Editors: **Robert R. Hoffman, Jeffrey M. Bradshaw, and Ken Ford,** Florida Institute for Human and Machine Cognition, rhoffman@ihmc.us

Seven Cardinal Virtues of Human-Machine Teamwork: Examples from the DARPA Robotic Challenge

Matthew Johnson, Jeffrey M. Bradshaw, Robert R. Hoffman, Paul J. Feltovich, and David D. Woods

This article plays counterpoint to our previous discussions of the "seven deadly myths" of autonomous systems.[1,2] The seven deadly myths are common design misconceptions to be acknowledged and avoided for the ills they breed. Here, we present seven design principles to be understood and embraced for the virtues they engender.

The cardinal virtues of classical antiquity that were adopted in Christian tradition included justice, prudence, temperance, and fortitude (courage). As we'll show in this essay, in effective human-machine teamwork we can also see virtues at play—namely clarity, humility, resilience, beneficence (helpfulness), cohesiveness, integrity, and thrift.

As we unfold the principles that enable these virtues to emerge, it will become clear that fully integrating them into the design of intelligent systems requires the participation of a broad range of stakeholders who aren't always included in such discussions, including workers, engineers, operators, and strategic visionaries developing research roadmaps. The principles aren't merely for the consumption of specialists in human factors or ergonomics.

We illustrate these principles and their resultant virtues by drawing on lessons learned in the US Defense Advanced Research Projects Agency (DARPA) Robotics Challenge (DRC).

DARPA Robotics Challenge

The primary goal of the DRC is to develop robots capable of assisting humans in responding to natural and man-made disasters. The robots are expected to use standard tools and equipment used by humans working in messy spaces engineered for humans to accomplish their mission—hence, a reliance on humanoid robots (www.darpa.mil/Our_Work/TTO/Programs/DARPA_Robotics_Challenge.aspx).

The first event of the DRC was the Virtual Robotics Challenge (VRC). Twenty-six teams from eight countries qualified for the competition. The VRC was carried out in a virtual environment and involved the remote operation of a simulated Atlas humanoid robot, created by Boston Dynamics. There were three tasks to complete: first, navigating complex terrain that included mud, hills, and debris; second, picking up a hose, attaching it to a spigot, and turning a valve; and third, entering a vehicle, driving on a road with turns and obstacles, and then getting out of the vehicle.

In the second competition event, the top eight teams of the VRC were provided with an actual Atlas robot. Additionally, eight other teams used robots they had purchased or designed themselves. Atlas is a hydraulically powered humanoid robot that weighs 150 kilograms (330 pounds) and stands 1.88 meters (6 feet, 2 inches) tall. The robot has 28 degrees of freedom, with six in each major appendage, one in the neck, and an additional three in the pelvis.

The trials consisted of eight tasks, including driving through an obstacle course, walking over slanted ramps, ascending a ladder, removing debris from a doorway, and attaching a fire hose to a spigot. Each task had to be completed in 30 minutes. The robot operator was required to be out of the line-of-sight of the robot at all times. Additionally, DARPA provided bandwidth constraints by introducing a network shaper that oscillated throughput between "good" (1 Mbps speed and 100 ms delay) and "bad" (100 Kbps speed and 1,000 ms delay) communications.

1541-1672/14/$31.00 © 2014 IEEE
Published by the IEEE Computer Society

In the VRC, IHMC was the top-scoring team. In the DRC Trial, IHMC placed second overall, and first among teams using the Atlas robot. Although this placing was due to several factors (including innovative walking control algorithms and agile software development practices), we're convinced that this success can be attributed to the "virtues" obtained by observance of the principles of effective human-machine teamwork described below. To ground these principles in experience, we refer to events from the DARPA competitions.

Lessons Learned: Principles of Teamwork

Here are the principles of teamwork that enable the seven virtues.

1. Focus on improving mission performance of the work system, not on maximizing autonomous capabilities. It's a natural temptation for technologists to focus their attention on technology. For autonomy enthusiasts, giving in to this temptation often leads to carrying out a full-court press to implement strategies that depend exclusively on autonomous capabilities for their success, even in cases where this goal may prove less reliable or more expensive than alternatives that leverage the potential strengths of human-robot teamwork.[3]

DARPA wanted to encourage as much reliance on autonomy as possible. One of the ways they hoped to accomplish this in the VRC was to impose limits on the amount of communication between operators and the robot—forcing the robot to do as much as possible on its own. Although DARPA dropped the communication limits during the subsequent DRC phase, many teams still succumbed to the temptation to focus effort on minimizing bandwidth to the neglect of maximizing overall mission performance. By way of contrast, a team could be better served by using all of the bandwidth available, so long as it contributed toward mission success.

In asserting this first principle, it's important not to be misunderstood. Increased investment in research and engineering for greater autonomous capabilities is undeniably worthwhile as part of an overall strategy to accomplish important work in an effective fashion. Moreover, there are situations where increased self-sufficiency and self-directedness (the basic dimensions of autonomy—see the related work[4]) are not only desirable but absolutely necessary (for example, distant planetary probes that must operate for long periods without communication from ground control).

However, many kinds of work aren't subject to such limitations. In these cases, it's evident that we should rely on whatever combination of human and machine abilities is most likely to get the job done reliably and cost-effectively. This is particularly true in light of the fact that today's so-called autonomous technology has proven fragile due to its deficiencies in anticipating and recognizing problems that may inhibit mission success[5,6] and in its inability to generate flexible alternatives to address them.[7] These are important characteristics for any work system that expects to accomplish complex work in the real world.

Moreover, for certain aspects of the DRC competition, pursuing flawless and genuinely self-sufficient and self-directed capability were indeed the path to successful task completion. Walking is an example of this. It's arguably impossible to develop a method for the human to assist a bipedal robot in balancing when controlling it remotely with a 500-ms delay. For this reason, we vigorously pursued robust dynamic balancing algorithms. However, the requirement for self-sufficiency and self-directedness in dynamic

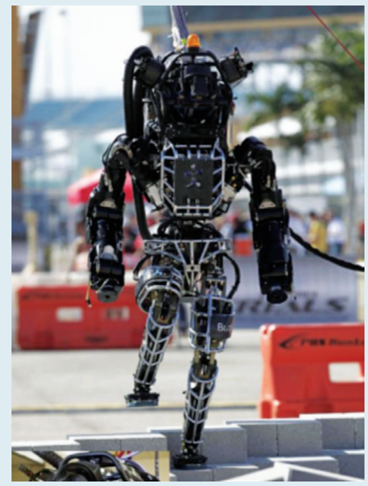

Figure 1. Atlas robot provided by Boston Dynamics. Atlas is a hydraulically powered humanoid robot that weighs 150 kilograms (330 pounds) and stands 1.88 meters (6 feet, 2 inches) tall. The robot has 28 degrees of freedom, with six in each major appendage, one in the neck, and an additional three in the pelvis.

balancing doesn't generalize to other aspects of walking, such as planning and obstacle avoidance.

Autonomous driving provides another example. In the VRC, the task involved driving in a simulated world on a well-marked road with consistent lighting. This was a task where an autonomous solution seemed viable and, for that reason, the Institute for Human and Machine Cognition (IHMC) developed two different driving algorithms to do the job. However, after a pre-competition evaluation, we determined that a teleoperation-style driving method gave us the best chance for mission success.

The Lesson: In competitions such as the DRC, carrying out the mission successfully in an effective, timely, frugal, and resilient manner is what counts—regardless of how much or how little autonomy is employed to do so. This was also true of the Mars Rover mission, the Deep Horizon oil spill ROV repair mission, and many other important occurrences. Focusing on mission performance makes it easy to determine how to allocate scarce resources to potential design and development tasks. We simply ask the question: "Does this approach increase or reduce the chances of mission success?"

The Virtue: Clarity—a focus on mission performance provides a straightforward criterion for decision making about the role autonomy can best play in a given situation.

2. Assess the sweet spot in development effort payoff. In our previous essays about the myths of autonomy, we argued that "complete" autonomy is hard or—perhaps more accurately—impossible in all but the most trivial tasks. In the DRC, we often heard the lament: "If we only had two more weeks, we could have gotten the technology working perfectly." However, experience has taught us that the perfect solution will always be two weeks later than our current deadline. Achieving the last fraction of autonomous capability in a robust and resilient manner is always a challenge, because it requires dealing with every variation in context and circumstance. From a practical perspective, it's a losing proposition because it wastes precious development time while simultaneously increasing system fragility. As evidence, even though the challenges in both the VRC and DRC Trials were specified in great detail prior to the competition, there were no fully autonomous solutions demonstrated for any of the tasks.

Designing for effective human-machine teamwork is frequently the most effective way to increase overall mission performance. People naturally fill many of the dimensions that prove most difficult for machines: perception, judgment, creativity, and so forth. People are amazingly flexible and adaptive. Their participation often allows machines to work out of otherwise fatal situations. While it's true that humans make mistakes, machines do, too.[6] And though machines can sometimes recover from their own mistakes, people are remarkably adept at doing so.

For example, the IHMC DRC team decided not to pursue machine perception. There were several tasks that required high-quality object recognition and localization. Given the limited number of items to recognize (such as the hose, the valve, and obstacles to walking) and the detailed advanced descriptions provided by DARPA, it seemed that autonomous perception routines might be workable. However, after analyzing the problem, we realized the perception tasks were both high risk (since the algorithms were likely to be frail) and low reward (since the task is trivial for a human). Our team saved significant time by not investing in complex perception and planning algorithms and instead devoting our resources to developing capabilities that would allow the human to be an effective teammate with the robot. By strategically ruling out the 100-percent solution (that is, full autonomy or full teleoperation) we could avoid some of the hardest problems.

Even when pursuing autonomy, it may be worth investing some of that time in an approach that will enable human-machine teamwork. As an example, in DRC we pursued an autonomous valve-turning capability. The goal wasn't to complete the entire valve task, but merely turn a single valve that's currently positioned directly in front of the robot and is in reach of the arms. The initial development took about a week, but it didn't work very well. Development and improvements continued through the following weeks with slow and steady progress. As the deadline approached, we began to develop a parallel approach that would rely on human-machine teamwork. This approach was ready in less than a week and ended up performing much more reliably than the autonomous alternative. The teamwork approach wasn't merely teleoperation. It used specific aspects of the autonomous solution, but didn't fully rely on it. Arm sequencing and general positioning were automated, but the most difficult part where the hand needed to contact the valve involved human oversight and adjustment to ensure success.

The Lesson: An 80-percent autonomy solution is always easier than a 100-percent one. Indeed, achieving the last 20 percent of capability is always the most challenging part of the problem. Relying on the human is a good way—sometimes the only way—to cover the remaining 20 percent.

The Virtue: Humility—by admitting that the solution probably won't be the desired fully autonomous robot, you can design better for the system that you will actually have.

3. If you don't plan to fail, you fail to plan. This is the contrapositive of Benjamin Franklin's saying, "If you fail to plan, you plan to fail." In robotics, if you do not plan to fail, then you are failing to plan. Uncertainty and unexpected events are part of reality in robotics, and the solution is to design-in resilience from the beginning.

As an example, to control a robotic arm, a position and orientation in 3D Cartesian space must be specified as a goal. However, the Atlas robot has six degrees of freedom in its joints for each arm. Inverse kinematics provides a mathematical method for figuring out how to control each joint angle so as to achieve the desired goal. Allowing the operator to simply drag a graphic depiction of a virtual arm to specify a desired position and orientation and using inverse kinematics to determine the joint angles is an easy way to position the arms. In fact, this method was used by IHMC as part of virtually all of the arm commands executed during the five hose task runs of the VRC.

Even though the virtual arms were extremely effective, we maintained support for other alternatives, specifically joint-level control. This was particularly important when the solution of the inverse kinetics equations resulted in a singularity, a position that didn't afford recovery by mathematical means. This happened less than one percent of the time—however, if we hadn't developed this alternative we would have failed in three out of the five runs that required it.

Another good example of planning to fail occurred during one of the driving tasks. The car went off the edge of the road during a sharp turn, and needed to perform a backup. We had an "autonomous" method of switching the car into reverse, but it failed because the robot had been jostled as the vehicle left the road and was no longer in the correct position. The support for observability that we had built into the work system involved observability and directability (see Virtue 5, below), which allowed the operator to correctly assess the problem and enabled the engagement of reverse by alternative means. The point to be made here is that our high score was obtained through resilient performance, not flawless performance.

The Lesson: Failure is the result of having only one way to succeed. Resilience, on the other hand, is the result of having many possible ways to recover from failure on the fly.

The Virtue: Resilience—recognizing inevitable problems and having flexible options to address them.[8]

4. Think "combine and succeed," not "divide and conquer." Most attempts at implementation of human-machine work systems rely on schemes whereby entire tasks are allocated to a person or a machine. This approach not only introduces a single point of failure for a given task, but also hinders others from contributing collaboratively to a teammate's task performance in helpful ways.

That said, it's not simply a matter of putting a human "in-the-loop." And it's certainly 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 the kind of interdependence that enables humans and machines to work effectively as teammates.

Approaches to human-machine teamwork typically handle the topic in a manner that's too abstract to be useful for design.[9] To address this need, we developed an approach to human-machine design we call *coactive design*.[10] It relies on a method that identifies areas where the human-machine work system can be made more flexible by providing alternative ways to recognize and handle unexpected situations.

For example, some parts of the VRC hose task required "full autonomy" from the robot (such as dynamic balancing during walking). On the other hand, since the robot had no autonomous perception, the human was required to perform all recognition tasks. However, we found that in several cases the main performer of a task (robot or human) could be assisted by the other supporting team members. Identifying and developing capabilities for support of this kind allows teammates to combine and succeed, especially when the teammates involved (humans and robot) have complementary strengths and weaknesses.

In connection with this principle, it's important to remember Principle 3, particularly focusing on how combining can mitigate some risk. After performing an analysis of interdependence of the human and the robot, it isn't enough merely to select one of the possible options for accomplishing a given task. In fact, our goal was to support as many options as time and money practically afford. Enabling multiple options is what provides the flexibility needed to ensure resilience.

The Lesson: In contrast to the function allocation approach, where the question is, "Which team member can perform the task best?" high-performing teams think through many different ways to perform the same task, and ask "How can each team member assist the other team member performing the task and what is required to support that assistance?"

The Virtue: Helpfulness—identifying the different ways that human and robot teammates can assist each other getting the job done.

5. Design for teamwork in addition to taskwork. What distinguishes joint activity from individual activity? It's the need to support the mutual observability, predictability, and directability (OPD) among team members necessary to support their interdependence. Because OPD is a two-way property, it should shape the design of both the interface of a human operator and the robot's capabilities.

Observability means making pertinent aspects of one's status, as well as one's knowledge of the team, task, and environment observable to others. Since interdependence is about complementary relations, observability also involves the ability to observe and interpret pertinent signals. This is one of the challenges for making machines a team player.[11] Work in the HRI domain[12] lists team knowledge as an important facet of human-agent interaction. Observability plays a role in many teamwork patterns (such as monitoring progress and providing backup behavior).

Predictability means one's actions should be sufficiently observable, reliable, and understandable that others can plan their own interdependent actions accordingly.[11,13] Predictability may involve the use of a priori agreements or it may involve the use of model-based regulation of activities.[14] Predictability is also essential to teamwork patterns such as synchronizing actions and achieving efficiency in team performance.

Team members can make use of what they observe and predict only to the degree that the other team members afford directability. Directability means one's ability to influence and be influenced by other teammates. Directability requires affordances for both explicit commands (such as task allocation or role assignment) and subtler influences. Such accommodations include the ability to productively incorporate "soft" commands in the form of guidance, suggestions, or warnings.[11,12] Teamwork patterns that require directability include requesting assistance and querying for input during decision making.

The need for observability directs designers to focus on questions such as "What information needs to be shared?" "Who needs to share with whom?" and "When is the information relevant?" It's important to remember that it isn't just about what information is shared, but also about what's not shared. Sometimes too much information can be just as big a problem. In each case—whether for observability, predictability, or directability—the goal of a designer isn't to maximize or minimize OPD. It's to attain sufficient OPD to support the interdependence among team members needed for resilient task performance.

An example from the work is the use of scripting. At first, the IHMC team was confident in their ability to automate completely the grasping and lifting portions of the hose task. We used a script recorder to capture successful executions of grasping and lifting movements so they could be played back later to provide full automation. The problem with this approach is that it prevented operators from providing assistance to the robot when necessary. For example, there was no capacity for the robot to verify its own grasp. And by automating the process, we removed the opportunity for the operator to verify that things were going well. Interdependence analysis revealed the brittleness of this approach. To counter this problem, we enabled step-by-step playback of the script with supporting visuals so that the upcoming action would be observable and predictable enough to allow the operator to verify the grasp and abort the task, if necessary. Though this was an improvement, it didn't solve the biggest problem: namely, that the robot couldn't always reliably position its hand for grasping. So we added support to allow the operator not just to verify the upcoming action but also to modify it if necessary, or replay it, or even skip it if desired—that is, to enable directability.

Designing support for OPD led to resilient performance. During the five hose tasks of the VRC, an average of 10 scripts were used per run. Only 50 percent of these were run without intervention. We averaged nine pauses in script behavior to verify performance and seven operator corrections to scripted actions per run. Even with operator intervention, eight of the 50 scripts failed to accomplish their purpose. Due to the flexibility in our system to retry, make adjustments, and use different approaches, we were successful in recovering from all eight failures.

The Lesson: Analysis of interdependence can provide insight into how design decisions, such as automating a task, might impact the overall work system. A solution designed in this manner allows for autonomous behavior, but with appropriate support for interdependence—meaning that the human can participate in the activity in a collaborative manner.

The Virtue: Cohesiveness—possessing the support mechanisms needed for the team to work together as one.

6. Designing for human-machine teamwork goes deeper than the user interface. There's growing sentiment that algorithms and the interface can't be designed separately.[15,16] If a machine algorithm is designed as a black box whose operations aren't observable, predictable, and directable at an appropriate level of granularity and timeliness, even the best operator interface will be to no avail. For this reason, the observability, predictability, and directability requirements of the human should be used not only to guide user interface design, but also the design of robot algorithms and behaviors. In fact, you can't effectively separate these two aspects of design, as we saw in the scripting example discussed previously.

In the design phase of DRC, this approach helped us appreciate and understand the constraints and opportunities as we tried to address eight new tasks. It also was invaluable in the redesign process, as it helped us understand how we needed to modify both our algorithms and our interface to account for changes in moving from simulation to hardware. While the outward result was a unique interface, the actual work involved integration of the algorithms and interface in a way that allowed the operator to make effective use of their perception, judgment, and creativity in concert with the underlying control algorithms.

Enrico Casini and his colleagues[17] ran squarely into this issue in their need to develop an approach to support human involvement in automated data processing pipelines—in this case, large-scale sensor networks that rely on the popular Hadoop MapReduce framework. The aim of the approach was to increase performance and throughput in the automated processing and delivery of data throughout the pipeline while also providing the advantages of human participation at key intervention points along the way. Ultimately, to design for interdependence, you would want to take advantage of specific ways in which all human capabilities—sensing, decision making, and acting—can be exploited to assist machines, and vice versa. These would constitute the universe of opportunities for human intervention.

The key finding of this study that relates to Principle 6 is that involving a human operator in an automated processing pipeline may require significant changes in machine algorithms that go beyond simply being able to make intermediate results available for human inspection. For example, due to the asynchronous nature of human intervention, care must be taken to ensure that once a user-made correction or assertion is introduced, all necessary adjustments and reprocessing are performed. In addition, to make the best use of limited resources and processing capabilities, reprocessing data in light of such corrections must be minimized.

The Lesson: Successful human-machine teamwork requires more than an engineering solution to the machine capabilities. It also requires a holistic view that takes into account the requirements of the entire human-machine work system in its application context. Unless algorithms and user interfaces evolve together, support for OPD is impossible.

The Virtue: Integrity—a marriage of machine functionality and user interface design with a deep commitment for mutual support.

7. Don't simply downsize human involvement; rightsize it. While the frequently touted promise of reduced manning through full autonomy has a natural appeal to organizations intent on reducing their bottom line, reality has a perverse way of shattering many such dreams.[3,18] Neither cost reduction nor improved performance is a guaranteed consequence of reduced manning through so-called autonomous operations. In the long run, the dual objectives of cost reduction and improved performance can only be achieved through greater team effectiveness—not by means of a false hope that the machine will do everything for you on the cheap.

Part of working better is having the means to dynamically rightsize human involvement according to the task and situation at hand. The VRC involved different kinds of walking challenges (such as mud, hills, debris, and flat, open ground). For each kind of walking, the approach and optimal degree of human participation varied. For example, the operator was relatively unburdened in handling walks over flat, open ground because the system could be allowed to work more or less "autonomously." However, when the robot was walking over more difficult terrain, more help from humans was required. Designing for these differences allowed us to seamlessly increase or decrease human involvement in the walking task without requiring a disruptive mode switch. The importance of this type of flexibility was further heightened because the specific nature of each task wasn't known in advance. It's worth emphasizing that in any work situation of sufficient complexity, task requirements will change over time and new kinds of tasks will emerge—making the flexibility and resilience enabled by effective human-machine teamwork an absolute necessity.

The Lesson: Working less comes most easily through working better. Honest assessment of machine and human capabilities and design of options to enable resilient mission performance are the mundane but essential elements of rightsizing human—and machine—involvement.

The Virtue: Thrift—cost reduction and resilient performance over the work system's lifetime through the ability to dynamically rightsize.

The "seven deadly myths" of autonomous systems were so named not only because they cripple the performance of human-machine work systems in their own right, but also because they engender a host of other serious consequences. In similar fashion, what makes the outlined principles of effective human-machine teamwork so powerful is that they allow a host of virtues to emerge and propagate. Chief among them is wisdom—knowing how to work smarter. ■

References

1. J.M. Bradshaw et al., "The Seven Deadly Myths of 'Autonomous Systems'," *IEEE Intelligent Systems*, vol. 28, no. 3, 2013, pp. 54–61; doi:10.1109/MIS.2013.70.
2. R.R. Hoffman, J.K. Hawley, and J.M. Bradshaw, "Myths of Automation, Part 2: Some Very Human Consequences," *IEEE Intelligent Systems*, vol. 29, no. 2, 2014, pp. 82–85.
3. R. Murphy et al., *The Role of Autonomy in DoD Systems*, Defense Science Board Task Force Report, 2012; http://fas.org/irp/agency/dod/dsb/autonomy.pdf.
4. J.M. Bradshaw et al., "Dimensions of Adjustable Autonomy and Mixed-Initiative Interaction," *Agents and Computational Autonomy: Potential, Risks, and Solutions*, LNCS 2969, M. Nickles, M. Rovatsos, and G. Weiss, eds., Springer-Verlag, 2004, pp. 19–39.
5. P.J. Feltovich et al., "Progress Appraisal as a Challenging Element of

Coordination in Human and Machine Joint Activity," *Eng. Societies in the Agents World VIII*, A. Artikis, G.M.P. O'Hare, K. Stathis, and G. Vouros, eds., Springer, 2008, pp. 124–141.
6. D.D. Woods and M. Branlat, "Basic Patterns in How Adaptive Systems Fail," *Resilience Eng. in Practice: A Guidebook*, J. Pariès, E. Hollnagel, J. Wreathall, and D.D. Woods, eds., Ashgate, 2010, pp. 127–144.
7. D.A. Norman, "The 'Problem' of Automation: Inappropriate Feedback and Interaction, Not 'Over-Automation'," D.E. Broadbent, A. Baddeley, and J.T. Reason, eds., *Human Factors in Hazardous Situations*, Oxford Univ. Press, 1990, pp. 585–593.
8. E. Hollnagel, D.D. Woods, and N. Leveson, eds., *Resilience Engineering: Concepts and Precepts*, Ashgate, 2006.
9. R.R. Hoffman and S.V. Deal, "Influencing Versus Informing Design, Part 1: A Gap Analysis," *IEEE Intelligent Systems*, vol. 23, no. 5, 2008, pp. 78–81.
10. M. Johnson et al., "Coactive Design: Designing Support for Interdependence in Joint Activity," *Human Robot-Interaction*, vol. 3, no. 1, 2014; http://dx.doi.org/10.5898/JHRI.3.1.Johnson.
11. 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; doi:http://dx.doi.org/10.1109/MIS.2004.74
12. K. Sycara and G. Sukthankar, *Literature Review of Teamwork Models*, Robotics Institute, Carnegie Mellon Univ., 2006.
13. A. Dragan and S. Srinivasa, "A Policy Blending Formalism for Shared Control," *Int'l J. Robotics Research*, 2013, vol. 32, no. 7, pp. 790–805.
14. P. Feltovich et al., "Toward an Ontology of Regulation: Support for Coordination in Human and Machine Joint Activity." *Eng. Societies for the Agents World VII*, G. O'Hare, M. O'Grady, Alessandro Ricci, and O. Dikenelli, eds., Springer-Verlag, 2007.
15. J.A. Adams et al., "Cognitive Task Analysis for Developing Unmanned Aerial Vehicle Wilderness Search Support," *J. Cognitive Eng. and Decision Making*, vol. 3, no. 1, 2009, pp. 1–26.
16. J.W. Crandall et al., "Validating Human-Robot Interaction Schemes in Multitasking Environments," *Systems, Man and Cybernetics*, vol. 35, no. 4, 2005, pp. 438–449; doi:10.1109/TSMCA.2005.850587.
17. E. Casini et al., "Enhancing Decision-Making by Leveraging Human Intervention in Large-Scale Sensor Networks," presentation at the Military Comm. Conf. (MILCOM 2014), AFCEA Int'l, 2014.
18. J.L. Blackhurst, J.S. Gresham, and M.O. Stone, "The Autonomy Paradox," *Armed Froces J.*, Oct. 2011, pp. 20–40; www.armedforcesjournal.com/2011/10/7604038.

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

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

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

Paul J. Feltovich is a research scientist at the Institute for Human and Machine Cognition. Contact him at pfeltovich@ihmc.us.

David D. Woods is a professor at The Ohio State University in the Institute for Ergonomics. Contact him at woods.2@osu.edu.