10–100K nodes and corresponding topological connections that are an organization's network map are not all of equal importance for any particular task. Rather it is typically a few (1–2) compute nodes that are of interest (e.g. a compute node that has been identified with malware) along with a small number of connecting nodes (e.g. 3–10 nodes that are local one-hop connections, organization boundary nodes, and external nodes communicating with the node(s) of interest). There are currently limited ways to communicate a select set of nodes to another application, let alone, to an autonomous agent that may be assisting a hunt activity. Furthermore, the way to internally represent a small subset of a network within an agent's limited capacity for representing knowledge is not well understood.

Also, important in human-machine teaming is the ability to communicate human intent and tasking to an agent [10]. Shared understanding arises between human and machine when the underlying knowledge *representation* supports a two-way dialogue, to include an agent receiving and incorporating a human's intent into its own internal representation and *presenting* the results of its internal processing through a natural interface and in a format that is comprehensible to humans.

An example of where approaches fall short in this regard, are deep learning agents. It is difficult to convey a human's intent to a deep learning architecture in order to direct the system to perform a specific task outside of the task the learned model was trained to recognize. The learned model is trained to classify a particular set of objects. Asking it to classify additional sets of objects or other classes of objects requires retraining the model.

The results of a deep learner's computations and how it inferred its results are not explainable to a human. This is of concern in situations where communication of human intent is important or in adversarial settings where the results may be in question [11]. So, although deep learning performs specific tasks very well (e.g. image recognition), it points to the need for other representations and processes that can support the incorporation of human intentions while explaining recommended actions.

Again, using our cyber threat hunting example, machine learning techniques are now showing up in commercial products to support actions such as anti-malware detection or intrusion detection with limited understanding as to how such systems earn a human's trust and ultimately team with them. There are cases where human analysts have ignored the results of a security product due to a lack of trust in the system's recommendations. For example, during the 2013 exfiltration of credit card data from Target [12], one of the cybersecurity systems warned of the breach but the humans monitoring the system chose to ignore its alerts and turned off its ability to automatically delete malware resulting in 40 million credit card number stolen and a loss of at least 61 million dollars. This is an example where the system was trustworthy, but the humans did not trust its warnings. To help serve as a basis for human-machine teaming in the cyberspace domain, the next section begins to describe some potential avenues to pursue in regards to representation and presentation challenges.

4 Representation and Interfaces for Cyberspace Operations

As stated above, human-machine teaming is centered around trust, with the agent's internal representations and external interface primary factors in supporting information exchange, communication, shared understanding, and communication of human intent. Currently, knowledge representations to support semantics for cyberspace operations is not well understood. Equally important are interfaces that support natural, two-way dialogue and presentation of relevant material, tailored to a human's work role. Within the context of cyber threat hunting, CMF operators and analysts typically prefer command line interfaces with analysts also using web-based tools to support data query, filtering, and prioritization. Visualization of network mapping technology is improving but still rudimentary and displays for situational awareness lacking [1]. Planners and leaders are mostly relegated to presentations and documents to record and convey information—formats that are not conducive to human-machine interaction in a domain such as cyberspace where agility and speed are paramount.

4.1 Knowledge Representations

There has been a plethora of research on knowledge representation to support decision-making, planning, reasoning, and communicating in physical domains. Many symbolic, rule-based systems support knowledge-rich problem spaces where the structure and processing is primarily hand-crafted knowledge based on elicitation from subject matter experts. Such approaches are brittle and do not scale in complex environments where reasoning over concrete representations, such as images or raw malware binaries are necessary. However, these symbolic approaches have been shown to be more explainable and support incorporation of human intent and tasking.

More recently, non-symbolic, deep learning architectures have shown significant progress where the system learns an internal knowledge representation scheme by training it to match its input to a desired output (i.e. supervised learning) or to cluster the input data into groups that are similar (i.e. unsupervised learning). The processing in these systems applies mathematical manipulations by combining affine transformations with continuous functions that are converted to probabilities with a softmax function to classify input. During training of the model, backward processing applies gradient adjustments to weight parameters in order to minimize loss. Despite showing great promise for classification tasks, deep neural networks have limited capacity to reason about their actions and suffer from shortfalls discussed in previously.

Rather than settling on either representation, we have found that support for mixed symbolic and non-symbolic approaches through fixed architectural mechanisms and perceptual interfaces are generalizable across multiple domains [13]. For example, in [14] we demonstrate mixed modality symbolic and non-symbolic representations for visual-spatial domains such as simulations or robotics, where the non-symbolic representations are manifested in the form or mental imagery processing (Table 1).

Amodal, *symbolic* representations are useful for general reasoning and explanations. In physical domains, symbols may denote an object, and visual properties of the object, and qualitative spatial relationships between objects. The first row in Table 1 represents two objects (tree, house) and some qualitative visual and spatial properties (green, left-of).

The *non-symbolic, spatial* representation is also amodal, although perceptual-based in that it is an interpretation of senses asserting the location, orientation, and rough shape of objects in space. Spatial processing is accomplished with sentential, mathematical equations. The second row in Table 1 represents the metric location, orientation, and rough shape of a tree and the house. Direction, distances between objects, size, and rough topology can be inferred implicitly from this information.

In contrast to the symbolic and spatial representation, both of which are sentential structures, space, including empty space, is inherent in the visual depictive representation that is based on the raw, perceived or stored data. Computationally, the depiction is a bitmap where the processing uses either mathematical manipulations (e.g., filters or affine transformations) or specialized processing that takes advantage of the topological structure. Both the symbolic and non-symbolic representations have functional and computational trade-offs that specific tasks often highlight. For example, given appropriate inference rules and the symbolic representation in Table 1, one can infer that the green object (tree) is to the left of the blue object (house). However, one cannot infer the distance between the tree and the house or that the top of the house is shaped like a triangle. One can infer these properties from a symbolic representation only when the relevant property is encoded explicitly or when task knowledge supports the inference. Thus symbolic, top-down processing, when augmented with bottom-up, data driven non-symbolic processing provides wider coverage to multiple classes of problems.

Table 1. Symbolic and non-symbolic representations for visual-spatial processing

Representation	Information	Processing	Example

Symbolic	• Object identities
	• Qualitative spatial and visual properties
		Symbolic manipulation/ productions	object(tree)
			color (tree, green)
			left-of(tree, house)

Spatial (non-symbolic)	• Object labels • Quantitative spatial and visual properties ○ Shape ○ Location, Direction, Orientation ○ Size ○ Topology	Mathematical manipulation	tree: location: <-2,4,0> orientation: 0 shape coordinates: <1,3,1>;<2,8,1>;<1,3,0>… house: location: <9,4,0> orientation: 0 shape coordinates: <8,3,1>;<2,3,1>;<4,3,0>…
Visual (non-symbolic)	• Object labels • Visual properties ○ Shape ○ Texture ○ Empty space • Spatial properties ○ Location, Direction ○ Size ○ Topology	Mathematical and depictive manipulations	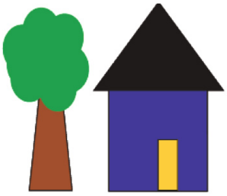

When applying this form of symbolic and non-symbolic representation to cyberspace operations some of the semantic interpretation breaks apart. For example, literal metric distance (*spatial*) between two compute nodes in cyberspace has little meaning, but distance in terms of latency (*temporal)* or the number of hops between two nodes (*logical and topological)* has relevant meaning both in a logical and in a geographic sense (i.e. geographic location may be inferred based on latency and other sources of information). In cyberspace operations topological, logical, and temporal relationships are first class constructs. The semantics of the visual and non-visual properties of compute nodes, their logical bindings (e.g. IP addresses), their software artifacts (e.g. files, processes) and spatial relationships with other nodes must be explicitly represented in any knowledge representation scheme.

Our hypothesis then is that the symbolic and non-symbolic representations used in physical domains apply in the cyberspace domain, but that the semantic interpretation of these features and relationships differ. Table 2 summarizes some of these differences. For example, symbolic objects in the cyberspace domain might be a hardware compute node, its operating system, applications running on that processing node (to include potential malicious applications, and human users interacting with that node such as normal users, system administrators, and remote adversaries. Topological relationships might include connectivity between nodes or qualitative spatial relationships such as the fact that a certain file is stored on a specific node. Such objects represent the physical, logical, and social-cognitive layers of cyberspace [2]. Note that these representations may not necessarily be stored within the agent's memories but rather may exist on an external system with which an agent interacts.

Quantitative spatial relationships may include distance, as previously discussed, or other relationships such as direction, described by a network interface (*logical*) vice a degrees or orientation (*spatial*). Location might imply a physical, medium access layer numeric (i.e. a MAC address), a logical address (e.g. an IP address), a listening port (e.g. a TCP port), a geographic location, or some combination. As discussed above, such quantitative relationships combine symbolic labels with concrete numeric (non-symbolic) information. Finally, visual representations in cyberspace operations

Table 2. Symbolic and non-symbolic representations for cyberspace processing

Representation	**Information**	**Processing**	**Example**

Symbolic	• Object identities
	• Qualitative spatial, visual, and non-visual properties
		Symbolic manipulation/ productions	node(node-1)
			binary(file-1); on (file-1, node-1)
			connected (node-1, node-2)

Spatial (non-symbolic)	• Object labels • Quantitative spatial, non-visual, and visual properties ○ Shape – not defined ○ Location (physical network) ○ Location (logical network) ○ Location (organization) ○ Location (geolocation) ○ Orientation ○ Size (e.g. file size, packet size) ○ Topology ○ Direction (e.g. network interface)	Mathematical manipulation	node-1: location (net): 192.168.1.1 location (geo): <1,3,1> direction: <eth0> direction: <eth1> node-2: location (net): 192.168.1.2 location (geo): <1,2,1> direction: <eth0> file-1: size: 215KB connection: <node-1, eth1> <node-2, eth0> distance: 5ms
Visual (non-symbolic)	• Object labels • Visual properties ○ Shape ○ Texture ○ Empty space • Spatial properties ○ Location ○ Size ○ Topology ○ Direction	Mathematical and depictive manipulations	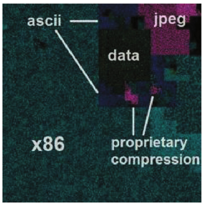 Executable 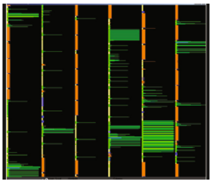 Network packets

that the agent might use to reason over for functional or efficiency gains or to present to the human user for further analysis could include visualizations of binary data such as executables, network packets, or file types within a directory structure [15].

As we have found in physical domains, our hypothesis is that the use of these hybrid approaches can afford efficient processing and provide additional functionality for a certain class of problems with cyberspace. For example, in cyber threat hunting operations, an agent may need to measure distance between communicating nodes by sending a *ping* request and measure latency. Non-symbolic, deep neural networks may provide some of the sensing infrastructure with the symbolic classifications received as perceptual input to the agent. The mix between symbolic and non-symbolic processing then provides support for decision-making and learning over multiple time scales while providing explanation-based representations in the form of symbolic knowledge. These representations are important not only for an agent's own internal processing but also supports interfacing with human teammates.

4.2 Interfaces to Support Two-Way Dialogue

We have found that developing usable human-agent interfaces for teaming requires not only an agent's internal knowledge representation as described above, but also maintenance of a model of the user to help understand their current information needs. This requires understanding users in context, making sense of the user's input, translating that input into a representation that the agent can process and store internally, and then taking the results of the agent's decision-making process across multiple time scales and presenting it to the user in natural ways.

Our research has provided much insight into how this interaction occurs in a human-machine teaming scenario involving unmanned systems [16]. However, we have not applied these lessons for cyberspace agents. Our hypothesis is that many of the techniques we have used for robot-human teaming will also apply here. For example, the interactive devices we have prototyped and employed, enable supervisory control providing the user with the ability to issue high-level commands to the robot with the robot providing feedback to maintain the user's situational awareness. These interactions are through natural interfaces, such as speech, gesture, sketch. To support such interaction, interface devices must have their own level of sophistication with modules to support *dialog management*, *human comprehension model,* and *planning and execution*.

In many cases the combination of multiple modes can help clarify the situation for the agent and build the human's trust that the agent understands the current task. For example, using a prototype interface in Fig. 3, a human cyber hunt analyst may task an agent by circling a node on a network graph and then stating "search for btw.z in the registry keys and identify any anomalous external nodes that *it* is communicating with". The agent interprets the *it* as the node that was circled by the analyst and, after conducting an DNS name lookup on the node may backbrief the analyst by stating, "searching web-1.acme.com for btw.z and calls to suspicious external nodes." As part of the feedback, that agent may visually show a subset of the network graph and highlight the communication path between the compromised internal node and an external command and control server.

We have also explored the use of augmented reality interfaces for human-machine interaction for robotics and Army battle staffs, finding that these interfaces work best when the overlay of control graphics or non-visible entities is important for the operation. Others have investigated the impact on cognitive workload when using augmented reality for SOCs [17]. Their research found that subjects wearing the devices reported reduced cognitive workload, performing the primary cyber-related tasks more efficiently, and responding to ancillary events more successfully. Such approaches may be useful in continuous monitoring situations where hunt operators or analysts need to move away periodically from the display to check on a physical computer.

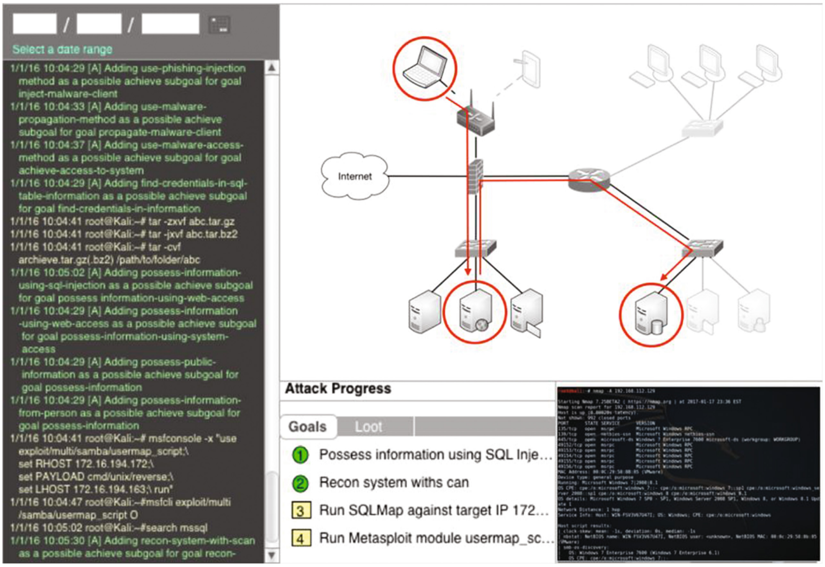

```
1/1/16 10:04:29 [A] Adding use-phishing-injection
  method as a possible achieve subgoal for goal
  inject-malware-client
1/1/16 10:04:33 [A] Adding use-malware-
  propagation-method as a possible achieve
  subgoal for goal propagate-malware-client
1/1/16 10:04:37 [A] Adding use-malware-access-
  method as a possible achieve subgoal for goal
  achieve-access-to-system
1/1/16 10:04:29 [A] Adding find-credentials-in-sql-
  table-information as a possible achieve subgoal
  for goal find-credentials-in-information
1/1/16 10:04:41 root@Kali:~# tar -zxvf abc.tar.gz
1/1/16 10:04:41 root@Kali:~# tar -jxvf abc.tar.bz2
1/1/16 10:04:41 root@Kali:~# tar -cvf
  archieve.tar.gz(.bz2) /path/to/folder/abc
1/1/16 10:05:02 [A] Adding possess-information-
  using-sql-injection as a possible achieve subgoal
  for goal possess information-using-web-access
1/1/16 10:04:29 [A] Adding possess-information
  -using-web-access as a possible achieve subgoal
   for goal possess-information-using-system-
   access
1/1/16 10:04:29 [A] Adding possess-public-
   information as a possible achieve subgoal for
   goal possess-information
1/1/16 10:04:29 [A] Adding possess-information-
  from-person as a possible achieve subgoal for
  goal possess-information
1/1/16 10:04:41 root@Kali:~# msfconsole -x "use
  exploit/multi/samba/usermap_script;\
  set RHOST 172.16.194.172;\
  set PAYLOAD cmd/unix/reverse;\
  set LHOST 172.16.194.163;\ run"
1/1/16 10:04:47 root@Kali:~#msfcli exploit/multi
  /samba/usermap_script O
1/1/16 10:05:02 root@Kali:~#search mssql
1/1/16 10:05:30 [A] Adding recon-system-with-scan
  as a possible achieve subgoal for goal recon-
```

1. Possess information using SQL Inje…
2. Recon system withs can
3. Run SQLMap against target IP 172…
4. Run Metasploit module usermap_sc…

```
root@kali:~# nmap -A 192.168.112.129

Starting Nmap 7.25BETA2 ( https://nmap.org ) at 2017-01-17 23:36 EST
Nmap scan report for 192.168.112.129
Host is up (0.00020s latency).
Not shown: 992 closed ports
PORT      STATE SERVICE      VERSION
135/tcp   open  msrpc        Microsoft Windows RPC
139/tcp   open  netbios-ssn  Microsoft Windows netbios-ssn
445/tcp   open  microsoft-ds Windows 7 Enterprise 7600 microsoft-ds (workgroup: WORKGROUP)
49152/tcp open  msrpc        Microsoft Windows RPC
49153/tcp open  msrpc        Microsoft Windows RPC
49154/tcp open  msrpc        Microsoft Windows RPC
49155/tcp open  msrpc        Microsoft Windows RPC
49156/tcp open  msrpc        Microsoft Windows RPC
MAC Address: 00:0C:29:58:0B:05 (VMware)
Device type: general purpose
Running: Microsoft Windows 7|2008|8.1
OS CPE: cpe:/o:microsoft:windows_7::- cpe:/o:microsoft:windows_7::sp1 cpe:/o:microsoft:windows_se
rver_2008::sp1 cpe:/o:microsoft:windows_8 cpe:/o:microsoft:windows_8.1
OS details: Microsoft Windows 7 SP0 - SP1, Windows Server 2008 SP1, Windows 8, or Windows 8.1 Upd
ate 1
Network Distance: 1 hop
Service Info: Host: WIN-FSV3V67O47I; OS: Windows; CPE: cpe:/o:microsoft:windows

Host script results:
|_clock-skew: mean: -1s, deviation: 0s, median: -1s
|_nbstat: NetBIOS name: WIN-FSV3V67O47I, NetBIOS user: <unknown>, NetBIOS MAC: 00:0c:29:58:0b:05
(VMware)
| smb-os-discovery:
|   OS: Windows 7 Enterprise 7600 (Windows 7 Enterprise 6.1)
|   OS CPE: cpe:/o:microsoft:windows_7::-
```

Fig. 3. Mockup use of sketch, visual, and speech to interface with cyberspace agent

5 Conclusion

This paper explores the interaction of humans with autonomous, intelligent agents working together as teammates in cyberspace operations. The ultimate goal of our research is to reduce workload requirements for cyberspace operators, analysts, and planners so that they can spend more time comprehending and responding to the broader threat.

To support communication, sharing of human intent, and explain ability, symbolic and non-symbolic knowledge representations were explored. Representational challenges unique to cyberspace operations are unlike physical domains where spatial and visual properties and relationships provide concrete interpretations of the world model.

Topological, logical, and temporal relationships are first class constructs in cyberspace, requiring a semantic interpretation of common properties and relationships such as distance, direction, location, and connectedness.

Future work will continue to investigate and prototype these agents with the proposed knowledge representation and natural interaction schemes for cyberspace operations while exploring how malicious adversaries can potential violate these mechanisms in support of their own goals.

References

1. Lathrop, S.D., Trent, S., Hoffman, R.: Applying human factors research towards cyberspace operations: a practitioner's perspective. In: Advances in Human Factors in Cybersecurity, pp. 281–293. Springer (2016)
2. Joint Publication 3–12, Cyberspace Operations (2013)
3. Lee, R.M., Assante, M.J., Conway, T.: Analysis of the cyber attack on the Ukrainian power grid (Traffic Light Protocol (TLP) White). Electrical Information Sharing and Analysis Center (2016)
4. The Department of Defense Cyber Strategy (2015). http://www.defense.gov/Portals/1/features/2015/0415_cyber-strategy/Final_2015_DoD_CYBER_STRATEGY_for_web.pdf
5. Toussain, M.: Home-Field Advantage: Using Indicators of Compromise to Hunt Down the Advanced Persistent Threat. SANS Institute InfoSec Reading Room (2014)
6. Cyber Threat Hunting. https://sqrrl.com/solutions/cyber-threat-hunting/
7. Newhouse, B., Keith, S., Schribner, B., Witte, G.: NIST SP 800-181, NICE Cybersecurity Workforce Framework, National Initiative for Cybersecurity Education (Draft) (2016)
8. Jones, R.M., O'Grady, R., Nicholson, D., Hoffman, R., Bunch, L., Bradshaw, J., Bolton, A.: Modeling and integrating cognitive agents within the emerging cyber domain. In: Interservice/Industry Training, Simulation, and Education Conference (I/ITSEC), vol. 20 (2015)
9. Abbass, H.A., Petraki, E., Merrick, K., Harvey, J., Barlow, M.: Trusted autonomy and cognitive cyber symbiosis: open challenges. Cogn. Comput. **8**(3), 385–408 (2016)
10. Sycara K, Lewis M. Integrating intelligent agents into human teams. In: Salas, E., Fiore, S., (eds.) Team Cognition: Process and Performance at the Inter and Intra-individual Level, Washington. American Psychological Association (2004)
11. Huang, L., Antony, J.D., Nelson, B., Rubinstein, B.I.P., Tygar, J.D.: Adversarial machine learning. In: The 4th ACM Workshop on Artificial Intelligence and Security, Chicago, IL (2011)
12. Riley, M., Elgin, B., Lawrence, D., Matlack, C.: Missed Alarms and 40 Million Stolen Credit Card Numbers: How Target Blew It, Bloomberg.com (2016). http://www.bloomberg.com/news/articles/2014-03-13/target-missed-warnings-in-epic-hack-of-credit-card-data. Accessed 9 Apr 2016
13. Laird, J.E.: The Soar Cognitive Architecture. MIT Press, Cambridge (2012)
14. Lathrop, S.D., Wintermute, S., Laird, J.E.: Exploring the functional advantages of spatial and visual cognition from an architectural perspective. Top. Cogn. Sci. **3**(4), 796–818 (2010)
15. Conti, G.: Security Data Visualization: Graphical Techniques for Network Analysis. No Starch Press, San Francisco (2007)

16. Taylor, G., Purman, B., Schermerhorn, P., Garcia-Sampedro, G., Lanting, M., Quist, M., Kawatsu, C.: Natural interaction for unmanned systems. In: SPIE Defense+Security, pp. 946805–946805. International Society for Optics and Photonics (2015)
17. Beitzel, S., Dykstra, J., Huver, S., Kaplan, M., Loushine, M., Youzwak, J.: Cognitive performance impact of augmented reality for network operations tasks. In: Advances in Human Factors in Cybersecurity, pp. 139–151. Springer (2016)