

# Joint activity in human-agent-robot teamwork

Paul J. Feltovich, Jeffrey M. Bradshaw, and Mathew Johnson

*Joint-activity theory enables system design to accommodate the demands of interdependence.*

Successful agent-autonomy research has fuelled demand for more sophisticated human-agent-robot teamwork (HART). As we strive to increase autonomous agent and robot capabilities, we necessarily wrest different aspects of control from humans. We thus reduce the human's knowledge and awareness, so that the need for effective coordination becomes greater. This need increases when multiple parties—human or machine—are involved. Coordination requires managing interdependences among activities. Any sort of real teamwork incurs a coordination cost and requires each party to manage some coordination needs. Part of the latter involves assuring that relevant aspects of the agents and the situation are observable and that parties are interacting effectively.

Joint-activity theory highlights three major requirements for coordination, including interpredictability (in highly interdependent activities, one can plan one's own actions only when what others will do can be predicted), common ground (the pertinent mutual knowledge, beliefs and assumptions that support interdependent actions) and directability (the capacity for deliberately assessing and modifying the actions of the other parties in a joint activity).

Following Geertz,<sup>1</sup> we argued that people create cultures and social conventions to provide order and predictability for effective coordination.<sup>2</sup> They construct elaborate regulatory structures, from formal legal systems to norms of proper, everyday behaviour. HART can exploit such mechanisms to support coordination of complex, interdependent activity.<sup>3</sup> People coordinate through signals and more complex messages (such as direct language or posture). Human signals are also mediated in many ways, for example through third parties or machines. Hence, direct and indirect party-to-party communication is one form of a 'coordination device' (coordination by *agreement*). There are three other common coordination devices, notably convention (guides for action, e.g., rules), precedent (how we have done it before) and salience (what the situation suggests).<sup>4</sup>



*Figure 1. Policies constitute an agent's 'rules of the road,' not its route plan.*

Roles are ways of packaging rights and obligations associated with parts that people play in joint activities, possibly including multiple roles for any one individual. Knowing one's own role(s) and those of others in a joint activity establishes expectations about how interactions will proceed. Collections of roles are often grouped to form organizations. The order needed for agents in joint activity is typically implemented in terms of formalized social regulations. Researchers have introduced the theme of social laws under two main headings, including norms and policies. In the multi-agent-system research community, investigators have argued that norms are variously described as guiding behaviour, goals or obligations.

While sharing much in common with norm-based approaches, policy-based perspectives differ in subtle ways. In contrast to the relatively descriptive nature and self-chosen adoption of norms, policies are prescriptive and externally imposed. While the former emerge gradually in everyday life, the latter are consciously designed and put into or out of force. Policy management should not be confused with planning. While policies are the 'rules of the road' (see Figure 1)—providing the stop signs, speed limits, etc. that serve to coordinate traffic and minimize mishaps—they are not sufficient for route planning.

*Continued on next page*

We coined the term ‘co-active design’ to characterise an approach to HART that takes interdependence as the central organising principle.<sup>5</sup> In addition to implying multiple parties, ‘co-active’ conveys the reciprocal and mutually constraining nature of actions that are conditioned by coordination. In joint activity, individuals must sacrifice, to a degree, their individual autonomy for group goals. Important considerations within co-active design include team- and task work, mutual affordances and obligations, soft dependencies, joint goals and mixed-initiative opportunities in all phases of the sense-plan-act cycle.

Co-active design complements task-focused approaches to HART, such as adjustable autonomy and mixed-initiative interaction, with a strong focus on teamwork. For example, the task work of playing football includes kicking and dribbling. The teamwork aspect focuses more on, e.g., allocating players to roles and synchronizing tactics.

Software agents are often described as assistants to people. While this one-way relationship is sometimes helpful, it is inadequate for describing joint human-agent activity. Joint activity implies greater parity of mutual assistance, enabled by webs of complementary, reciprocal affordances and obligations.

Co-active design emphasizes the importance of both hard and soft dependencies in coordinating interrelated activities. Hard dependencies are necessary, or the joint activity could not happen at all, such as the passing of a baton in a relay race. Soft dependencies are not strictly necessary but are helpful. Attending to soft dependencies is a subtle, but no less significant process. For instance, the first runner may shout something to the second runner before the handoff to convey a warning about a slippery section of track ahead.

Multi-agent teamwork research typically has held a simple view of joint goals, based on unification of symbols common to all parties. However, more sophistication is needed when humans are involved. Apart from the problem of establishing and maintaining common ground on complex goals and the best means to achieve them, team goals sometimes compete with the goals of individuals.

Mixed-initiative interaction, where the roles and actions of people and agents are opportunistically negotiated during problem solving, has typically been limited to planning and command generation. To these can be added aspects of perception and cognition. Co-active design extends earlier work in all phases of the sense-plan-act cycle, consistent with the contention that any required resource or capability within an agent’s action-perception loop affords a possible point for interdependence.

Co-active design accentuates designing agents with the capabilities they need to be interdependent from the start.

Based on these considerations, we summarise some of the key characteristics of a good agent—human or artificial—with regard to joint activity. A good agent is observable (it makes its pertinent state and intentions apparent), attuned to the requirement of progress appraisal (it enables others to stay informed about the status of its tasks and any potential trouble spots), informative and polite (it knows enough about others and their situations to tailor its messages to be useful), predictable and dependable, directable at all points of the sense-plan-act cycle (it can be directed to what is most important in a given situation), coordinated (it helps communicate, manage and deconflict interdependencies among activities, knowledge and resources that are prerequisites to effective task performance and ‘common ground’) and knows its limits (it knows when to take the initiative and when it needs outside direction). We continue to study co-active system-design issues that affect the success of coordination within human-agent-robot teams.

---

## Author Information

**Paul J. Feltovich, Jeffrey M. Bradshaw, and Mathew Johnson**  
Florida Institute of Human and Machine Cognition  
Pensacola, FL, USA  
<http://www.ihmc.us/groups/jbradshaw/>

Paul Feltovich is a research scientist. He holds a PhD from the University of Minnesota (USA) and was a professor at Southern Illinois University School of Medicine (USA) from 1982 to 2001. He has conducted research in human expertise, difficult learning and human-agent teamwork.

Jeffrey Bradshaw is a senior research scientist. He leads the research group developing the Knowledgeable Agent-oriented System (KAoS) policy framework. Among many other publications, he edited *Knowledge Acquisition as a Modeling Activity* (with Ken Ford, Wiley, 1993) and *Software Agents* (Association for the Advancement of Artificial Intelligence Press/The MIT Press, 1997).

Matthew Johnson has been working on human-robotic systems at the Institute for Human and Machine Cognition for eight years. He has a BS in aerospace engineering and an MS in computer science. He has worked on advanced cockpit

*Continued on next page*

displays, augmented cognition and several human-robot coordination projects.

## References

1. C. Geertz, **The Interpretation of Cultures**, Basic Books, New York, 1973.
2. P. Feltovich, J. M. Bradshaw, W. J. Clancey, and M. Johnson, *Toward an ontology of regulation: support for coordination in human and machine joint activity*, **Engineering Societies in the Agents World. VII. (Lect. Notes Comput. Sci.)**, pp. 175–192, 2007.
3. J. M. Bradshaw, P. J. Feltovich, and M. Johnson, *Human-agent interaction*, in G. Boy ed., **Handbook of Human-Computer Interaction**. In press.
4. H. H. Clark, **Using Language**, Cambridge Univ. Press, Cambridge, UK, 1996.
5. M. Johnson, J. M. Bradshaw, P. Feltovich, C. Jonker, B. van Riemsdijk, and M. Sierhuis, *Coactive design: why interdependence must shape autonomy*, **Coordination, Organizations, Institutions, and Norms in Agent Systems. VII. (Lect. Notes Comput. Sci.)**. In press.