

Cognition, Distributed

Like all other branches of cognitive science, distributed cognition seeks to understand the organization of cognitive systems. Like most of cognitive science, it takes cognitive processes to be those that are involved in memory, decision making, inference, reasoning, learning, and so on. Also following mainstream cognitive science, it characterizes cognitive processes in terms of the propagation and transformation of representations.

What distinguishes distributed cognition from other approaches is the commitment to two related theoretical principles. The first concerns the boundaries of the unit of analysis for cognition. While boundaries are often a matter of tradition in a field, there are some general rules one can follow. Bateson (1972) says one should bound the unit so that things are not left inexplicable. This usually means putting boundaries on units where the traffic is low. The second principle concerns the range of mechanisms that may be assumed to participate in cognitive processes. While mainstream cognitive science looks for cognitive events in the manipulation of symbols (Newell et al. 1989), or more recently, patterns of activation across arrays of processing units (Rumelhart et al. 1986, McClelland et al. 1986) inside individual actors, distributed cognition looks for a broader class of cognitive events and does not expect all such events to be encompassed by the skin or skull of an individual.

When one applies these principles to the observation of human activity 'in the wild,' at least three interesting kinds of distribution of cognitive process become apparent: cognitive processes may be distributed across the members of a social group, cognitive processes may be distributed in the sense that the operation of the cognitive system involves coordination between internal and external (material or environmental) structure, and processes may be distributed through time in such a way that the products of earlier events can transform the nature of later events. The effects of these kinds of distribution of process are extremely important to an understanding of human cognition.

The roots of distributed cognition are deep, but the field came into being under its current name in the mid-1980s. In 1978, Vygotsky's *Mind in Society* was published in English. Minsky published his *Society of Mind* in 1985. At the same time, Parallel Distributed Processing was making a comeback as a model of cognition (Rumelhart et al. 1986). The nearly perfect mirror symmetry of the titles of Vygotsky's and Minsky's books suggests that something special might be happening in systems of distributed processing, whether the processors are neurons, connectionist nodes, areas of a brain, whole persons, groups of persons, or groups of groups of persons.

1. *Mind in Society*

For many people, distributed cognition means cognitive processes that are distributed across the members of a social group (Salomon 1993). The fundamental question here is how the cognitive processes we normally associate with an individual mind can be implemented in a group of individuals? A wide range of disciplines in the social sciences has explored this question.

Treating memory as a socially distributed cognitive function has a long history in sociology and anthropology. Durkheim, and his students, especially Halbwachs (1925), maintained that memory could not even be coherently discussed as a property of an isolated individual. Roberts (1964) proposed that social organization could be read as a sort of architecture of cognition at the community level. He characterized the cognitive properties of a society (its memory capacity and ability to manage and retrieve information) by looking at what information there is, where it is located, and how it can move in a society. Schwartz (1978) proposed a distributional model of culture that emphasized the distribution of beliefs across the members of a society. Romney et al. (1986) created quantitative models of the patterns of cultural consensus. The identification of patterns raised the question of why such patterns form. Sperber (1985) introduced the idea of an epidemiology of representations. He suggested an analogy in which anthropology is to psychology as epidemiology is to pathology. In the same way that epidemiology addresses the distribution of pathogens in a population, anthropology should treat questions about the distribution of representations in a community. A similar set of developments followed from Dawkins' (1976) discussion of 'memes' as the cultural analog of genes. These ideas have now coalesced in the field of memetics (Blackmore 1999). March and Simon (1958) argued that organizations can be understood as adaptive cognitive systems. Juries are an important class of distributed problem solving organization and they have been intensely studied by social psychologists (Hastie 1993). Of course, in social psychology there is a vast literature on small-group decision making, some of which discusses the properties of aggregates.

Scientific communities have received special attention because the work of science is fundamentally cognitive and distributed. The phenomena that have been explored include how the organization of communication media in a scientific community affect the kinds of things the community can learn (Thagard 1993), how conditions external to the individual scientists can affect their individual choices in ways that lead to different high-level structures to emerge (Kitcher 1990), how the distribution of cognitive activity within social networks and between people and inscriptions accounts for much of the work of science (Latour 1987), and how scientific facts are

created by communities in a process that simply could not be fit into the mind of an individual (Fleck 1935).

Economists have been interested in the tension between what is individually rational and what is rational at the aggregate level. This theme has been explored in game theory under the rubric of the Prisoner's Dilemma, the paradox of the commons, and other cases where individual rationality and group rationality diverge (Von Neumann and Morgenstern 1964).

Anthropologists and sociologists studying knowledge and memory, social psychologists studying small-group problem solving and jury decision making, organizational scientists studying organizational learning, philosophers of science studying discovery processes, and economists and political scientists exploring the relations of individual and group rationality, all have taken stances that lead them to a consideration of the cognitive properties of societies of individuals. There is ample evidence that the cognitive properties of a group can differ from the cognitive properties of the members of the group.

2. *The Society of Mind*

The work described above looks for mind-like properties in social groups. This is the *Mind in Society* reading. The metaphor can be run the other way as is done in Minsky's *Society of Mind* (1985). Rather than using the language of mind to describe what is happening in a social group, the language of social groups can be used to describe what is happening in a mind.

Minsky argued that to explain intelligence we need to consider a large system of experts or agencies that can be assembled together in various configurations to get things done. Minsky also allowed that a high-level agency itself could be composed of low-level agencies. With Papert (Minsky and Papert 1988), he argued that the low-level agencies (the ones that take on 'toy-sized problems') could be implemented as distributed computations in connectionist nets. Minsky said, '... each brain contains hundreds of different types of machines, interconnected in specific ways which predestine that brain to become a large, diverse society of partially specialized agencies' (1988). What this means of course is that the cognition of an individual is distributed cognition too.

A problem that remained unsolved by Minsky's work was 'how such systems could develop managers for deciding, in different circumstances, which of those diverse procedures to use' (1988). That is, how can the relations among the agencies get organized to perform new functional skills? To solve this problem, Minsky and Papert invoked biological maturation. An alternative way to approach this problem is to note that each 'society of mind' resides and develops in a community

of similar societies of mind. This means, of course, that both what's in the mind, and what the mind is in are societies. Getting internal agencies into coordination with external structure can provide the organization of the relations between the internal agencies that is required to perform the new functional skill.

Vygotsky developed this idea of the social origins of individual psychological functions in *Society of Mind* (Vygotsky 1978, Wertsch 1985). Vygotsky argued that every high-level cognitive function appears twice: first as an interpsychological process and only later as an intrapsychological process. The new functional system inside the child is brought into existence in the interaction of the child with others (typically adults) and with artifacts. As a consequence of the experience of interactions with others, the child eventually may become able to create the functional system in the absence of the others. This could be seen in Minsky's terms as a mechanism for the propagation of a functional skill from one society of mind to another. From the perspective of distributed cognition, this sort of individual learning is seen as the propagation of a particular sort of pattern through a community. Cultural practices assemble agencies into working assemblages and put the assemblages to work. Some of these assemblages may be entirely contained in an individual, and some may span several individuals and material artifacts. The patterns of activity that are repeatedly created in cultural practices may lead to the consolidation of functional assemblages, the atrophy of agencies that are rarely used, and the hypertrophy of agencies that are frequently employed. The result can be individual learning or organizational learning, or both.

2.1 *Interaction as a Source of Novel Structure*

An important property of aggregate systems is that they may give rise to forms of organization that cannot develop in the component parts. Freyd (1983) argued that some of the features of language that are identified as linguistic universals could arise out of the necessity of sharing the linguistic code. For instance, the reason that linguistic categories tend to approximate discrete structures may have little to do with the organization of the brain, and everything to do with the problem of pushing a complex representation through a very narrow channel. As Minsky and Papert point out, symbols can be expected to arise where there are bottlenecks in communication. That means we should look for the origins of symbols at places where the information 'traffic' is relatively low—or at the boundaries of our various units of analysis.

The phenomena related to the social distribution of cognition are most often investigated using ethnographic methods. In some cases, however, simulation models may be used to test hypotheses about the behavior of such distributed systems. For example, Hutchins and Hazlehurst (1995) explored Freyd's

ideas in a series of simulation models in which individuals (modeled by connectionist networks) interact with one another. They developed a robust procedure in which a shared lexicon emerges from the interactions of individuals. Hazlehurst and Hutchins (1998) demonstrated the emergence of reduced conventional sequences of lexical items—which they take to be the beginnings of syntax. These conventional sequences arise only in the condition of negotiated learning where the representing structures must simultaneously come to accurately represent the world and be shared among individuals, that is, be able to pass the communication bottleneck between individuals. Representations that are learned inside an individual, without the requirement of sharing them with others, come to represent the world, but do not show the reduced conventional code aspects that are the hallmarks of language and syntax.

By simultaneously considering the society of mind and mind in society, the distributed cognition approach provides a new place to look for the origins of complexity. Phenomena that are not predictable from the organization of any individual taken in isolation may arise in the interactions among individuals. Once having developed in this larger system, they may become elements of cultural practices and thereby become available for appropriation by individuals. This sort of scheme may be a partial solution to the paradox of how simple systems can lead to more complex ones.

3. *The Material Environment*

A second major thread in the fabric of distributed cognition concerns the role of the material environment in cognitive activity. Again, the question of where to bound the unit of analysis arises. The potential of the material environment to support memory is very widely recognized. But the environment can be more than a memory. Cognitive activity is sometimes situated in the material world in such a way that the environment is a computational medium.

Cognitive artifacts are the *Things that Make Us Smart* in the title of Norman's (1993) book. The notion that cognitive artifacts amplify the cognition of the artifact user is fairly commonplace. If one focuses on the products of cognitive activity, cognitive artifacts do seem to amplify human abilities. A calculator seems to amplify one's ability to do arithmetic, writing down something one wants to remember seems to amplify one's memory. Cole and Griffin (1980) point out that this is not quite correct. When I remember something by writing it down and reading it later, my memory has not been amplified. Rather, I am using a different set of functional skills to do the memory task. Cognitive artifacts are involved in a process of organizing functional skills into cognitive functional systems.

Consider an example from the world of ship navigation (Hutchins 1995). Navigators frequently face the problem of computing the ship's speed from distance traveled over a given period of time. If a ship travels 1,500 yards in 3 minutes, what is the speed of the ship in knots? There are many ways to solve this problem. Most readers of this article would probably attempt to use a paper and pencil plus their knowledge of algebra to solve it. That procedure is effective, but not nearly as efficient as the 'so-called' 3-minute rule. An experienced navigator need only see the problem stated to see that the answer is 15 knots. The speed in knots equals the number of hundreds of yards covered in 3 minutes. The use of this rule is a case of situated seeing. The rule itself is an internal cognitive artifact. But suppose the ship covered 4,000 yards in 7 minutes? For that problem a material artifact called the three-scale nomogram is more appropriate. A nomogram has three logarithmic scales: one each for distance, time, and speed. If the values of any two variables in a distance/rate/time problem are known, the other can be determined by laying a straight edge on the nomogram so that it touches the known values. The straight edge will touch the third scale at the value of the answer. It is clear that cognitive work is being done, but it is also clear that the processes inside the person are not, by themselves, sufficient to accomplish the computation. A larger unit of analysis must be considered. The skills of scale reading and interpolation are coordinated with the manipulation of objects to establish a particular state of coordination between the straight edge and the nomogram. This is a very different set of agencies than was involved in doing the problem via algebra and paper and pencil. In fact, the skills that are needed to use the nomogram are the things that humans are good at: pattern matching, manipulation of objects in the world, and mental simulation of simple dynamics (Rumelhart et al. 1986).

A computation was performed via the manipulation of a straight edge and nomogram. The nomogram was designed in such a way that the errors that were possible in algebra are impossible when using the nomogram. It is essential to distinguish the cognitive properties required to manipulate the artifact from the computation that is achieved via the manipulation of the artifact. This is a key point, and the failure to see it clearly has been the source of many difficulties in cognitive science.

4. *Distributing Cognition in Time*

Simon (1998) offered a parable as a way of emphasizing the importance of the environment for cognition. He argued that, as we watch the complicated movements of an ant on a beach, we might be tempted to attribute to the ant some complicated program for constructing the path taken. In fact, Simon says, that trajectory tells us more about the beach than about the

ant. Similarly, in watching people thinking in everyday settings, we may be learning as much about their environment for thinking as about what is inside them. The environments of human thinking are not 'natural' environments. They are artificial through and through. They develop over time. The crystallization of partial solutions to frequently encountered problems in artifacts such as the 3-minute rule and the nomogram is a ubiquitous strategy for the stabilization of knowledge and practice. Humans create their cognitive powers in part by creating the environments in which they exercise those powers.

5. Conclusion

It does not seem possible to account for the cognitive accomplishments of our species by reference to what is inside our heads alone. One must also consider the cognitive roles of the social and material world. But, how shall we understand the relationships of the social and the material to cognitive processes that take place inside individual human actors? This is the problem that distributed cognition attempts to solve.

According to Gardner (1985), a more or less explicit decision was made in cognitive science to leave culture, context, history, and emotion out of the early work. These were recognized as important phenomena, but their inclusion made the problem of understanding cognition too complex. The 'Classical' vision of cognition that emerged was built from the inside out, starting with the idea that the mind is a central logic engine. From that starting point, it followed that memory could be seen as retrieval from a stored symbolic database, that problem solving is a form of logical inference, that the environment is a problem domain, and that the body is an input device (Clark 1996). Attempts to reintegrate culture, context, and history into this model of cognition have proved very frustrating. The distributed cognition perspective aspires to rebuild cognitive science from the outside in, beginning with the social and material organization of cognitive activity.

See also: Cognitive Science: Overview; Situated Cognition: Origins

Bibliography

- Bateson G 1972 *Steps to an Ecology of Mind*. Balantine Books, Location
 Blackmore S 1999 *The Meme Machine*. Oxford University Press, Oxford
 Clark A 1996 *Being There: Putting Brain, Body and World Together Again*. MIT Press, Cambridge, MA
 Cole M, Griffin M 1980 Cultural amplifiers reconsidered. In: Olson D (ed.) *The Social Foundations of Language and Thought*. Norton, London

- Dawkins R 1976 *The Selfish Gene*. Oxford University Press, Oxford
 Fleck L 1979 *The Genesis and Development of a Scientific Fact*. University of Chicago Press, Chicago
 Freyd J 1983 Shareability: the social psychology of epistemology. *Cognitive Science* 7: 191–220
 Gardner H 1985 *The Mind's New Science*. Basic Books
 Halbwachs M 1925 *Les Cadres Sociaux de la Mémoire*. Librairie Felix Alcan, Paris
 Hastie R 1993 *Inside the Juror: The Psychology of Juror Decision Making*. Cambridge University Press, Cambridge, UK
 Hazlehurst B, Hutchins E 1998 The emergence of propositions from the coordination of talk and action in a shared word. *Language and Cognitive Process* (special issue on Connections Approaches to Language Development edited by Kim Plunkett). 2 & 3 April/May
 Hutchins E 1995 *Cognition in the Wild*. MIT Press, Cambridge, MA
 Hutchins E, Hazlehurst B 1995 How to invent a lexicon: the development of shared symbols in interaction. In: Gilbert N, Conte R (eds.) *Artificial Societies: The Computer Simulation of Social Life*. UCL Press, London
 Kitcher P 1990 The distribution of cognitive labor. *The Journal of Philosophy* 87(1): 5–22
 Latour B 1987 *Science in Action*. Harvard University Press, Cambridge, MA
 March J, Simon H 1958 *Organizations*. Wiley, London
 McClelland J, Rumerlhart D, PDP research group 1986 *Parallel Distributed Processing: Explorations in the Microstructure of Cognition. Vol 2: Psychological and Biological Models*. MIT Press, Cambridge, MA
 Minsky M 1985 *Society of Mind*. Simon and Schuster, Hemel Hempstead, UK
 Minsky M, Papert S 1988 *Perceptrons*. MIT Press, Cambridge, MA
 Newell A, Rosenbloom P, Laird J 1989 Symbolic architectures for cognition. In: Posner M (ed.) *Foundations of Cognitive Science*. MIT Press, Cambridge, MA
 Norman D 1993 *Things that Make Us Smart*. Addison Wesley, Location
 Roberts J 1964 The self-management of cultures. In: Good-enough V (ed.) *Explorations in Cultural Anthropology: Essays in honor of George Peter Murdock*. McGraw-Hill
 Romney A K, Weller S, Batchelder W 1986 Culture consensus: A theory of culture and informant accuracy. *American Anthropologist* 88(2): 313–38
 Rumelhart D, McClelland J, PDP research group 1986 *Parallel Distributed Processing: explorations in the Microstructure of Cognition. Vol 1: Foundations*. MIT Press, Cambridge, MA
 Rumelhart D, Smolensky P, McClelland J, Hinton C 1986 Schemata and sequential thought processes in PDP models. In: McClelland J, Rumelhart D, PDP research group (eds.) *Parallel Distributed Processing: Explorations in the Microstructure of Cognition. Vol 2: Psychological and Biological Models*. MIT Press, Cambridge, MA
 Salomon G 1993 *Distributed Cognitions*. MIT Press, Cambridge, MA
 Simon H 1998 *The Sciences of the Artificial* 3rd edn. MIT Press, Cambridge, MA
 Schwartz T 1978 The size and shape of a culture. In: Barth F (ed.) *Scale and Social Organization*. Universitetsforlaget
 Sperber D 1985 Anthropology and psychology: towards an epidemiology of representations. *Man* 20: 73–89

- Thagard P 1993 Societies of minds: Science as distributed computing. *Studies in History and Philosophy of Science* **24**: 49–67
- Von Neumann J, Morgenstern O 1964 *Theory of Games and Economic Behavior*. Wiley, Chichester
- Vygotsky L 1978 *Mind in Society: The Development of Higher Psychological Processes*. Harvard University Press, Cambridge, MA
- Wertsch J 1985 *Vygotsky and the Social Formation of Mind*. Harvard University Press, Cambridge, MA

E. Hutchins

Cognitive Aging

Cognitive aging is concerned with age-related changes in adulthood in the basic processes of learning and memory, as well as the complex higher order processes of language and intellectual competence or executive functioning. Although most of the literature has been concerned with explaining the mechanism of cognitive decline, there is also a substantial interest in issues such as compensation and the role of external support, including collaborative problem solving.

1. Definition of Cognitive Aging

There have been two distinct traditions in the study of cognitive aging. The first grew out of experimental child psychology while the second derives from psychometric roots.

1.1 Experimental Study of Memory Functions and Language

The concern in this literature is to explicate possible causal variables that would explain why many adults suffer memory loss and decline in the complex manipulation of language variables such as text processing. The typical approach here is to design experiments testing for the effects of single variables in carefully controlled laboratory settings requiring only limited numbers of subjects. Because there is often little interest in individual differences, or population parameters, study participants are typically drawn from convenience samples (McKay and Abrams 1996).

1.2 Descriptive Study of Adult Intellectual Development

Descriptive studies of adult intellectual development often stem from the longitudinal follow up of samples first assessed in childhood or adolescence. Other such studies may represent carefully stratified samples from

defined populations, first assessed at a particular life stage, whether in early adulthood or in early old age. Although descriptive studies often begin as cross-sectional inquiries, they are most frequently conducted as longitudinal analyses since the interest is often in individual differences in intraindividual change, or in the elucidation of typologies of individuals who follow different growth trajectories. These are frequently large-sample studies, and the use of correlational or quasi-experimental approaches is typical (Baltes et al. 1999, Schaie 1996b).

2. Methodological Issues

2.1 Age-comparative vs. Age Change Designs

Much of the experimental cognitive aging literature is based on age-comparative studies, which typically contrast a group of young adults (typically college students) with convenience samples of community-dwelling older adults in their sixties and seventies. It should be recognized that such comparisons are fraught by the problem that it is often unreasonable to assume that the two age groups can be adequately matched for other status variables that might provide a rival explanation for any observed age difference on the dependent variable. This creates particular problems for identifying the mechanisms that may be implicated in age-related decline from young adulthood into old age. Age-comparative designs are also inadequate in explaining individual differences in age changes. The latter can only be investigated by means of longitudinal paradigms (Schaie 1965).

2.2 The Role of Response Speed

A number of theorists have argued that changes in the central nervous system are the primary common cause for the observed age-related declines in cognitive performance. In fact, there have been many published analyses that show a substantial reduction in age differences, if some measure or measures of reaction time or perceptual speed is partialled out of the relation between a given cognitive process and chronological age (Salthouse 1999). This issue is of particular concern because it is not clear whether the observed average increase in reaction time (generally assumed to be of the magnitude of approximately 1.6 from the early twenties to the late sixties), while reliably demonstrable in the laboratory, is of significance in many or most tasks of daily living.

3. Basic Findings from the Experimental Literature on Cognitive Aging

Most of this literature is cross-sectional in nature and usually consists of a comparison of convenience samples of young adults (often sophomore psychology

2072