HUMAN-CENTERED COMPUTING

The Dynamics of Trust in Cyberdomains

Robert R. Hoffman, *Institute for Human and Machine Cognition*
John D. Lee, *University of Wisconsin–Madison*
David D. Woods, *The Ohio State University*
Nigel Shadbolt, *University of Southampton*
Janet Miller, *US Air Force Research Laboratory*
Jeffrey M. Bradshaw, *Institute for Human and Machine Cognition*

All economic, social, and legal interactions are based on assumptions that individuals can verify identities; that they can rely on rules, institutions, and normative practice; and that they can be assured that their private space will remain protected. Monitoring, managing, verifying, auditing, and enforcing these assumptions are difficult online.

How are "cydentities" (cyber-identities) affirmed when they are constantly changing and are sometimes human, sometimes machine, and sometimes virtual? How can individuals and organizations derive persistent and valuable digital identities? How does digital activity change when threats to privacy and trust assumptions arise? How can trust be repaired? Is trust actually stronger once repaired? Is it trust we actually want from our digital systems, or is it accountability and the opportunity for recourse or redress?

Many such empirical and technological questions are raised by cyberstories that have become well known.[1–4] Some of the more widespread cyberattacks are phishing and e-mail scams. Research from Get Safe Online (www.getsafeonline.org), published 9 February 2009, found that 23 percent of Internet users in the United Kingdom had either been a victim of a phishing scam in the preceding 12 months or knew someone who had been affected by the crime. Recent months have witnessed denial-of-service cyberattacks such as those on Facebook and Twitter.[5] A Google search reveals numerous Web sites that provide advice about avoiding scams at Web sites.[6] "Spoofs abound on the Internet. Web sites about hoaxes cover bogus science and technology … and report hoax computer viruses."[7] Thus, we see the advent of "trust evaluation tools" and protocols for a common language for online trust.[8]

An account in the *Washington Post* of the recent cyberattacks on South Korea and the United States points to how difficult it is to protect from attacks and even to identify the perpetrators.[9]

> Cyberspace is a primary medium for the way the Air Force does business, whether it is used for command, control, communications … Almost everything I do is either on the Internet, an Intranet, or some type of network—terrestrial, airborne or spaceborne … Yet, everyone out there knows that hackers can potentially get into my network and slow down or corrupt it or cause me to lose faith in the networks or shut them down completely.[10]

The Lesson and Its Implication

For some decades now, cybersecurity has been a continuous game of catch-up. Clever hackers find some new scheme for malware, the malware is sent out and does its thing, the malware is sniffed out, and then the security people come up with a new software or hardware patch. This "clever game" is not likely to end. Indeed, attempts to improve software and hardware must continue.

But the real lesson is this: *the window of vulnerability never closes*. We must ask, therefore, what else might be done in addition to playing the clever

game? With this question we find issues that are directly pertinent to human-centered computing, both in leveraging what we know about humans, and in leveraging technology to amplify humans.

The cultural, social, psychological, and computational are merging into a Networld. Trust in and through technology will likely mediate the effectiveness of software and hardware in maintaining security. In this essay, we examine some human-centering issues for the Networld, placed primarily for convenience into five categories: antitrust in technology, a consensus on what "trust" is, interpersonal trust versus trust in automation, trusting as a dynamic process, and resilience engineering for the active management of trust.

Antitrust

Previous essays in this department have discussed how so-called intelligent technology triggers frustration and other types of negative affect. Frustration leads users to create kludges and workarounds.[11,12] Antitrust—the skeptical assessment of a technology-reliant work system—is another form of negative affect that can trigger an effort to adapt. In a study of weather forecasters the experts were asked, "Do you trust your technology?" Frequently, they immediately responded, "Never!"[13] This is antitrust: confidence (to the point of certainty) that the technology *will* choke, *will* frustrate, and *will* trigger a need for workarounds. It is no surprise, therefore, that psychologists and human-factors engineers have devoted considerable effort to understanding the circumstances of automation underuse, misuse, and even abuse.[14,15]

This suggests a combinatoric of justified versus unjustified trust and justified versus unjustified distrust. The calibration of trust in automation involves finding the sweet shifts within this constantly morphing space.[16] For instance, novice users sometimes assume that computers are infallible. While they can adapt to the fallibilities of a human teacher, they are stymied when the computer gives what seems to be wrong feedback. People sometimes need help to understand how and when to shift from unjustified trust closer to justified distrust. On the other hand, experienced domain practitioners can become jaded, and are more likely to benefit from information that helps them calibrate their trust by shifting them from unjustified distrust to justified trust.

> This is antitrust: confidence (to the point of certainty) that the technology *will* choke, *will* frustrate, and *will* trigger a need for workarounds.

Trust calibration is captured in one of the laws of macrocognitive work systems,[17,18] which we dub Mr. Weasley's Law, after that fictional character's admonition of his wizarding daughter, "Never trust anything that can think for itself, if you can't see where it keeps its brain."[19] Formally stated, the law is this:

> Mr. Weasley's Law: *Workers in macrocognitive work systems develop unjustified trust and unjustified mistrust in their technology and the work system as a whole when the factors governing the technology's activity are not visible.*

What the cyberresearch community needs is a set of powerful principles to guide people in reaching for desired states on such polarities as antitrust (unjustified mistrust) versus skeptical trust (justified trust); and contingent trust (conditional trust) versus unconditional trust (faith).

We also need to try to define some terms, as impossible as that often is.

A Consensus on What "Trust" Is

The concept of trust has been a topic of analysis in many disciplines, including philosophy (especially ethics), sociology, management science, and psychology. As Kieron O'Hara has noted, any comprehensive account of trust would have to "plunder many sources; the philosophy of Socrates and Aristotle, Hobbes and Kant; the sociology of Durkheim, Weber and Putnam; literature; economics; scientific methodology; the most ancient of history and the most current of current affairs."[20]

Nevertheless, we find it interesting that, of all the fuzzy and abstract concepts inhabiting terra cognita (and ripe for debate in human-centered computing), there actually seems to be a consensus on how to define the multifaceted concept of trust.[16,21,22] Trust has aspects of

- an attitude (of the trustor about the trustee),
- an attribution (that the trustee is trustworthy),
- an expectation (about the trustee's future behavior),
- a feeling or belief (faith in the trustee, or a feeling that the trustee is benevolent, or a feeling that the trustee is directable),
- an intention (of the trustee to act in the trustor's interests), and
- a trait (some people are trusting and more able to trust appropriately).

Table 1. Some of the possible trust relationships.

Type of trust	Beliefs for interpersonal trust	Beliefs for trust in automation	Trustor certainty that the trustee (or automation) will carry out the directives
Unconditional trust (faith)	The trustor takes the trustee's assertions as true.	The trustor is confident that the automation is directable. The trustor takes the automation's actions as correct.	Certain.
Skeptical trust	The trustor takes some of the trustee's assertions as possibly true.	The trustor takes some of the automation's actions as possibly correct.	Somewhat certain.
Circumscribed trust	The trustor takes the trustee's assertions as true, for the time being, or with respect to a certain class of activities.	The trustor takes the automation's actions as correct for the time being, or correct with respect to a class of functions.	Certain for the time being, or certain with respect to certain activities.
Contingent trust	The trustor takes the trustee's assertions as true, depending on the circumstances.	The trustor takes the automation's actions as correct, depending on the circumstances.	Certain depending on the circumstances.
Antitrust	The trustor takes all of the trustee's assertions as false and potentially misleading.	The trustor might take some of the automation's actions as possibly correct, but also anticipates some to be wrong.	The trustor is certain that the technology will choke, will frustrate, and will trigger a need for workarounds.
Swift trust	The trustor has to take the trustee's assertions as true, because of urgent circumstances, often on the basis of trustee authority or position.	The trustor has to take the automation's actions as correct, because of urgent circumstances.	Somewhat certain.
Swift antitrust	The trustor takes the trustee's assertions as false and potentially misleading, because of emerging events or circumstances that reveal the trustee's genuine intentions.	The trustor takes any of the automation's actions as possibly wrong.	The trustor is certain that the technology will provide misleading information.

There are many perspectives on how trust plays a role in human relations,[21,22] in human-machine interaction,[16] and in decentralized sensing and networked systems.[23,24] Analyses generally converge on the concept that *trust reflects an assessment of the trustee's capabilities and competencies to respond to uncertain situations to meet common goals.* The outcome of the assessment is not always a single trust/do-not-trust judgment, but rather an assessment of what particular roles and tasks the trustee can be counted upon to accomplish successfully. The assessment is contingent; it is dynamic and evolves (or degrades) as events occur in the world and information and outcomes feed back to influence the trustor-trustee relation.

Trust can be thought of as a family of relations, certainly not a single relation. For instance, trust can be *about* different things: beliefs (information, data, knowledge), resources (such as valuables), or actions. Trust can be directional (the trustor trusts the trustee) or reciprocal (each party is both trustor and trustee). Trust can be general or contingent. Table 1 expresses some of the many relations.

There is also a consensus on the danger of anthropomorphism in generalizing ideas or research findings about interpersonal trust to the domain of trust in automation.[14,16]

Interpersonal Trust versus Trust in Automation

Interpersonal trust is both fundamentally and subtly different from trust in automation. Trust between people typically involves expectations of intent and reciprocity. But while such basic aspects of social relationships might seem irrelevant to human-automation relationships, people often behave as if technology is a social actor.[25] A direct comparison of trust in humans to trust in automation showed the dynamics of trust to be qualitatively similar.[26] For instance, an aspect of interpersonal trust is directability—that is, that the trustor can direct the trustee. Directability also contributes to human-computer relations.[27]

The feeling that the automation is a partner or a personality is sometimes enhanced by graphic personas. These can promote unjustified trust, and can be an undercurrent in automation abuse (that is, wanting an uncooperative machine to feel pain). The starkest differences between interpersonal trust and trust in automation have to do with the many and powerful effects of menus, graphic objects, and the like on automation misuse and underuse—effects that have no direct analog in interpersonal trust.[28]

How trust (in either a human or a technological agent) develops depends on context and experience. Early evidence that a trustee might be untrustworthy (or that automation might have a high false-alarm rate) can subsequently make it hard for the trustor to develop trust in the trustee (or in the automation). On the other hand, if the trustee (or the automation) provides explanations of the untrustworthy behavior (or why the automation makes false alarms), the effect of fallible performance can be mitigated.

These empirical findings link trust concepts to two additional laws of macrocognitive work systems. First, people are explanation generators. If they are not given sufficient information to make satisfying explanations, they will make explanations anyway.

> The Cognitive Vacuum Law: *Workers develop mental models of the macrocognitive work system, including the technology.*

Thus, in most cases, a worker's mental model of the automation is likely incomplete and not entirely veridical with respect to the designer's intent. The potential for error is increased when the designer's intent actually mismatches the worker's goals and needs. This latter situation is often brought about by the reliance on a designer-centered design approach instead of human-centered design approach. Operators will develop a level of trust in the automation, but it can be over- or under-trust if the designers do not fill the vacuum to make the automation trustable. This leads us to another law:

> Law of Surrogate Systems: *Macrocognitive work systems embody the stances, agendas, and goals of the designers.*

Across all the research on trust in automation, a finding that percolates up consistently is that information allowing the user to understand what the automation does, why it does what it does, and what the designer's intent is significantly promotes trust in the automation. Such information promotes a better understanding of the automation and leads to an appropriate level of trust, which in turn mediates between the feeling that the automation is trustworthy and the actual intention to go ahead and rely on the automation.[14,16,29,30] Without this information, trust or distrust will still

---

Repaired trust relationships might be stronger than relationships in which a breakdown of trust has not occurred.

---

develop, but it will likely be disconnected from the automation's actual capability.

A third perspective on trust is to think of trust using a verb rather than a noun form. What Table 1 points to is not just that trust is a family of relations, but that any given trustor–trustee relation is a dynamic thing.

Trusting

Both interpersonal trust and trust in automation take time to develop. Early experiences contribute to understanding, which eventually can become a more stable relation of faith.[16] On the other hand, sometimes there can be swift trust, where a trustor automatically trusts a trustee on the

basis of authority, confession, profession, or even exigency (see Table 1). Novice belief in the infallibility of computers is an instance of swift (and unjustified) trust. Swift trust can be prominent early in a relationship, with contingent trust developing over time, as people experience the automation in different circumstances.

Trust, whether human–human or human–machine, is dynamic, though it can be temporarily stable. Trust is always context-dependent, though it can be temporarily invariant. Trust can appear to be insulated, though it is actually contingent. Trust, like common ground, must be maintained and even managed. Likewise, mistrust is dynamic, and it too can be maintained (which is unfortunate) or managed.

The dynamics of trust are complex. Factors that contribute to this complexity include, first, threshold effects relating the level of trust to changes in reliance, and, second, contingencies between reliance and the information that guides the evolution of trust.[31]

Threshold effects are reflected in the tendency to maintain a fixed level of reliance even as the level of trust changes, resulting in a dichotomous pattern of reliance. This pattern of reliance can in turn affect the information an operator has regarding the automation's performance. In some cases, the automation's performance is only perceivable when the person is relying on the automation. Threshold effects and contingent information availability complicate the dynamics of trust, and both can undermine the calibration of trust.

There are also interesting findings on the repair of trust.[32] The role of apology and forgiveness as a means to repair trust have been extended to online transaction systems. This line of inquiry entails the possibility that

repaired trust relationships might be stronger than relationships in which a breakdown of trust has not occurred.

Opportunities and Dangers

Cybersecurity and defense issues, including social deception, misdirection, influence, and manipulation, span all the venues of macrocognitive work in the Networld. How can we analyze networks to detect poorly calibrated trust, or increases and declines in trust? All networks are vulnerable to malicious attack, motivating our earlier comment, that the window of vulnerability never closes.

A direct implication of this is that the vulnerability of macrocognitive work systems to malware must be continuously evaluated and experimented on *within the operational context* where networks are being used, generated, grown, and adapted. How can net work continue when one believes that the network has been corrupted? Indeed, clever adversaries seek to maintain networks' overall integrity so that they can manipulate data or applications and go undetected. For the individual net worker, how do we calibrate trust and distrust in technology when the technology is potentially compromised? For macrocognitive work systems, how do we create work methods and work processes for active management of trustworthy communities and work groups?

We might change our perspective from that of victim to that of influencer, from defense to offense. If you remove the highly connected individual nodes in a network you can do much damage, but you can also do damage by removing nodes that are "weakly" connected to other nodes. To concretize this, consider law enforcement combating organized crime, which by tradition has tight and secure networks. If the law

can inject swift mistrust or antitrust somewhere into an organized crime network, tweaking a weak link, the mistrust might spread or cascade, to the benefit of law enforcement.

All the challenges and questions we have raised point to the human element and the cognitive terrain of decision making, as much as they point to challenges for computer science and intelligent systems.

Trust in Macrocognitive Work Systems

Trust emerges from knowledge about the resilience (or brittleness) of the macrocognitive work system—how work systems composed of humans

> To achieve resilience, the technology and work methods must be created to support directability, responsiveness, reciprocity, and responsibility.

and machines adapt when events challenge their boundary conditions.[33] To achieve resilience, the technology and work methods must be created to support directability, responsiveness, reciprocity, and responsibility.[27]

Workers must appropriately trust those parts of the work system that will respond adaptively to disrupting events, events that alter plans and activities in progress. Coordinating the adaptive responses makes decentralized systems resilient.[34] Thus, trust in work systems can be thought of as confidence in this ability of different

units at different echelons to act resiliently. Low levels of trust among units of a work system could play a critical role in maintaining a resilient system by signaling the need for additional resources or reconfiguration.

In macrocognitive work systems, trust can also be thought of as the expectation of reciprocity from others.[35] The parties involved in joint activity enter into a "basic compact," an agreement (often tacit) to facilitate coordination, work toward shared goals, and prevent coordination breakdowns.[27] In a reciprocal cooperative relationship, one human or machine agent relies on another to reciprocate in the future by taking an action that might give up some benefit in order to make both agents better off than they were at the starting point. Coordination and synchronization in distributed systems require reciprocity; otherwise, distributed systems are brittle and exhibit one or more maladaptive behavior patterns.

Both responsiveness and reciprocity emphasize anticipation, which introduces a forward temporal dimension of trust interactions. Both concepts point out that strategies change with changes in trust. For example, when one risks counting on others but anticipates little reciprocity or responsiveness, the result will be unstable, and the responsible party will shift to more conservative, independent strategies.

All the workers in macrocognitive work systems benefit if they can coordinate their activities, but coordination has costs. The challenge is to sustain the commitment to coordination and counteract any tendency to take advantage of others so as to benefit only one role. Decision making always occurs in the context of an expectation that one might be called to account for decisions. Expectations about what are considered adequate accounts, and the consequences for people whose

accounts are judged inadequate, are critical parts of a cycle of accountability. Breakdowns in responsiveness or reciprocity encourage role retreat and break the basic compact.

Models of resilient work systems emphasize how interdependent activities are co-adapted to each other and to changing short-term and long-term conditions.[17] We need new control architectures to manage these resilient, distributed, multi-echelon, human–automation systems.

Finally, the issue of information accountability is arousing new interest in computer science. As we noted at the start of this essay, it is difficult to guarantee privacy when the technologies for information storage, aggregation, and analysis develop so rapidly. We live in an increasingly open information environment, in an increasingly linked Networld. Perhaps we need our technologies to be accountable such that the use of information is apparent, thus making it possible to determine whether a use is appropriate or legitimate in a particular context.[36]

The laws of macrocognitive work systems, which we have presented in this and previous essays in this department, can be thought of as signposts along the path to resilience. The approach of designing for resilience offers some guidance and foundation for the active management of trust in cyberdomains. □

References

1. K. Raina and A. Harsh, *eCommerce Security: A Beginner's Guide,* McGraw Hill, 2001.
2. B. Ortutay, "Don't Post That: Web Etiquette Evolves," Associated Press, 4 Sept. 2009, www.dailynews.com/news/ci_13273951.
3. C. Rhoads and L. Chao, "Iran's Web Spying Aided by Western Technology," *Wall Street Journal,* 22 June 2009, p. 2, http://online.wsj.com/article/SB124562668777335653.html.
4. B. Stelter and B. Stone, "Web Pries Lid of Iranian Censorship," *New York Times,* 22 June 2009, www.nytimes.com/2009/06/23/world/middleeast/23censor.html.
5. N. Gammeltoft and R. Garner, "Twitter, Facebook Resume Online Services after Cyber Attacks," Bloomberg.com, 7 Aug. 2009, www.bloomberg.com/apps/news?pid=20601102&sid=aEVBXs9eH8Gg.
6. "Finding Information on the Internet: A Guide," Teaching Library Internet Workshops, Univ. of California, Berkeley, 2009, www.lib.berkeley.edu/TeachingLib/Guides/Internet.
7. N. Shadbolt, "A Matter of Trust," *IEEE Intelligent Systems,* vol. 17, no. 1, 2002, pp. 2–3.
8. E. Naone, "Adding Trust to Wikipedia, and Beyond," *Tech. Rev.,* 4 Sept. 2009, www.technologyreview.com/web/23355/?a=f.
9. L.C. Baldor, "US Officials Eye North Korea in Cyber Attack," Associated Press, 8 July 2009, www.usatoday.com/news/washington/2009-07-08-hacking-washington-nkorea_N.htm.
10. J. Weckerlein, "Cyberspace Warfare Remains Serious Business," *Air Force Print News Today,* 2 Mar. 2007, www.af.mil/news/story_print.asp?id=123043232.
11. P. Koopman and R.R. Hoffman, "Work-Arounds, Make-Work, and Kludges," *IEEE Intelligent Systems,* vol. 18, no. 6, pp. 70–75.
12. R.R. Hoffman, M. Marx, and P.A. Hancock, "Metrics, Metrics, Metrics: Negative Hedonicity," *IEEE Intelligent Systems,* vol. 23, no. 2, pp. 69–73.
13. R.R. Hoffman et al., "A Method for Eliciting, Preserving, and Sharing the Knowledge of Forecasters," *Weather and Forecasting,* vol. 21, no. 3, 2006, pp. 416–428.
14. M.T. Dzindolet et al., "The Role of Trust in Automation Reliance," *Int'l J. Human-Computer Studies,* vol. 58, no. 6, 2003, pp. 697–718.
15. R. Parasuraman and V. Riley, "Humans and Automation: Use, Misuse, Disuse, Abuse," *Human Factors,* vol. 39, no. 2, 1997, pp. 230–253.
16. J.D. Lee and K.A. See, "Trust in Automation: Designing for Appropriate Reliance," *Human Factors,* vol. 46, no. 1, 2004, pp. 50–80.
17. D.D. Woods and E. Hollnagel, *Joint Cognitive Systems: Patterns in Cognitive Systems Engineering,* CRC Press, 2006.
18. R.R. Hoffman and D.D. Woods, "Steps Toward a Theory of Complex and Cognitive Systems," *IEEE Intelligent Systems,* vol. 20, no. 1, 2005, pp. 76–79.
19. J.K. Rowling, *Harry Potter and the Chamber of Secrets,* Scholastic Press, 1999, p. 329.
20. K. O'Hara, *Trust: From Socrates to Spin,* Icon Books, 2004, p. 5.
21. R.C. Mayer, J.H. Davis, and F.D. Schoorman, "An Integrative Model of Organizational Trust," *Academy of Management Rev.,* vol. 20, no. 3, 1995, pp. 709–734.
22. J.A. Simpson, "Psychological Foundations of Trust," *Current Directions in Psychological Science,* vol. 16, no. 5, 2007, pp. 264–268.
23. T.D. Huynh, N.R. Jennings, and N.R. Shadbolt, "An Integrated Trust and Reputation Model for Open Multi-Agent Systems," *Autonomous Agents and Multi-Agent Systems,* vol. 13, no. 2, 2006, pp. 119–154.
24. R. Stephens, A. Morison, and D.D. Woods, "Trust, ATR, and Layered Sensing: Models, Metrics, and Directions for Design," report to the US Air Force Research Laboratory, Sensors Directorate, from the Cognitive Systems Engineering Laboratory, Inst. for Ergonomics, Ohio State Univ., 2009.

25. C. Nass and Y. Moon, "Machines and Mindlessness: Social Responses to Computers," *J. Social Issues,* vol. 56, no. 1, 2000, pp. 81–103.

26. S. Lewandowsky, M. Mundy, and G. Tan, "The Dynamics of Trust: Comparing Humans to Automation," *J. Experimental Psychology: Applied,* vol. 6, no. 2, 2000, pp. 104–123.

27. G. Klein et al., "Ten Challenges for Making Automation a 'Team Player' in Joint Human-Agent Activity," *IEEE Intelligent Systems,* vol. 19, no. 6, 2004, pp. 91–95.

28. C.L. Corritore, B. Kracher, and S. Wiedenbeck, "On-Line Trust: Concepts, Evolving Themes, a Model," *Int'l. J. Human-Computer Studies,* vol. 58, no. 6, 2003, pp. 737–758.

29. A.M. Bisantz and Y. Seong, "Assessment of Operator Trust in and Utilization of Automated Decision-Aids under Different Framing Conditions," *Int'l J. Industrial Ergonomics,* vol. 28, no. 2, 2001, pp. 85–97.

30. V. Riley, "Operator Reliance on Automation: Theory and Data," *Automation Theory and Applications,* R. Parasuraman and M. Mouloua, eds., Erlbaum, 1996, pp. 19–35.

31. J. Gao and J.D. Lee, "Extending the Decision Field Theory to Model Operators' Reliance on Automation in Supervisory Control Situations," *IEEE Trans. Systems, Man, and Cybernetics,* vol. 36, no. 5, 2006, pp. 943–959.

32. A. Vasalou, A. Hopfensitz, and J. Pitt, "In Praise of Forgiveness: Ways for Repairing Trust Breakdowns in One-Off Online Interactions," *Int'l J. Human-Computer Studies,* vol. 66, no. 6, 2008, pp. 466–480.

33. E. Hollnagel, D.D. Woods, and N. Leveson, eds., *Resilience Engineering: Concepts and Precepts,* Ashgate, 2006.

34. J. Watts-Perotti and D.D. Woods, "Cooperative Advocacy: A Strategy for Integrating Diverse Perspectives in Anomaly Response," *J. Collaborative Computing,* vol. 18, nos. 2–3, 2009, pp. 175–198.

35. E. Ostrom, "Toward a Behavioral Theory Linking Trust, Reciprocity, and Reputation," *Trust and Reciprocity: Interdisciplinary Lessons from Experimental Research,* E. Ostrom and J. Walker, eds., Russell Sage Foundation, 2003.

36. D. Weitzner et al., "Information Accountability," *Comm. ACM,* vol. 51, no. 6, 2008, pp. 82–87.

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

John D. Lee is a professor in the Department of Industrial and Systems Engineering at the University of Wisconsin–Madison. Contact him at jdlee@engr.wisc.edu.

David D. Woods is a professor in the Cognitive Systems Engineering Laboratory at the Ohio State University. Contact him at woods.2@osu.edu.

Nigel Shadbolt is a professor of artificial intelligence in the School for Electronics and Computer Science at the University of Southampton, UK. Contact him at nrs@ecs.soton.ac.uk.

Janet Miller is senior cognitive systems engineer at the US Air Force Research Laboratory. Contact her at janet.miller3@wpafb.af.mil.

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

---

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