

# The Airspace as a Cognitive System

Gavan Lintern

*Cognitive Systems Design, Melbourne, Australia*

The theme for this special issue, the airspace as a cognitive system, stimulates these 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 conclude that this notion of the airspace as a cognitive system offers 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-focused emphasis that will aid development of robust and effective systems.

The airspace is a distributed and heterogeneous system that includes diverse human and technological functions. In this article, I argue that the airspace can be viewed as a cognitive system and from that perspective, I consider the nature of a cognitive system and discuss why it is useful to characterize the airspace as such.

## IS THE AIRSPACE A COGNITIVE SYSTEM?

The defining characteristic of a system versus an assemblage is that the constituent parts or subsystems work together. In physical systems, that is accomplished by exchanges of physical energy as constrained by force fields. In cognitive systems, it is accomplished by exchanges of information as constrained by information fields.

We normally think of cognition as something that happens in the head of a single individual, but a cognitive system is more than that. A whole person, a perceiv-

ing, thinking, acting entity, is a cognitive system. Each human in a system is, himself or herself, a cognitive system, but the larger entity of humans working in collaboration with each other and with the support of capabilities provided by various technologies is also a cognitive system.

A cognitive system is one that performs cognitive work via cognitive functions such as communicating, deciding, planning, and problem solving (Figure 1) as, for example, in military command and control, transportation, health care, and air traffic management. These sorts of cognitive functions are supported by cognitive processes such as perceiving, analyzing, exchanging information, and manipulating. The characterization of the airspace as a cognitive system represents a claim that the airspace is an entity that does cognitive work.

### DISTRIBUTED COGNITION

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.




FIGURE 1 Cognitive work functions are supported by cognitive processes.

## The Theory of Distributed Cognition

This view is aligned with the theory of distributed cognition outlined 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 cognitive work 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.

## A Defining Illustration

In the early 1990s, 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 it. That description was grounded in the activities of a navigation team as it guided a U.S. 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 permits computation of ship speed and direction (Figure 2). A plotting cycle is initiated by the bearing recorder, located in the pilothouse, who advises the

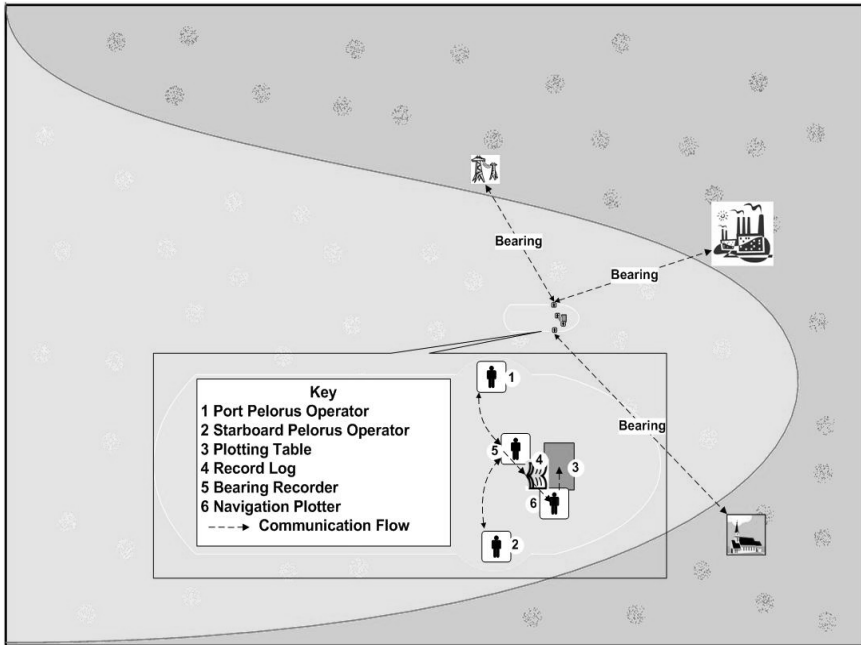


FIGURE 2 Navigation in enclosed waters as a distributed cognitive system.

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, and he or she 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. The course and land-reference speed of the ship is established via repeated position plots.

This style of navigation is a product of a distributed cognitive system in which 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 can be achieved by any member of the system individually. However, this is not just a matter of more power from greater numbers. The system has cognitive properties that differ from the cognitive properties of the individuals and the cognitive potential of the group depends more on its social organization than on any combination of the individual cognitive potentials of its members. Thus the navigational system performs computations that might not be entirely within the grasp of all (or even any) of its members.

## COGNITIVE SYSTEMS FOR MISSILE DEFENSE

It is often said that humans have limited processing capacity and are error prone and, as a result, we should seek to design humans out of systems (Seng, Wai, Chan, & Weng, 2009). Here I offer a contrasting view: Systems work because of, not in spite of, their human participants. The unique communicating, deciding, planning, and problem solving that human participants provide are essential to safe and effective operation. I illustrate the power of human cognition by reference to issues surrounding missile defense.

### The Power of Human Cognition

Klein (1999) described an incident that occurred during the first Gulf War (August 2, 1990–February 28, 1991) in which a British warship, the *HMS Gloucester*, shot down an incoming Iraqi Silkworm missile. The *Gloucester*, stationed some 20 miles off the coast of Kuwait, had not, to that point in the war, encountered incoming missiles. It had, however, encountered many allied aircraft returning from bombing runs over Iraq and Kuwait. Ostensibly, the radar return of the missile was indistinguishable from radar returns of allied aircraft, yet the tactical officer identified the missile and ordered a counteraction that resulted in the *Gloucester's* own missiles destroying it.

The incident lasted around 90 seconds. Subsequently, there was a concern that the *Gloucester* had destroyed an allied aircraft. That concern proved to be unfounded. The incoming object had indeed been a Silkworm missile, although the tactical officer could not explain how he knew that. This puzzle remained unresolved for several months. What was known was that, within this scenario, allied fighter aircraft and Silkworm missiles traveled at similar speeds and were around the same size, therefore presenting a similar radar profile. It was also known that they traveled at different altitudes, around 3,000 ft for the aircraft and 1,000 ft for the missile, although that information was not directly available from the radar.

The resolution of the puzzle does, however, lie in the altitude difference. The aircraft and the missile became visible on the radar only when they emerged from ground clutter as they crossed the coastline. Given its lower altitude, the missile was closer to the *Gloucester* when it first emerged from ground clutter. This difference, not consciously recognized by the tactical officer, sensitized him to the fact that this was probably not an aircraft and was, in all likelihood, a hostile missile.

### The Fragility of Technical Logic

Let us now move forward to the second Gulf War (beginning March 20, 2003). The Patriot missile system was deployed to defend allied installations against tactical ballistic missiles. It proved itself in its defensive role, accounting for all nine

enemy tactical ballistic missiles launched against areas defended by the Patriot system.

With a decade or so having elapsed since the first Gulf War, we might imagine that advances in technology would have reduced the potential for fratricide from allied missiles. Regrettably, that is not the case. Two allied aircraft were destroyed by the Patriot system, resulting in the deaths of three aviators. The Defense Science Board (2005) concluded that the operating philosophy of the Patriot system did not match the conditions of this conflict. In contrast to the situation on board the *Gloucester* approximately a decade earlier, the Patriot system was heavily automated. Operators were trained to trust the system software and, most notably for the argument I develop here, they were not accorded the opportunity to develop cognitive skills specific to the conditions of this demanding work environment.

## WHERE IS THE COGNITION?

Nevertheless, technological artifacts remain as important elements of distributed cognitive systems; they facilitate cognitive processing and at the very least, cognitive systems that are widely distributed in time or space could not function without them. Do we need to treat technological artifacts as if they are also performing cognitive processes or should we view them as merely facilitating the cognitive processing of the humans within the system? That is, can technological artifacts perceive, analyze, exchange information, and manipulate (Figure 1), or do they play only a supporting role.

### Conceptions of Cognition

Traditionally, cognition is equated to information processing and more recently to computation. Both have cognition being realized through creation, transformation, and propagation of representational states.<sup>1</sup> Although it is helpful to think of cognition in these terms, we should recognize that information processing and computation are broad concepts and that not all information processing and not all computation are cognitive processes. The common home thermostat, for example, is a computational device that processes information but it is not, by itself, a cognitive system.

Blomberg has, in this special issue, reviewed various conceptions of cognition:

---

<sup>1</sup>It is often said that these conceptions are metaphors for cognition (i.e., an implied comparison) but cognitive scientists treat them more as analogues (i.e., an equivalence relationship used as a basis for explanation).

- An internal view in which the cognitive processing is restricted to individual humans<sup>2</sup> although what happens externally might facilitate cognition but is not, in itself, cognitive.
- An extended view in which human cognitive processing is coupled with distributed, external entities that, by playing an active causal role, jointly govern behavior in the same sort of way that cognition usually does.

Notably, the internal view constrains cognitive processing to the humans in the system whereas the extended view does not. In what follows, and in contrast to Blomberg (this issue), I promote the internal view by arguing that cognition is a uniquely human enterprise.

### The Artificial Intelligence Debate

A coordinated system made up of multiple entities might be intelligent, but 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. Rather, it emerges from the coordinated collaboration of distributed human agents via their interactions with each other and their interactions with functionally heterogeneous technological artifacts. The internal view of cognition to which I subscribe explicitly rejects the possibility of intelligent technological artifacts. This has, however, been a point of contention within the scientific and philosophical literature for decades.

The argument for artificial intelligence goes something like this. Intelligence is a reflection of effective cognition as realized through creation, transformation, and propagation of representational states (i.e., information processing or computation). The representational states denote knowledge in the form of facts, contingency relationships, and contextual associations. New knowledge can be created from existing knowledge through transformations and propagations. We can create an intelligent agent by programming all of this into a computer.

Critics of the artificial intelligence program counter that we should be skeptical; demonstrations that purport to exemplify artificial intelligence are largely constrained to closed, formal (rule-based) systems such as chess. The natural world is more complex and beyond the scope of any artificial intelligence program. Haugeland (1979), for example, argued that we, as natural cognitive systems, are very good at resolving situational ambiguity. He offers an excerpt from a folk tale to illustrate:

One evening, Khoja looked down into a well, and was startled to find the moon shining up at him. It won't help anyone down there, he thought, and he quickly fetched a

---

<sup>2</sup>Ignoring, for the purpose of this article, the possibility that other biological systems also perform cognitive processing.

hook on a rope. But when he threw it in, the hook snagged on a hidden rock. Khoja pulled and pulled. Then suddenly it broke loose, and he went right on his back with a thump. From where he lay, however, he could see the moon, finally back where it belonged—and he was proud of the good job he had done. (Haugeland, 1979, p. 625)

We, as readers, recognize at least implicitly that this narrative switches between an imagined, misconceived situation and the real one. The humor emerges from Khoja's naive interpretation of the world and his pride in his imagined accomplishment. The implication of Haugeland's argument is that this sort of situational complexity is pervasive in literature and in the world and that a seamless understanding of it is something that could not reasonably be expected of anything other than a natural cognitive system.

This is not, however, a compelling argument. A technological system might track the situational switches if programmed with appropriate facts, contingency relationships, and contextual associations. To program the appropriate knowledge for all books ever written or for all world situations would be a huge and probably impossible task, but that is an unrealistic benchmark for a cognitive system; no individual human understands all of the nuances and implications of all books ever written or all world situations and we often do not understand all of the nuances and implications of any particular one of them.

### How Is Cognition Uniquely Human?

The more serious challenge for artificial intelligence is exemplified in the *Gloucester* incident where the tactical officer developed a new capability to perceive an important situational property. That new capability was exquisitely tuned to the demands of the work situation and that tuning illustrates a unique and powerful capacity of the human cognitive system. We adapt continuously to situational demands. Some of that adaptation is based on new understandings (facts, contingency relationships, contextual associations) but some is based on changes in system capability (new property detectors in our perceptual systems, new cognitive processing strategies, new coordinative patterns).

In technological parlance, we change our functionality, which is a characteristic that engineers have historically sought to design out of technological artifacts. Whereas designers of technology seek to develop stationary systems (those that maintain their functionality over time), the human cognitive system is non-stationary (its functionality changes over time). Of particular significance is that many changes in functionality are not driven by a developmental time line (at least not in the human adult) or by intent, but emerge spontaneously in response to situational exigencies. Thus the functionality of the human cognitive system adapts, changing in ways that prepare it for new, unanticipated situational demands.



We could, of course, add functionality to an artificial intelligence system, but the functional changes in the human cognitive system emerge via a self-organizing process and the new functionality is specific to situational exigencies. It is this self-organizing capability of the human cognitive system that is at the basis of adaptability in the face of novel events and it is a capability that is uniquely human.

### Why Does This Debate Persist?

The standard approach to the artificial intelligence debate has two sides. On one side are the technologists who regard this as an objectively logical problem of mechanism. On the other side are those who think of the human cognitive system as so complex as to be beyond any conceivable artifact. The fundamental flaw in all of this is that the argument is based on outdated psychological theories; a science of the artificial (Simon, 1981) as found in communication, information, and computation theories.

As a result, the discipline of psychology continues to ask the wrong questions (Reed, 1996). We remain overly concerned about the rational properties of knowledge when we should rather be concerned with the nature of how cognition supports our adaptive interactions with the world. We need to reorient from being concerned with storage of facts, contingency relationships, and contextual associations to being concerned with how we regulate our functional action as we encounter the world (Reed, 1996). In short, we need to subscribe to a science of the natural as found in many areas of biology and in contemporary evolutionary theory (Clancey, 1997; Edelman, 1987; Freeman, 1995).

## COGNITION IN SOCIO-TECHNICAL SYSTEMS

Cognitive systems are intentional; that is, they have goals or purposes. In contrast, technological artifacts do not have goals or purposes; they do not care about anything (Haugeland, 1979). How might we clarify the relationship among technological functionality, cognitive functionality, and goals or purposes. An answer to this question can be found in reference to the abstraction hierarchy, which has been used to good effect in three other articles in this special issue (Millen et al., this issue; Neal et al., this issue; van Marwijk, Borst, Mulder, Mulder, & van Paassen, this issue).

The abstraction hierarchy is a structural description of a work domain over five levels (domain purpose, domain values, domain functions, physical functions, and physical resources; see Figure 3). Means–ends relations identify enabling relationships between levels. A means–ends relation is two-way. It shows which structural elements at one level support a particular structural element at the level above and also shows which structural elements at any particular level are supported by a se-

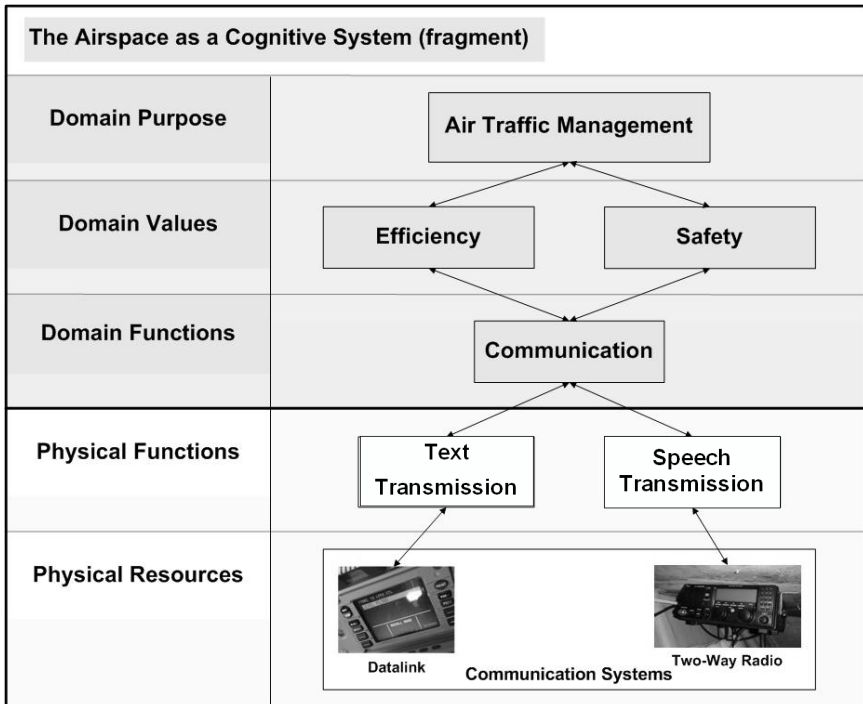


FIGURE 3 A fragment of an abstraction hierarchy for the commercial airspace as a cognitive system

lected structural element from the level below. Mappings between levels can be one-to-one, one-to-many, many-to-one, or many-to-many.

Figure 3 depicts a fragment of an abstraction hierarchy for the commercial airspace as a cognitive system. This fragment features communication<sup>3</sup> as the domain function. It is supported by the physical (noncognitive) functions of speech and text transmission, which, in turn, are supported by physical resources (data link, two-way radio). Communication serves to support realization of values (efficiency and safety) that are important to the domain purpose of air traffic management.

## 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 expecta-

<sup>3</sup>A cognitive work function (see Figure 1).

tions from traffic density projections that our current systems will soon be overloaded. Many of us have a role to play in these developments. From the perspective that the airspace is a cognitive system (Figure 4), it is imperative that we build the essential functionality into the system. As we design the airspace, we can do many things in terms of substituting types of devices (e.g., communication systems of various types). However, at least some of the functionality has to be cognitive. We cannot replace all of the devices that do cognitive processing with devices that do no cognitive processing, and the only entity that we know has cognitive functionality is us humans.

There remains, however, a rational imperative to rely predominantly on technological development. The problems of overreliance on technological solutions together with neglect of the human role have been cogently illustrated in the early developments of highly automated cockpits where the groundbreaking work of Sarter and Woods (1994) should give us pause. Although no one would wish to return to the precomputer days of mechanical and hard-wired systems, technological dominance in design of socio-technical systems has produced solutions that are elegant and efficient, but also brittle. Most troubling are those technologically inspired solutions that impose high cognitive load on the human participants in the system at the worst possible times. There is no doubt that dramatic advances in




FIGURE 4 Air traffic management as a joint, distributed cognitive system.

technology offer new opportunities that were not available during development of previous generation air traffic control and air management systems. Nevertheless, this is not just a matter of building better technical artifacts.

## CONCLUSION

The lessons of cognitive engineering, particularly from investigations of distributed cognition, emphasize the crucial, integrative role that human agents play in complex socio-technical systems. Much has been said within the cognitive systems engineering community about how we might proceed to build better cognitive systems through emphasis on the coordinating, adaptive, and sense-making roles played by the human participants and I do 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. To do that, we must develop an airspace system that is robust and intelligent principally because it amplifies rather than replaces the cognitive and coordination capabilities of its human participants.

## REFERENCES

- Clancey, W. J. (1997). *Situated cognition: On human knowledge and computer representations*. Cambridge, UK: Cambridge University Press.
- Defense Science Board. (2005). *Defense Science Board task force on Patriot system performance*. Washington, DC: Office of the Undersecretary of Defense for Acquisition, Technology, and Logistics.
- Edelman, G. M. (1987). *Neural Darwinism: The theory of neuronal group Selection*. New York, NY: Basic Books.
- Freeman, W. J. (1995). *Societies of brains: A study in the neuroscience of love and hate*. Mahwah, NJ: Lawrence Erlbaum Associates, Inc.
- Haugeland, J. (1979). Understanding natural language. *The Journal of Philosophy*, 76(11), 619–632.
- Hollan, J., Hutchins, E., & Kirsh, D. (2000). Distributed cognition: Toward a new foundation for human–computer interaction research. *ACM Transactions on Computer–Human Interaction*, 7(2), 174–196.
- Hutchins, E. (1995). *Cognition in the wild*. Cambridge, MA: MIT Press.
- Klein, G. (1999). *Sources of power: How people make decisions*. Cambridge, MA: MIT Press.
- Reed, E. S. (1996). *Encountering the world: Toward an ecological psychology*. New York, NY: Oxford University Press.
- 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.

- Seng, Y. K., Wai, N. H., Chan, S., and Weng, L. K. (2009, July). *Automated metro—Ensuring safety and reliability with minimum human intervention*. Paper presented at the Nineteenth Annual International Symposium of the International Council on Systems Engineering (INCOSE), San Diego, CA.
- Simon, H. A. (1981). *The sciences of the artificial* (2nd ed.). Cambridge, MA: MIT Press.

Manuscript first received: September 2010