Available online at www.sciencedirect.com
SciVerse ScienceDirect
Int. J. Human-Computer Studies 71 (2013) 171–176
International Journal of Human-Computer Studies
www.elsevier.com/locate/ijhcs

From knowledge science to symbiosis science

Jeffrey M. Bradshaw

Florida Institute for Human and Machine Cognition (IHMC) 40 So. Alcaniz Street, Pensacola, FL 32571, USA

Accepted 25 September 2012
Available online 16 October 2012

Abstract

In the mid-1980s, Brian Gaines first developed a model to predict the trajectory of progress in human–computer relationships, including how the knowledge science research programme would naturally transform itself over time into something he called ''symbiosis science.'' In this article, we reflect both on the extraordinary prescience of this model, and the contributions and challenges faced by researchers intent on progressive achievement toward the aspirations it inspires.

Keywords: Symbiosis science; BRETAM; Autonomous systems; Software agents; Personal assistants; Multi-agent systems; Human-agent-robot teamwork; Sensemaking; Wisdom

1. Introduction

Brian Gaines was always thinking decades ahead of the rest of us. His BRETAM diagram brilliantly predicted the trajectory of progress in human–computer relationships, including how the knowledge science research programme would naturally transform itself over time into something he called ''symbiosis science'' (see Gaines, 2013). The term ''symbiosis'' hearkens back to a 1960 article on *man-computer symbiosis* written by J.C.R. Licklider, the first director of the Information Processing Technology Office of the US Advanced Research Projects Agency—now DARPA (Licklider, 1960). In the ultimate form of such symbiosis, human capabilities would be transparently augmented by cognitive prostheses—computational systems that would leverage and extend human intellectual, perceptual, and collaborative capacities, just as a steam shovel is a sort of muscular prosthesis or eyeglasses are a sort of visual prosthesis (Ford et al., 1997; Ford, 1998; Hoffman et al., 2012).

This vision of symbiosis can be contrasted with early efforts in knowledge acquisition where our intelligent systems were somewhat like the disembodied brains shown in low-budget black-and-white science fiction movies: entities that ruled the world while floating in a glass jar tethered by wires.[1] While potentially rich in knowledge models and inferential power, their only direct experience of the world arrived through the impoverished modes of keyboard input and video display output. As a result these intelligent systems were virtually blind and helpless, having little they could realistically *learn about* and even less that they could directly *act upon*. As others in this special issue have observed, the rise of the Internet as the largest repository of knowledge on the planet has given intelligent systems immeasurably richer means to sense, learn, and interact with humans and with the myriad specialized interactive devices, sensors, and services on which people routinely rely.

However, this accumulation of human knowledge in machine interpretable form is only the beginning. Brian Gaines proposed four additional steps that would be necessary to bring the notion of symbiosis science to full fruition:

- the development of goal-directed autonomous knowledge-creating processes;
- the increasing coupling of knowledge processing entities in social networks;

E-mail address: jbradshaw@ihmc.us

[1] For a sampler of such movies, see http://grandoldmovies.wordpress.com/2011/08/20/there-ain't-nothing-like-a-brain-some-favorite-brain-movies/

1071-5819/$ - see front matter © 2012 Jeffrey M. Bradshaw. Published by Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.ijhcs.2012.10.003

- the development of techniques to facilitate the synergy between human and computer knowledge processes;
- the synthesis of both into a unified system.

Let's look at progress on these steps in more detail.

2. The promise and problems of autonomous systems

Addressing the first step of developing "goal-directed autonomous knowledge-creating processes," one of Brian Gaines' students proposed in 1997 a conception of the future Internet as a "cyberorganism" consisting of "distributed intelligent agents," both human and software (Chen, 1997). Subsequently, proponents of the Semantic Web (Berners-Lee et al., 2001) envisioned that such agents would, as Mark Musen expresses it, "comb the Internet and would reason about user goals and how to achieve them" (Musen, 2013). In setting their sights on this goal, agent researchers abandoned the metaphor of the intelligent system qua disembodied brain and adopted the vision of software robots operating in a world of networked computing resources. In this change of metaphor, the research emphasis made an important shift from deliberation to doing, from reasoning to remote action.

Much of the early research on autonomous systems was motivated, not by cyber applications, but by situations in the physical world in which autonomous systems were required to "replace" human participation, thus minimizing the need for considering the human aspects of such solutions. For example, one of the earliest high-consequence applications of sophisticated agent technologies was in NASA's Remote Agent Architecture (RAA), designed to direct the activities of unmanned spacecraft engaged in distant planetary exploration (Muscettola et al., 1998). RAA was expressly designed for use in human-out-of-the-loop situations where response latencies in the transmission of round-trip control sequences from earth would have impaired a spacecraft's ability to respond to urgent problems or to take advantage of unexpected science opportunities.

Sadly, since those early days, most researchers in autonomous systems have continued to pursue their work in a technology-centric fashion, as if full autonomy—complete independence and self-sufficiency of each system—were the holy grail in every situation. Of course, there are problems like deep-space exploration where the goal of minimizing human involvement with autonomous systems can be argued effectively. However, reflection on the nature of human work reveals the shortsightedness of such a singular focus: What could be more troublesome to a group of individuals engaged in dynamic, fast-paced, real-world collaboration than a colleague who is perfectly able to perform tasks alone but lacks the skills required to coordinate his or her activities with those of others? Despite a widespread perception to the contrary, it should be noted that virtually all of the significant deployments of autonomous systems to date—e.g., military UAVs, NASA rovers, oil spill UUVs, and disaster inspection robots—have involved people in important roles, and that such involvement was not merely to make up for the current inadequacy of autonomous capabilities, but also because their jointly coordinated efforts with humans were—or should have been—intrinsically part of the mission planning and operations itself.

In view of the shortcomings of standalone autonomy for complex situations, interest has grown in the topic of "cooperative" or "collaborative" autonomy. Unfortunately, however, this research has a fundamental limitation—namely, that the kind of "collaboration" usually imagined encompasses solely the autonomous systems themselves, regrettably excluding the role of humans as potential collaborators. For example, the United States Department of Defense Unmanned Systems Roadmap stated the goal of pursuing "greater autonomy in order to improve the ability of unmanned systems to operate *independently* [i.e., without need for human intervention], either individually or collaboratively, to execute complex missions in a dynamic environment." Similar briefs have complained of the fact that because UxVs are not truly autonomous, their operation requires substantial input from remote operators. They ask whether additional research in cooperative autonomous behavior—referring to cooperation between the autonomous systems without any human element—could address this "problem."

3. Social machines and human–computer synergy

In contrast to such views, Brian Gaines never saw standalone agent autonomy as the end of the journey. He recognized that just as machine intelligence is hobbled without autonomy, so machine autonomy without sociality is reduced to mere autism. Thus, as a next step, he predicted "the increasing coupling of knowledge processing entities in social networks," a topic deftly summarized by Nigel Shadbolt in his discussion of "social machines" that embody new kinds of emergent and collective large-scale problem-solving by people who are supported by socially-contextualized machines (Shadbolt, 2013). My personal focus, however, has been primarily on the subsequent step in Brian Gaines' model, namely "the development of techniques to facilitate the synergy between human[s] and computer[s]," with the machines acting in the role of differently-abled teammates rather than of sophisticated tools.

Increased synergy between humans and autonomous systems as teammates requires a better understanding of how they become interdependent as part of joint activity. Regrettably, most methodologies for autonomous system design have not been formulated with a sufficient appreciation for the essential role of interdependence in joint human-machine activity (Johnson et al., 2010). While certain approaches to cooperative interaction between humans and machines have become widely known (e.g., dynamic function allocation, supervisory control, adaptive

Table 1
An "un-Fitts" list, © 2002 IEEE.

Machines Are constrained in that:	Need people to:
Sensitivity to context is low and is ontology-limited	Keep them aligned to context
Sensitivity to change is low and recognition of anomaly is ontology-limited	Keep them stable given the variability and change inherent in the world
Adaptability to change is low and is ontology-limited	Repair their ontologies
They are not "aware" of the fact that the model of the world is itself in the world	Keep the model aligned with the world
People	
Are not limited in that:	**Yet they create machines to:**
Sensitivity to context is high and is knowledge- and attention-driven	Help them stay informed of ongoing events
Sensitivity to change is high and is driven by the recognition of anomaly	Help them align and repair their perceptions because they rely on mediated stimuli
Adaptability to change is high and is goal-driven	Effect positive change following situation change
They are aware of the fact that the model of the world is itself in the world	Computationally instantiate their models of the world

automation, and adjustable autonomy), all of them share a common flaw: namely, that they rely on some notion of "levels of autonomy" as a basis for their effectiveness (Johnson et al., 2011).[2] The problem with such approaches is their singular focus on managing human-machine work by varying which tasks are assigned to an agent or robot based on some (usually context-free) assessment of its *independent* capabilities for executing that task. However, decades of studies have shown that successful collaboration in everyday human interaction is largely a matter of managing the context-dependent complexities of *interdependence* among tasks and teammates. To counter the limitations of the well-known Fitts' HABA–MABA (Humans-Are-Better-At; Machines-Are-Better-At) list (Fitts, 1951), which was intended to summarize what humans and machines each do well on their own, Robert Hoffman has summarized the findings of David Woods in an "un-Fitts list" (Table 1), which emphasizes how the competencies of humans and machines can be enhanced through appropriate forms of mutual interaction (Hoffman et al., 2002).

None of this is to say that the pursuit of greater machine autonomy should be abandoned. However, though continuing research to make machines more active, adaptive, and functional is essential, the point of increasing such proficiencies is not merely to make the machines more *independent* during times when unsupervised activity is desirable or necessary (i.e., *autonomy*), but also to make them more capable of sophisticated *interdependent* joint activity with people and other machines when such is required—i.e., *teamwork*. The mention of *joint* activity highlights the need for autonomous systems to support not only fluid orchestration of task handoffs among different people and machines, but also combined participation on shared tasks requiring continuous and close interaction—i.e., *coactivity*.

In contrast to the human-out-of-the-loop autonomy of the RAA, NASA's Portable Satellite Assistant (PSA) prototype is an example of an autonomous system that required close and continuous interaction with people (Bradshaw et al., 2001; Gawdiak et al., 2000). The PSA is a softball-sized flying robot prototype that was designed to operate onboard manned and unmanned spacecraft, collaborating with the limited number of crew members to maintain complex systems, assist with life-critical environmental health monitoring and regulation, coordinate dozens of major simultaneous payload experiments, and perform general housekeeping. Apple's Siri, discussed by Gruber in 2013 (Gruber, 2013), is another successful instance of a collaborative agent that will continue to incorporate an increasing range of autonomous capabilities as it seeks to assist people with their everyday tasks. In addition to such personal assistants, my colleagues and I have been interested in exploring the potential of multi-agent systems (Bradshaw, 1997) in collaborative tasks ranging from coordinated operations of people carrying out semi-structured missions with heterogeneous unmanned robots (Johnson et al., 2008) to sensemaking in cyber operations, where software agents and analysts jointly engage in a process of progressive convergence to identify emerging threats (Bradshaw et al., under review; Bunch et al., 2012). We like to think of the latter as a form of joint human-machine modeling that is consistent with the constructivist thinking of George Kelly and the elaborations of those of us who were inspired by his work, including, among others, John Boose, Guy Boy, Ken Ford, Brian Gaines, and Mildred Shaw (Bradshaw et al., 1993; Ford and Bradshaw, 1993).

4. Teamwork knowledge

Building on the insights of Bill Clancey and Paul Compton, Brian Gaines rightfully pointed out the importance of

[2]In a significant step that has reversed years of precedent in autonomy research, a 2012 US Task Force recommended "that the DoD abandon the use of 'levels of autonomy'" and instead focus their efforts to develop a reference framework that emphasizes the importance of human–computer collaboration (United States Department of Defense — Defense Science Board, 2012), p. 4.

"practical knowledge" (Gaines, 1993) in models of human expertise that drive the task-related behavior of many agent systems. This kind of knowledge is usually represented in the form of heuristics that serve to avoid disasters and to weakly direct goals. Despite the shallow nature of such models, Brian Gaines concluded that practical knowledge "can exhibit robust… strategies" and "remarkable 'adaptivity' in that it is insensitive to major changes in the domain in which it is operating."

But this is only half the story. One of the most important contributions of research on human-agent collaboration is the finding that many aspects of effective joint activity rely not only on the practical knowledge needed to execute a task in isolation, but also on teamwork knowledge in the form of principles, heuristics, and mechanisms for coordinating joint work effectively. Pioneers in agent teamwork research such as Cohen, Levesque, and Tambe concluded early on that teamwork knowledge tends to be more generic and reusable across different applications than taskwork knowledge (Cohen and Levesque, 1991; Tambe et al., 1999; Tambe et al., Kaminka). For this reason, many kinds of teamwork knowledge can be modeled somewhat separately from taskwork knowledge per se.

Teamwork knowledge is typically conceived in terms of formalized social regulations. The idea of building strong social regulation into intelligent systems can be traced at least as far back as the 1940s to the science fiction writings of Isaac Asimov (Asimov, 1942). Shoham and Tennenholtz (Shoham and Tennenholtz, 1992) later introduced the theme of social "laws" into the agent research community. In addition to applying policy constraints to avoid disasters in multi-agent systems, my colleagues and I have attempted to develop reusable policies and mechanisms to guide teamwork behavior (Feltovich et al., 2004; Feltovich et al., 2006; Klein et al., 2004; Sierhuis et al., 2003). Like Web-based knowledge used for human and machine deliberation and like practical knowledge used by software and robotic agents to perform taskwork, it can be convenient to represent significant portions of this teamwork knowledge within ontologies (Bradshaw et al., 2011; Bunch et al., Uszok; Uszok et al., 2008; Uszok et al., 2011).

5. Synthesis into a unified (and wise?) system

In 1985, Brian Gaines produced an early version of a diagram predicting the future of knowledge systems, and showing Wisdom as the pinnacle of that evolution (Fig. 1). Four years later, I was honored to join Brian, along with my mentor and friend John Boose (Boose, 1986; Boose and Bradshaw, 1987; Bradshaw and Boose, 1990; Bradshaw et al., 1991; Shema et al., 1990), in sending out a call for papers for a Workshop on Wisdom-Based Systems that was to take place on June 22–24, 1989 at the Rosario Resort on Orcas Island in Washington state. Here is a paragraph from that call:

Knowledge-based systems are now being applied to a wide spectrum of applications. As new applications in diplomacy, environmental management, jurisprudence, corporate strategy, and others are developed, there is a critical need to understand the limitations and potential of future automated systems. Based on what we have learned from attempts to represent knowledge in computer form, what can we say about the possibility of representing wisdom? Can knowledge-based systems recognize the limits and proper application of their knowledge? Can human values be used to enhance the effectiveness of such systems during judgment and decision making? Will

	EVOLUTION	DEFINITION	BUZZWORDS	OUR VIEW OF COMPUTERS	THEIR VIEW OF US
'56	Data	Uninterpreted numbers	Data processing	Fast, accurate, and reliable	Slow,inaccurate, and unreliable
'64	Information	Structured data	Information technology	Keep track of complex structures	Easily swamped by complexity
'72	Decision	Selecting information	Decision support	Always correct, but rigid	Sometimes wrong, but flexible
'80	Knowledge	Reasoning underlying decision	Knowledge science	First baby thoughts	How do they know so much?
'88	Learning	Acquiring knowledge	Inductive inference	Slow, inaccurate, and unreliable	Slow and social, but necessarily so
'96	Purpose	Directing learning	Autonomous systems	No sense of direction	How do they know what they want?
'04	Power	Achieving purpose	?	Not in my lifetime	Their ultimate goal
'12	Wisdom	Applying power reasonably	?	They don't have it	They don't have it

Fig. 1. The evolution of knowledge systems (Gaines, 1985).

human/computer participatory systems be developed that improve the prospects for wisdom?

For a variety of reasons, perhaps in part due to a recognition of our hubris in making the proposal in the first place, the workshop never materialized. Note however that, according to the figure, 2012 is the year of Wisdom. Has the time now arrived to put out another call for papers?

References

Asimov, I. Runaround .In I, Robot, edited by I. Asimov, 33–51. London, England: Grafton Books. Originally published in Astounding Science Fiction, 1942, pp. 94–103, 1942/1968.

Berners-Lee, T., Hendler, J., Lassila, O., 2001. The semantic web. Scientific American 36, 35–43.

Boose, J.H., 1986. Expertise Transfer for Expert System Design. Elsevier, New York.

Boose, J.H., Bradshaw, J.M., 1987. Expertise transfer and complex problems: using Aquinas as a knowledge-acquisition workbench for knowledge-based systems. International Journal of Man-Machine Studies 26, 3–28.

Bradshaw, J.M., Boose, J.H., 1990. Decision analysis techniques for knowledge acquisition: combining information and preferences using *Aquinas* and *Axotl*. International Journal of Man-Machine Studies 32 (2), 121–186.

Bradshaw, J.M., Covington, S.P., Russo, P.J., Boose, J.H., 1991. Knowledge acquisition techniques for decision analysis using *Axotl* and *Aquinas*. Knowledge Acquisition 3 (1), 49–77.

Bradshaw, J.M., Ford, K.M., Adams-Webber, J.R., Boose, J.H., 1993. Beyond the repertory grid: new approaches to constructivist knowledge acquisition tool development. In: Ford, K.M., Bradshaw, J.M. (Eds.), Knowledge Acquisition as Modeling. John Wiley, New York, pp. 287–333.

Bradshaw, J.M. (Ed.), 1997. Software Agents. The AAAI Press/The MIT Press, Cambridge, MA.

Bradshaw, J.M., Sierhuis, M., Gawdiak, Y., Jeffers, R., Suri, N., and Greaves, M. Teamwork and adjustable autonomy for the Personal Satellite Assistant. Presented at the Workshop on Autonomy, Delegation and Control: Interacting with Autonomous Agents. Seventeenth International Joint Conference on Artificial Inelligence (IJCAI-2001), Seattle, WA, 6 August 2001.

Bradshaw, J.M., Feltovich, P., Johnson, M., 2011. Human-agent interaction. In: Boy, G. (Ed.), Handbook of Human-Machine Interaction. Ashgate, pp. 283–302.

Bradshaw, J.M., Carvalho, M., Bunch, L., Eskridge, T., Feltovich, P.J., Forsythe, C., Hoffman, R.R., Johnson, M., Kidwell, D., Woods, D.D. Coactive emergence as a sensemaking strategy for cyber operations. Manuscript submitted for publication, available at: //www.jeffreym bradshaw.net.

Bunch, L., Bradshaw, J.M., Carvalho, M., Eskridge, T., Feltovich, P.J., Lott, J., and Uszok, A. .Human-agent teamwork in cyber operations: supporting co-evolution of tasks and artifacts with Luna. Presented at the Tenth German Conference on Multiagent System Technologies (MATES 2012) (LNAI 7598), Trier, Germany, October 10–12, 2012, pp. 53–67.

Chen, L.L.-J., 1997. Modeling the Internet as cyberorganism: a living systems framework and investigative methodologies for virtual cooperative interaction. Doctoral Dissertation, University of Calgary.

Cohen, P.R., Levesque, H.J., 1991. Teamwork. SRI International, Menlo Park, CA.

Feltovich, P., Bradshaw, J.M., Jeffers, R., Suri, N., Uszok, A., 2004. Social order and adaptability in animal and human cultures as an analogue for agent communities: toward a policy-based approach. Engineering Societies in the Agents World IV. LNAI 3071. Lecture Notes in Computer Science, 21–48. Springer-Verlag, Berlin, Germany.

Feltovich, P.J., Bradshaw, J.M., Clancey, W.J., Johnson, M., 2006. We regulate to coordinate: limits to human and machine joint activity. Presented at the Proceedings of ESAW 2006, Dublin, Ireland 6–8.

Fitts, P.M. (Ed.), National Research Council, Washington, D.C.

Ford, K.M., Bradshaw, J.M., 1993. Introduction: knowledge acquisition as modeling. In: Ford, K.M., Bradshaw, J.M. (Eds.), Knowledge Acquisition as Modeling, 1–8. John Wiley, New York.

Ford, K.M., Glymour, C., Hayes, P., 1997. Cognitive prostheses. AI Magazine 18 (3), 104 Fall.

Ford, K.M., 1998. On computational wings: rethinking the goals of artificial intelligence. Scientific American (Special Issue on Intelligence), 78–83.

Gaines, B.R., 1985. Personal Communication.

Gaines, B.R., 1993. Modeling practical reasoning. In: Ford, K.M., Bradshaw, J.M. (Eds.), Knowledge Acquisition as Modeling. John Wiley, New York, pp. 51–70.

Gaines, B.R., 2013. Knowledge acquisition: past, Present and future. International Journal of Human-Computer Studies 71 (2), 135–156.

Gawdiak, Y., Bradshaw, J.M., Williams, B., and Thomas, H. .R2D2 in a softball: The Personal Satellite Assistant." Presented at the Proceedings of the ACM Conference on Intelligent User Interfaces (IUI 2000), New Orleans, LA 2000, pp. 125–128.

Gruber, T., 2013. Siri Inc. Nature, Nurture, and Knowledge Acquisition. International Journal of Human-Computer Studies 71 (2), 191–194.

Hoffman, R., Feltovich, P., Ford, K.M., Woods, D.D., Klein, G., Feltovich, A., 2002. A rose by any other name… would probably be given an acronym. IEEE Intelligent Systems 17, 72–80.

Hoffman, R.R., Bradshaw, J.M., Ford, K.M., 2012. Introduction. In: Hoffman, R.R. (Ed.), In Essays in Human-Centered Computing, 2001–2011. IEEE Press, New York City, NY.

Johnson, M., Intlekofer, Jr. K., Jung, H., Bradshaw, J.M., Allen, J., Suri, N., Carvalho, M., 2008. Coordinated operations in mixed teams of humans and robots. In: Marik, V., Bradshaw, J.M., Meyer, J. (Eds.) Proceedings of the First IEEE Conference on Distributed Human-Machine Systems (DHMS2008), Athens, Greece.

Johnson, M., Bradshaw, J.M., Feltovich, P., Jonker, C., van Riemsdijk, B., Sierhuis, M., 2010. The fundamental principle of coactive design: Interdependence must shape autonomy. Coordination, Organizations, Institutions, and Norms in Agent Systems VI, De Vos, M., Fornara, N., Pitt, J., Vouros, G. (Eds.), Springer Berlin, Heidelberg, pp. 172–191.

Johnson, M., Bradshaw, J.M., Feltovich, P.J., Hoffman, R.R., Jonker, C., van Riemsdijk, B., Sierhuis, M., 2011. Beyond cooperative robotics: the central role of interdependence in coactive design. IIEEE Intelligent Systems 26 (3), 81–88.

Klein, G., Woods, D.D., Bradshaw, J.M., Hoffman, R., Feltovich, P., 2004. Ten challenges for making automation a team player in joint human-agent activity. IEEE Intelligent Systems 19 (6), 91–95.

Licklider, J.C.R., 1960. Man-computer symbiosis, IRE Transactions in Electronics. Institute of Radio Engineers, New York 4–11.

Muscettola, N., Nayak, P.P., Pell, B., Williams, B.C., 1998. Remote agent: to boldly go where no AI system has gone before. Artificial Intelligence 103 (1–2), 5–48.

Musen, M.A., 2013. The knowledge acquisition workshops: a remarkable convergence of ideas. International Journal of Human-Computer Studies 71 (2), 195–199.

Shadbolt, N., 2013. Knowledge acquisition and the rise of social machines. International Journal of Human-Computer Studies 71 (2), 200–205.

Shema, D.B., Bradshaw, J.M., Covington, S.P., Boose, J.H., 1990. Design knowledge capture and alternatives generation using possibility tables in Canard. Knowledge Acquisition 2 (4), 345–364.

Shoham, Y. and Tennenholtz, M., On the synthesis of useful social laws for artificial agent societies. Presented at the Proceedings of the Tenth National Conference on Artificial Intelligence, San Jose, CA 1992, pp. 276–281.

Sierhuis, M., Bradshaw, J.M., Acquisti, A., Van Hoof, R., Jeffers, R., and Uszok, A., Human-agent teamwork and adjustable autonomy in practice. Presented at the Proceedings of the Seventh International

Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS), Nara, Japan, 19–23 May, 2003.

Tambe, M., Shen, W., Mataric, M., Pynadath, D.V., Goldberg, D., Modi, P.J., Qiu, Z., and Salemi, B., Teamwork in cyberspace: Using TEAMCORE to make agents team-ready. Presented at the Proceedings of the AAAI Spring Symposium on Agents in Cyberspace, Menlo Park, CA 1999.

Tambe, M., Pynadath, D., Chauvat, C., Das, A., and Kaminka, G.. Adaptive agent architectures for heterogeneous team members. Presented at the International Conference on Multi-Agent Systems (ICMAS 2000) 2000.

United States Department of Defense — Defense Science Board. Task Force Report: The Role of Autonomy in DoD Systems. 2012.

Uszok, A., Bradshaw, J.M., Breedy, M.R., Bunch, L., Feltovich, P., Johnson, M., and Jung, H., New developments in ontology-based policy management: Increasing the practicality and comprehensiveness of KAoS. In: *Proceedings of the 2008 IEEE Conference on Policy*. Palisades, NY, 2008.

Uszok, A., Bradshaw, J.M., Lott, J., Johnson, M., Breedy, M., Vignati, M., Whittaker, K., Jakubowski, K., and Bowcock, J., Toward a flexible ontology-based policy approach for network operations using the KAoS framework. Presented at the Military Communications Conference (MILCOM 2011) 2011, pp. 1108–1114.