

WHAT IS A COGNITIVE SYSTEM?

Gavan Lintern
General Dynamics – Advanced Information Systems
Dayton, Ohio 54431
glintern@earthlink.net

The theme for this year's symposium, *The Airspace as a Cognitive System*, stimulates the questions; what is a Cognitive System and in what sense can we characterize the airspace as a cognitive system? I discuss these questions by reviewing ideas promoted in discussions of distributed cognition. I also contrast the concept of *distributed* with the similar concepts of shared and joint as they are discussed in the literature on cognitive systems, team training and situation awareness. I conclude that the notions of distributed cognition and joint cognitive systems offer considerable leverage for addressing the anticipated design challenges in airspace systems but that we need to avoid the distortions engendered by the pervasive techno-centric emphasis in systems design in favor of a human-centric emphasis that will aid development of robust and effective systems.

Introduction

The airspace is a distributed and heterogeneous system comprising diverse human and technological functions. The theme for this year's symposium, *The Airspace as a Cognitive System*, stimulates the questions; what is a Cognitive System and in what sense can we characterize the airspace as a cognitive system?

A cognitive system is a one that performs the cognitive work of knowing, understanding, planning, deciding, problem solving, analyzing, synthesizing, assessing, and judging as they are fully integrated with perceiving and acting. The characterization of the airspace as a cognitive system represents a claim that the airspace is an entity that does cognitive work.

The claim that the airspace does cognitive work expands the view of what is *cognitive* beyond the individual mind to encompass coordination between people and their use of resources and materials. This view is aligned with the theory of distributed cognition enunciated by Hutchins (1995) and further described by Hollan, Hutchins and Kirsh (2000). A foremost claim of this theory is that distributed cognition is not a theory about a special type of cognition but rather a theory about fundamental cognitive structures and processes (Hollan et al, 2000). Thus, all cognition is distributed.

Traditionally, we are used to thinking that cognition is an activity of individual minds but from the perspective of distributed cognition, it is a joint activity that is distributed across the members of a work or social group and their artifacts. Cognition is distributed spatially so that diverse artifacts shape cognitive processes. It is also distributed temporally so that products of earlier cognitive processes can shape later cognitive processes. Most significantly,

cognitive processes of different workers can interact so that cognitive capabilities emerge via the mutual and dynamic interplay resulting from both spatial and temporal coordination among distributed human agents.

A distributed cognitive system is one that dynamically reconfigures itself to bring subsystems into functional coordination. Many of the subsystems lie outside individual minds; in distributed cognition, interactions between people as they work with external resources are as important as the processes of individual cognition. Both internal mental activity and external interactions play important roles as do physical resources that reveal relationships and act as reminders. A distributed system that involves many people and diverse artifacts in the performance of cognitive work is therefore properly viewed as a cognitive system.

The theory of distributed cognition forces a shift in how we think about the relationship between minds, social interactions and physical resources. Interactions between internal and external processes are complex and unfold over different spatial and time scales and neither internal nor external resources assume privileged status.

The Defining Example

In the early 90s, the concept of distributed cognition stimulated considerable interest. Nevertheless, different commentators had different views of what that concept encompassed. Furthermore, these diverse views were typically not well grounded in reality. Within that scientific environment, the approach taken by Hutchins (1995) was refreshing. He developed a narrative description of distributed cognition in action that illustrated, with exceptional clarity, how he thought about distributed cognition.

That description was grounded in the activities of a shipboard navigation team as they navigated a US Navy ship through enclosed waters. Hutchins argued that the navigation team, together with accompanying navigational artifacts and procedures, is a cognitive system that performs the computations underlying navigation.

For enclosed waters, navigation involves successive plots of position, which permit inference of ship speed and direction (Figure 1). A plotting cycle is initiated by the bearing recorder, located in the pilothouse, who advises the pelorus operators on the wings of the bridge of the time to take sightings. The pelorus operators advise the bearing recorder of the landmark bearings, who records them. The navigation plotter, also located in the pilothouse, reads the bearings and plots the position of the ship at the time of the observations. Via repeated position plots, the course and land-reference speed of the ship is established.

This style of navigation is a product of a distributed cognitive system in that various elements of the computations are carried out over time and in different locations. The results of early computations are passed to another location and then integrated into a further computation.

Such a distribution of processes underlying cognition can result in a computation of greater complexity than is achievable by any member of the system individually. However, this is not just a matter of

more power from greater numbers. Hutchins argues that the system has cognitive properties that differ from the cognitive properties of the individuals and that the cognitive potential of the group depends more on its social organization than on the cognitive potentials of its members (See Box 1). Thus the navigational system performs computations that need not necessarily be within the grasp of all (or even any) of its members.

Cognition is Emergent

Cognitive capabilities emerge from activity in relation to a Cognitive System's architecture and are shaped by that architecture. This cognitive architecture is a synergy of the functional structure of the physical environment, the social organization of the work place and the functional structure of individual minds. As implied by the word, *synergy*, cognitive capability is not merely a sum of capabilities of parts; the interaction and interplay between subsystems generates a cognitive power beyond that of any subsystem, whether artifact or individual. This view of Cognitive Systems forces a shift in how we think about the relationship between minds, social interactions and physical resources. Interactions between internal and external processes are complex and unfold over different spatial and time scales. If collaborations are more coordinated, more effective, more robust, and more meaningful, the distributed Cognitive System will be cognitively more capable.

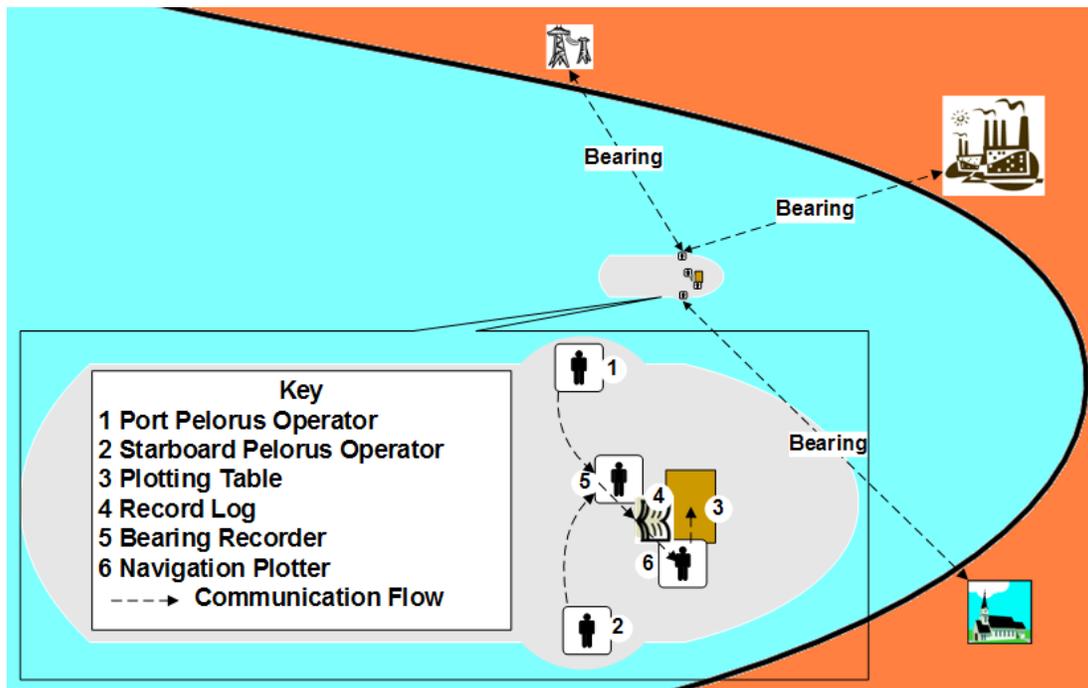


Figure 1; Navigation in enclosed waters as a distributed cognitive system

Box 1: Foraging Ants

The meaning of the claim that the cognitive properties of a social system can differ from those of its individuals might be clarified by an example from Franks (1989). Army ants forage in distinctive patterns. Over successive days, they rotate their radial direction of foraging by approximately 22.5 degrees with the result that the area around their bivouac site is systematically depleted after 16 raids. On the 16th raid, instead of returning to their established bivouac site, the ants travel to a new site. However, no single ant has a plan to employ this efficient foraging strategy. Ants deposit pheromone trails and then respond to the odor thus generated. This apparently intelligent behavior cannot be referenced to the intelligence of any of the participants. It is system intelligence but not shared intelligence.

The nature of emergent properties is poorly understood in general as are the processes of self-organization that generate emergent properties. An emergent property is one that has no identifiable description of its form in its microstructures or processes (Box 1). Thus thermal convection rolls in a heated fluid are said to be emergent (Figure 2) while a construction from a detailed plan is not. Theories that posit a mental image, a mental model or a mental schema as a formative cause of cognition eschew self-organization (e.g., Johnson-Laird, 1983). In contrast, some argue that an understanding of self-organization is central to understanding cognition (e.g., Lintern & Kugler, 1991; Lintern 2001).

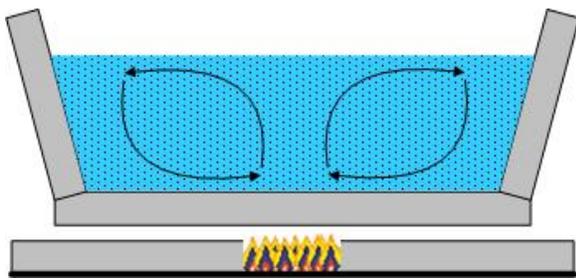


Figure 2; A Rayleigh Bernard Convection offers a common illustration of self-organization.

Those who promote self-organization as an explanation of cognition typically emphasize the role of local interactions in the development of patterns and might offer self-organization as a bottom-up, emergent view in contrast to the top-down view of mental imagery as the shaping influence on cognition. Some caution is needed here. While local constraints

play an important role in self-organization, it is the interaction between local and global constraints that generate the emergent patterns. For example, while local interactions between molecules are important, thermal convection rolls would not emerge without the global constraints of a heat gradient and the containment vessel (Figure 2).

Similarly, cognitive emergence owes as much to the functional layout of the environment as it does to the local interactions of individuals with each other and with artifacts. The cognitive architecture determines the way information flows through the system. This architecture encompasses the functional structure of the physical environment, the social organization of the work place and the functional structure of individual minds. New cognitive capabilities emerge from activity undertaken within the constraints imposed by the cognitive architecture and are shaped by those architectural constraints.

Technological Function & Cognition

The theory of distributed cognition is consistent with the view that a cognitive system is a thinking (or intelligent) information system. However, the enhanced intelligence is not generated by the activity of intelligent technological functions as many in the discipline of Artificial Intelligence will want to claim, but emerges from the coordinated collaboration of distributed human agents via their interactions with each other and with functionally heterogeneous technological artifacts. In the sense that collaborations between human agents and their use of technological artifacts are coordinated, effective, robust, and meaningful, the distributed system is intelligent.

It is sometimes argued that computer based agents can be employed to reason about the beliefs of human participants in teams (D'Inverno, Luck, Georgeff, Kinny & Wooldridge, 2004). From the perspective of distributed cognition, technological devices do not reason; people reason. Two people in coordination can possibly reason more effectively than either in isolation, and if they (as a coordinated dyad) avail themselves of the opportunities presented by technological devices that can compute logical relationships, find and organize information, and probably offer a number of as yet unimagined supporting functions, these entities (the two people together with the technological devices) constitute a reasoning system.

Joint, Shared & Distributed Cognition

The allied but not identical adjectives, *joint* and *shared*, vie for dominance with *distributed* in the team, situation awareness and cognitive systems

literature (e.g., Endsley, Hansman & Farley 1999; Hollnagel and Woods, 2005; Stout, Cannon-Bowers, Salas, & Milanovich, 1999). No distinction has yet been drawn between these three modifiers and some clarification is overdue.

To be distributed is to be spread or diffused spatially and temporally. To be joint is to be coordinated in united action. To be shared is to be owned or possessed by all participants. At first glance, it might seem that these three modifiers could each add something to our view of cognition. They address different dimensions: a distributed system has spatial and temporal extent, a joint system is an integration of sub-systems and a shared system is one that can be accessed equally by all participants.

Nevertheless, this is not an entirely comfortable rapprochement. The notion of sharing suggests a common core shared by all coordinating entities. While the notions of distributed and joint allow elements of mutual understanding, neither suggests the desirability of a common core shared by all. They undoubtedly imply shared understandings but neither implies system wide, equal access to a common core. They rather suggest opportunistic, asymmetric and fragmented but functionally meaningful sharing by subgroups. On the other hand, the terms distributed and joint are complementary; a distributed system is made up of coordinated sub-systems while a joint system is necessarily distributed.

In scientific discourse, this compatibility has unfortunately become strained. When restricted to human agents, the intentional parity implied in discussions of Joint Cognitive Systems remains compatible with Hutchins' discussions of distributed cognition. The strain emerges in the views expressed by Hollnagel and Woods (2005) that technological agents can be viewed as team players. That suggests intentional parity between technological artifacts and humans. Discussions of distributed cognition do not suggest intentional parity between human and technological agents. In particular, the navigation narrative offered by Hutchins (1995) suggests that technological artifacts are necessarily subordinate to human agents. In this view, technological artifacts are tools that support and extend the cognitive capabilities of the humans who guide and direct the system.

While the concept of sharing has little to offer, the concepts of distributed and joint cognition can both strengthen our conceptualization of the airspace as a cognitive system. The airspace is inevitably a system of coordinated sub-systems, thereby conforming to the fundamental nature of joint systems. Its functions and activities are also inevitably distributed spatially

and temporally so that it conforms to the fundamental nature of a distributed system. These ideas, if brought to the fore, will benefit our conceptualization of the airspace as a cognitive system.

The Airspace as a Cognitive System

Ongoing developments in air traffic control and air management systems are motivated largely by obsolescence of previous generation technology and by expectations from traffic density projections that our current systems will soon be overloaded. The tendency is to emphasize technology as a solution and there is no doubt that dramatic advances in technology offer new opportunities that were not available during development of previous generation air traffic control and air management systems.

Nevertheless, the lessons of cognitive engineering, particularly from investigations of distributed cognition and joint cognitive systems emphasize the crucial, integrative role that human agents play in complex socio-technical systems. The problems of over-reliance on technological solutions together with neglect of the human role have been cogently illustrated in the early developments of highly automated cockpits. The groundbreaking work of Sarter and Woods (1994) should give us pause. While no one would wish to return to the pre-computer days of mechanical and hard-wired systems, it should now be evident that the design of a distributed cognitive system is not just a matter of building better technical artifacts.

There remains however, a substantial rational imperative to rely predominantly on technological development. The strong field of artificial intelligence is at the forefront in promoting that rational imperative but despite lavish promises from that quarter (Brighton, 2004), it is not just a matter of building intelligent devices. That will result inevitably in a human role that is subservient to technological functionality. A typical result of such technological dominance is a system that is elegant and efficient but also brittle. Most troubling is that technologically inspired solutions for socio-technical systems impose a high cognitive load on the human participants in the system at the worst possible times. Thus the term *clumsy automation* has become an evocative catch phrase.

As a member of this symposium's organizing committee, I concur with the view that the airspace can be viewed as a cognitive system (Figure 3). It behooves all of us to ensure that the lessons emerging out of treatments of distributed cognition and joint cognitive systems are heeded so that we develop an airspace system that is robust and more intelligent

principally because it amplifies rather than replaces the cognitive and coordinative capabilities of the human participants.

We must build on the ideas expressed in treatments of distributed cognition and joint cognitive systems to identify how new technological functionality can be used to support the cognitive work undertaken by human agents and how it can be used to facilitate better collaboration between distributed human agents. Much has been said within the community of Cognitive Systems Engineering about how we might proceed to build better cognitive systems through emphasis of the coordinating, adaptive and sense-making roles played by the human participants and I will not repeat it here. However, the principal lesson is that we need to develop a coordinated system of human agents and technological functionality in which there are effective communication tools to support collaboration between human agents and effective interfaces that support their use of the technological functionality.

It is imperative that we are not seduced by the techno-centric aura that constrains current development of socio-technical systems and it is important that our discussions do not encourage a techno-centric focus. There is a danger that technologists will find, in the notion of technological artifacts as team players, justification for the perverse and fruitless pursuit of technological solutions at the expense of integrating and supporting unique and critical human functionality. From the cognitive engineering perspective, we must combat this science fiction fantasy that technologists can somehow automate all critical human functions in case we end up with a system in which humans have no more than a peripheral role or even no role at all.

References

Brighton, Henry (2004). *Introducing Artificial Intelligence*. Totem Books (Paperback), ISBN-10: 184046463.

Endsley, Mica R., Hansman, R. John & Farley, Todd C. (1999). *Shared Situation Awareness in the Flight Deck – ATC System*. Digital Aviation Systems Conference, Seattle, Washington.

D'Inverno, Mark; Luck, Michael; Georgeff, Michael; Kinny, David & Wooldridge, Michael (2004). *The dMARS Architecture: A Specification of the Distributed Multi-Agent Reasoning System*. *Autonomous Agents & Multi-Agent Systems*, 9 (5), 53.

Franks, N. R. (1989). *Army ants: A collective intelligence*. *American Scientist*, 77, 139-145.

Hollan, James; Hutchins, Edwin & Kirsh, David (2000). *Distributed Cognition: Toward a New Foundation for Human-Computer Interaction Research*. *ACM Transactions on Computer-Human Interaction*, Vol. 7, No. 2, Pages 174–196.

Hollnagel, Erik & Woods, David D. (2005). *Joint Cognitive Systems: Patterns in Cognitive Systems Engineering*. Boca Rotan, FL: CRC Press. ISBN: 0849328217.

Hutchins, E. 1994. *Cognition in the Wild*. MIT Press, Cambridge, MA.

Johnson-Laird, P.N. (1983). *Mental models*. Cambridge, MA: Harvard.

Lintern, G. (2001). *Distributed mission training: Issues and directions*. In R. Jensen, (Ed), *Proceedings of the Eleventh International Symposium on Aviation Psychology, March 5 – 8, 2001, Columbus, Ohio* [CD-ROM]. Columbus, OH: The Ohio State University Department of Aerospace Engineering, Applied Mechanics and Aviation.

Lintern, G., & Kugler, P. N. (1991). *Self organization in connectionist models: Associative memory, dissipative structures, and Thermodynamic Law*. *Human Movement Science*, 10, 447-483.

Sarter, N. B., & Woods, D. D. (1994). *Pilot interaction with cockpit automation II: Operational experiences with the flight management system*. *The International Journal of Aviation Psychology*, 4, 1-28.

Stout, R. J., Cannon-Bowers, J.A., Salas, E. & Milanovich, D.M. (1999). *Planning, shared mental models, and coordinated performance: An empirical link is established*. *Human Factors*, 41, 61-71.

Woods, David D. & Hollnagel, Erik (2006). *Joint Cognitive Systems: Patterns in Cognitive Systems Engineering*. Boca Rotan, FL: CRC Press. ISBN: 0849339332.

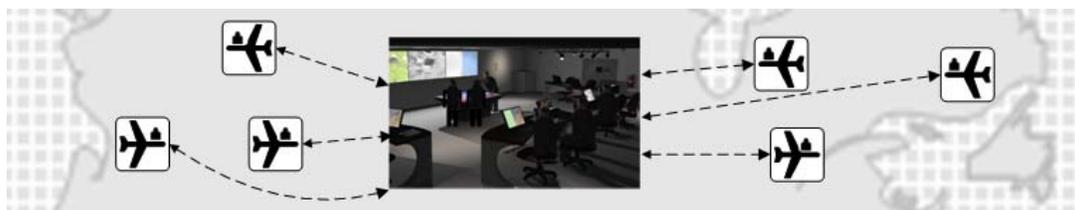


Figure 3; Air Traffic management as a joint, distributed cognitive system