

CONCEPTS OF COGNITION AND CONSCIOUSNESS: FOUR VOICES

Bonnie A. Nardi

Apple Research Laboratories
Apple Computer
1 Infinite Loop
Cupertino, CA 95014
phone: +1-408-974-8708 fax: +1-408-974-8414
http://www.atg.apple.com/personal/Bonnie_Nardi/

ABSTRACT

This paper considers theories of cognition and consciousness in four traditions: neuroscience, cognitive science, activity theory and the distributed cognition approach. It is most concerned with social theories of consciousness—activity theory and distributed cognition—but briefly considers biological and computational models as a foil or backdrop against which the social theories stand out more clearly.

INTRODUCTION

Of all the many artifacts we've produced in our two million years on the planet, computers are the most human-like. Our encounter with computers has given a new edge to ancient philosophical questions about what it means to be human. While taking up such questions might seem a long way from information systems design and evaluation, as an anthropologist and designer, I believe that the way we describe what it means to be human is ultimately reflected in our designs. Thus I have found it important and useful to try to keep up with the major attempts to define who we humans might really be. While there are a great many such attempts, I would like to analyze four powerful paradigms whose vocabulary and concepts permeate today's discourse: neuroscience, cognitive science, activity theory and the distributed cognition approach. Philosophers also have much to say about cognition and consciousness, but I'll concentrate here on the scientific approaches. I will not deal with the practical matter of exactly how to apply these theories to design and evaluation (but see Nardi, 1996 for the application of activity theory to problems of human-computer interaction); rather I hope to raise consciousness about the implications of the concepts and vocabulary that become part of the way we talk about the people we study or design for, whether in a research or applied context.

Two concepts that capture modern cutting edge efforts to describe the nature of humankind are *cognition* and *consciousness*. Neuroscience, cognitive science, artificial intelligence, psychology, philosophy and anthropology have all taken as one of their greatest challenges explaining consciousness or cognition. The question asked in this paper is: What is the view of humanity that emerges in each paradigm as the difficult problems of consciousness and cognition are tackled? Reflecting on this question will, I hope, help to move our dialogue forward as we define who we are, so that we can reflect that definition in our designs. I offer these analyses of cognition and consciousness in the spirit of what Donald Schön calls the "reflective practitioner" (Schön, 1983)—they are meant to raise our level of awareness about the concepts we employ and the assumptions we make about people as we study and design.

CONSCIOUSNESS, PART I

Neuroscience has lately set itself the task of explaining consciousness, marshaling its resources to try to understand "those states of sentience and awareness" that are an "inner, first-person" phenomenon, as John Searle defines it in his review of research on consciousness (Searle, 1996). According to this view, the seat of consciousness is the brain, and the way to understand consciousness is to understand how the brain works. Consciousness is a fundamentally biological phenomenon; its secrets will yield to biological analyses. As Searle puts it, "It is an amazing fact that everything in our conscious life, from feeling pains, tickles, and itches to—pick your favorite—feeling the angst of postindustrial man under late capitalism or experiencing the ecstasy of skiing in deep powder—is caused by brain processes." Francis Crick (1993) tells us that our joys, sorrows, memories, ambitions, identity and free will are "no more than the behavior of a vast assembly of nerve cells and their associated molecules."

Consciousness is as biological as "growth, digestion, or the secretion of bile," deadpans Searle. A very strong claim is made that *the brain causes consciousness*.

Under a biological framework, the problems of understanding consciousness lie in the complexity of the brain itself (those mind-boggling 100 billion neurons and their complicated connections), the moral impossibility of controlled experiments and the difficulty of talking rationally about what it feels like to be pinched (or to ski in deep powder, or whatever). The task before neuroscience is to find a brain architecture that explains how

something that fits in the palm of the hand can ski and remember and plot and plan and feel sad and try to come up with architectures to describe itself. Crick, for example, gives a go at sketching an architecture by describing the brain as possessed of "emergent properties" that "arise in the brain from the interactions of its many parts." Gerald Edelman (1992) is more specific in detailing a brain that actively recategorizes memories, learns by valuing some stimuli over others, discriminates self from others, categorizes temporal events, forms concepts, and provides reentrant connections between memory and perceptual categorization.

In all of these architectures, the brain and its physical structures are the locus of study, the home of consciousness. As Edelman (1992) says, "...mind depend[s] on special arrangements of matter." "We need not reach beyond biology itself to mount any exotic explanations of the mind" (Edelman, 1992). In the neuroscientific view, the problems of consciousness are tidily contained beneath the skull, in one known place, with physical structures that can be unambiguously defined and studied. Sure the brain is complex, but at least we know where to look.

COGNITION, PART I

The cognitive scientists, by contrast, have chosen to concentrate on something far more abstract than neurons and synapses. They take up not consciousness (which would sound a bit mystical to most of them), but cognition, defining symbolic representations as the mainstay of cognition. No secretions of bile here. For cognitive science, cognition is, to give a talk show definition, all the neat mental stuff that people can do. Problem solving is a big favorite of the cognitivists (and they love games and logic puzzles, the kind of thing the math whizzes were always best at), but they embrace any kind of mental activity as their turf: classification, remembering, decision making, judging, calculating and so forth. Their game is to spell out how people represent and manipulate representations. This clean disembodied activity can be simulated on computers, and showing that a program can solve a problem is an important activity for cognitivists. They don't worry too much about consciousness per se because their heart belongs to the more precise image of the person as rational problem solver, systematically crunching away at those representations and "mental models." A focus on consciousness, rather than cognition, forces one to the ambitions and sorrows and free will and skiing in deep powder that a Francis Crick or Gerald Edelman will take a crack at, but these squishy things quickly prove intractable in a representations-centric perspective. The cognitivists' strategy has been to first figure out human problem solving capabilities (as they define them), and then to someday move on to the rest of life, after the cognitive ship has come in.

Like the neuroscientists, the cognitive scientists confine their attention to what's under the hood, that is, to what goes on inside the head of a single individual—who is supposedly the locus of cognition. They firmly locate cognition "beneath the skin" as Lave (1988) puts it, and there is a neat input-output loop between external stimuli and the resulting representations that the individual creates, stores and manipulates. The body, other people and things like tools are not especially interesting for a cognitivist account, except as "stimuli." The theory of cognition provided by cognitive science is profoundly Cartesian, separating "mind" from everything else.

CONSCIOUSNESS, PART II

What if there were a theory of mind that was *not* Cartesian, that *did* take account of the role of other people and tools and the environment in which cognition takes place? What if there were a *social theory of consciousness* in which consciousness was seen as inextricably woven together with the social, the social being defined as interactions with other real people, as well as interactions with the tools other real people have designed and left for others as part of their culture? Such a theory would most certainly look beyond "biology itself" to understand consciousness and would not regard thinking about other people and tools as "exotic," as our neuroscience friends would have it.

There is such a theory, and it is called activity theory. I'll sketch out its five main principles and then move to a discussion of another more recent theory—"distributed cognition"—that shares some insights with activity theory but which does not, in my opinion, capture its depth or richness. The reader is likely to run across the recent work of Ed Hutchins or Jean Lave and perhaps less likely to find Vygotsky or Leont'ev, and the other activity theorists, so I'll make the pitch here that Vygotsky and Leont'ev and their colleagues have already covered most of the ground today's distributed cognition theorists have, and in fact gone well past it. The vagaries of Russian history, the home of activity theory, have meant that activity theory has been slow to reach the West.

For those who do know this work, I'll note that I use the term "activity theory" broadly to include the cultural-historical school founded by Vygotsky in the 1920's, as well as the activity theory work of Luria and Leont'ev

and their students, and the continuing work in these traditions in many countries. Key activity theorists include the Russian psychologists Vygotsky (1978, 1934/1986), Rubinshtein (1940), and Leont'ev (1978). There is a thriving activity theory tradition with computer-science related work in Scandinavia (Bertelsen, 1994; Bødker, 1989, 1991; Christiansen, 1996; Engeström and Escalante, 1996; Kaptelinin, 1993, 1996; Kuutti, 1991, 1996; Saarelma, 1993; Sjöberg, 1996), as well as increasing interest in activity theory applied to information systems the U.S. (Bellamy, 1996; Blumenthal, 1995; 1996; Holland and Reeves, 1996; Nardi, 1996a), Europe (Bannon, 1990; Bannon and Bødker, 1991; Draper, 1993; Imaz and Benyon, 1996; Raeithel, 1992; Raeithel and Velichkovsky, 1996), Canada (Cohen, Candland and Lee, 1995), and Australia (Bourke, Verenikina, and Gould, 1993), and also continuing work in Russia (Tikhomirov, 1991; Zinchenko and Munipov, 1989; Zinchenko, 1996).

A key point I wish to alert the reader to is that activity theory insists that consciousness is *social*. Consciousness can never be reduced to brain mechanisms or representations inside one person's head, because without other people and artifacts, *there is no consciousness*. Other people and artifacts are not just "stimuli" as cognitive science would have it, nor are they epiphenomenal to the firing of synapses; rather they are *crucial* to any kind of consciousness. If it weren't for a social world, there would be no you. Your "mind" does not belong strictly to you; it is not really located under your skull; rather it functions only so long as you are in a social matrix in which it can find expression. Remember the old joke about the person who answers the door and says, "There's no one home." They really mean it; that's why it's funny. According to activity theory, the mind is like that: without the rest of the family or roommates around, there is literally "no one home," no consciousness. I may be sitting here by myself in California struggling to explain activity theory, but this would be impossible if I did not have some notion of you, dear readers, in my consciousness to motivate and channel my activity. My mind is really stretched across the ocean, all the way to Australia, in an absolutely real sense.

Activity theory does not in any way deny the importance of the brain or its structures. In fact a number of eminent activity theorists such as Luria were neuropsychologists. Activity theory sees the brain as "the material basis for all complex psychological processes" (Luria, 1972). Activity theory insists however that the brain does not *cause* consciousness; rather it provides the physical matrix in which it can develop and change. And, most crucially, *we must look beyond the brain itself to understand consciousness*.

One more point before we plunge in to the principles of activity theory: activity theory says *you are what you do*. If you spend your days digging ditches, your consciousness will be thusly shaped. If you find yourself writing computer programs as your life's work, that is quite another kind of consciousness. "Doing" is very broadly defined in activity theory to include things like speaking, meditating, remembering, as well as activities more centered in the body and its movements. This flattening of "doing" to include all kinds of things is disconcerting at first, but it scores big points in helping to make us aware that our consciousness will be shaped just as much by praying as by bungee jumping. It helps to reduce the seduction of Cartesianism which wants to yank the mind away from everyday life and set up a life of its own.

Activity theory posits that you are a unique individual because you've had and chosen for yourself a unique set of "doings" in your life. Saying you are what you do is a lot different than saying you are your neurons, or you are the mental representations that sit under your skull. Saying you are what you do includes the way you physically move through the world, and the tools you pick up and put in your hands and use, and the hugs and kisses you give and get and everything that makes you human. Activity theory asserts that the consciousness created by your everyday activities spills over into other seemingly unrelated activities, so that activities like formal education and reading have tremendously powerful effects on consciousness.

Now let's look systematically at the five principles of activity theory: hierarchical structure of activity, object-orientedness, internalization/externalization, tool mediation, and development.

Hierarchical Structure of Activity

In activity theory the unit of analysis is an activity. Leont'ev, one of the chief architects of activity theory, described an activity as being composed of subject, object, actions, and operations (1974). A subject is a person or a group engaged in an activity. An object (in the sense of "objective" or motive) is held by the subject, and motivates activity, giving it a specific direction: "... behind the object there always stands a need or a desire, to which [the activity] always answers" (Leont'ev, 1974). The term "objectified motive" is a useful mnemonic. One might also think of the "object of the game" or an "object lesson."

Actions are goal-directed processes that must be undertaken to fulfill the object. They are conscious (because one holds a goal in mind), and different actions may be undertaken to meet the same goal. For example, "A person may have the object of obtaining food, but to do so he must carry out actions not immediately directed

at obtaining food... His goal may be to make a hunting weapon. Does he subsequently use the weapon he made, or does he pass it on to someone else and receive a portion of the total catch? In both cases, that which energizes his activity and that to which his action is directed do not coincide" (Leont'ev, 1974). Goals can have lower level goals, which can have lower level goals, and so forth, much like the concept of goals/subgoals in AI and other traditions. For example, making the hunting weapon is an action which then entails finding suitable materials and tools for the manufacture of the weapon, and so on.

Objects can be transformed in the course of an activity; they are not rigid structures. As Kuutti (1996) noted, "It is possible that an object itself will undergo changes during the process of an activity." Objects do not, however, change on a moment-by-moment basis (see Holland and Reeves, 1996).

Actions are similar to what are often referred to in the human computer interaction literature as tasks (e.g., Norman, 1991). Activities may overlap in that different subjects engaged together in a set of coordinated actions may have multiple or conflicting objects (Kuutti, 1991).

Moving down the hierarchy of actions we cross the border between conscious and automatic processes. Actions have operational aspects, that is, the way the action is actually carried out. Operations become routinized and unconscious with practice. Operations do not have their own goals; rather they provide an adjustment of actions to current situations. When learning to drive a car, the shifting of the gears is an action with an explicit goal which must be consciously attended to. Later, shifting gears becomes operational, and "can no longer be picked out as a special goal-directed process: its goal is not picked out and discerned by the driver" (Leont'ev, 1974). Operations depend on the conditions under which the action is being carried out. If a goal remains the same while the conditions under which it is to be carried out change, then "only the operational structure of the action will be changed" (Leont'ev, 1974).

Activity theory holds that the constituents of activity are not fixed, but can dynamically change as conditions change. This is an important distinction between activity theory and cognitive science-based techniques such as GOMS (Card, Moran and Newell, 1983). In activity theory, all levels can move both up and down (Leont'ev, 1974). As we saw with gear-shifting, actions become operations as the driver habituates to them. An operation can become an action when "conditions impede an action's execution through previously formed operations" (Leont'ev, 1974). For example, if one's mail program ceases to work, one continues to send mail by substituting another mailer, but it is now necessary to pay conscious attention to using an unfamiliar set of commands. Notice that here the object remains fixed, but goals, actions and operations change as conditions change.

Object-orientedness

The principle of "object-orientedness" (not to be confused in any way with object-oriented programming) is one of the most important principles of activity theory. Every motive is an object (as defined above) but there is another related sense of the word object, i.e., a prospective outcome toward which activity is directed, around which activity is coordinated, and which will be crystallized in a final form when the activity is complete. For example, a computer program is an object of a programmer's activity. It is not her "motive"; the software program is that toward which she directs her activity so she can attain a motive, say, becoming the best programmer in California. Line by line she writes the code; the program itself channels her attention and thought and motion.

Objects can be physical things (such as the bull's eye on a target) or ideal objects ("I want to become a brain surgeon."). Leont'ev noted that the concept of object is not limited to physical, chemical, and biological properties; it crucially includes social properties as well.

Internalization/externalization

Activity theory differentiates between internal and external activities. The traditional notion of mental processes in cognitive science corresponds to internal activities. Activity theory emphasizes that internal activities cannot be understood if they are analyzed separately, in isolation from external activities, because there are mutual transformations between these two kinds of activities. Internalization is the transformation of external activities into internal ones. Activity theory argues that it is not just mental representations that get placed in someone's head; it is the holistic activity including motor activity and the use of artifacts that are crucial for internalization. For example, learning to calculate may involve counting on the fingers, in the early stages of learning simple arithmetic. Once the arithmetic is internalized, the calculations can be performed in the head without external aids.

Internalization provides a means for people to try potential interactions with reality without performing actual manipulation with real objects (mental simulations, imaginings, considering alternative plans, etc.). In some cases external components can be omitted in order to make an action more efficient, e.g., in the case of

performing calculations in the head. Internalization can help to identify an optimal way to perform action before performing the action externally.

Externalization transforms internal activities into external ones. Externalization is often necessary when an internalized action needs to be "repaired," or scaled, e.g., when a calculation is not coming out right when done mentally, or is too large to perform without pencil and paper or calculator (or some external artifact). Externalization is also important when a collaboration between several people requires their activities to be performed externally in order to be coordinated. While the concept of internalization shares much with traditional cognitive science's notions of information processing, externalization is not emphasized in cognitive science. Furthermore, activity theory says that it is the *constant transformation between external and internal* that is the very basis of human cognition and activity.

Mediation

Activity theory's emphasis on social factors and on the interaction between people and their environments explains why the principle of tool mediation plays a central role within the approach. First of all, tools shape the way human beings interact with reality. And, according to the principle of internalization/externalization, shaping external activities results in shaping internal ones. Second, tools reflect the experience of other people who encountered and solved similar problems and invented or modified a tool to make it effective and efficient. This experience is accumulated in the structural properties of tools (their shape, size, material, and so forth) as well as in the knowledge of how the tool should be used. The use of tools constitutes an accumulation and transmission of social knowledge. Tools influence the nature not only of external behavior but also internal mental functioning.

Vygotsky identified two kinds of tools: technical tools and psychological tools. Technical tools manipulate physical objects (e.g., a hammer) while psychological tools are used to influence other people or oneself (e.g., the multiplication table, a calendar, or an advertisement).

Development

Finally, activity theory requires that activity be analyzed in the context of development. Development continuously reforms and develops practice.

In activity theory development is not only an object of study, it is also a general research methodology. When activity theorists do experiments, they always hang around long enough to see what happens as the person gets used to the tool being tested, or the material being internalized. For example, in a simple but classic study, Kaptelinin (1993) investigated how ordinary computer users adapted to two different forms of a menu. Users were first presented with a menu with all menu items spelled out fully and given in the same order on repeated trials. Then one group of users got the menu with the items fully spelled out but in scrambled order on successive trials. The second group got a menu with dots for the menu items but with the items always in the same order. The second menu was harder to use at first, but proved to be faster and easier once people got used to it (they could look up the bindings till they learned them). A typical cognitive science analysis might have stopped the experiment prematurely, or worse, counted on a theory of visual "affordances" (Norman, 1991) for which the first menu clearly was superior (i.e., because you could identify the actual items and didn't have to take time to learn them; they could be directly "picked up" from the environment as Gibson (1979) would have said). But developmentally, users were able to process the ordered dotted items better over time.

This experiment also shows the relationship of the brain to real activity: clearly there is something about the brain that does linear processing efficiently, *once a person has had the experience of a particular activity*; here using the menu over time. Activity theory is quite happy to look at the wetware aspects of the brain as neuroscience does, but it insists on viewing those aspects *in the context of some real activity*, and *developmentally, over time*, and not as a question reduced to nerve cells and molecules in the way Crick talks about it.

Integration of the Principles

These basic principles of activity theory should be considered as an integrated system. A systematic application of any of the principles makes it eventually necessary to engage all the other ones, just as we saw that mediation calls upon internalization/externalization. Activity theory insists on the unity of these principles and does not abstract out any single process (e.g., internalization), because then the whole activity—and consciousness itself—cannot be understood.

Summary

In activity theory then, what it means to have a human consciousness is to be a part of a web of social activities and to live and act in a culturally elaborated environment that is profoundly *artificial*, populated by a wealth of tools, including language. As Vygotsky noted,

The use of artificial means, the transition to mediated activity, fundamentally changes all psychological operations just as the use of tools limitlessly broadens the range of activities within which the new psychological functions may operate. In this context, we can use the term *higher* psychological function, or *higher behavior* as referring to the combination of tool and sign in psychological activity (1930/1978).

Vygotsky reckoned that people start out with a set of basic capabilities such as attention, will, intention, and then, through human activity, develop what he called the higher psychological functions. These functions include language, decision making, abstraction, generalization, classification, problem solving, and more advanced forms of will, attention, memory, intention. Vygotsky saw the potential for development as limitless. He very much believed in formal, especially scientific, education, as a way to create a person possessed of the higher mental functions (seeing education as a necessary but not sufficient condition for development).

Vygotsky's very definition of consciousness, then, is the active processes of the higher and lower psychological functions. Because these functions arise, develop and change within a social matrix (as detailed in the principles of mediation, internalization/externalization, etc.), they cannot be seen as residing strictly "under the skull"—consciousness is social, simultaneously beyond and within the individual. Because of externalization, there is not a one-way path from outer to inner activity; rather, inner activity gets concretized and made public through external activity.

The emphasis on the artificial pits activity theory directly against the neuroscientific view which is fundamentally biological. Edelman (1992) talks about "putting the mind back into nature," while activity theorists (in direct descent from Marx) see the mind as fundamentally the product of the social relations in which people spend their everyday lives. Activity theory wants to get the mind into *culture*, while neuroscience wants to nail it down in nature.

The neuroscientific view does not account for the role of artificial tools, including language, in the development of consciousness. (The ability to learn language obviously has a biological substrate, but words and the concepts they convey are patently manufactured, artificial things). While it might be argued that a neuroscientist would of course think that language is important to consciousness, there is in fact an unambiguous divide between the two approaches: a theory that is profoundly biological does not find causality in the artificial. The emphasis on biology leads one away from the artificial. One never gets from brain architecture as cause to activity as cause. Activity never gets into the discussion at all. I argue that it is necessary to adopt the activity theory strategy of acknowledging activity as the primary shaper of consciousness, and then to get to the brain and biology. Activity theory declares that consciousness depends *directly and profoundly* on the mediation provided by human activity. To say that consciousness is fundamentally about nerve cells and molecules would be impossible in activity theory.

A limitation of the activity theory approach is that while activity theorists posited the importance of the wider culture in shaping consciousness, their analyses in practice have not included culture in very interesting ways. This is an area for further development within activity theory. It is not an easy area as 150 years of anthropology has shown the difficulty of elaborating what culture is all about, and especially how culture affects the individual. Understanding culture remains a future challenge for activity theory.

Another limitation of activity theory is that it provides few tools for understanding social organization and its impact on activity. This is an area of active research in Europe and progress is being made (see especially Engeström, 1987 and Raiethel, 1992).

Activity theory is optimistic about the potential for human development. While recognizing that everyone starts out with some basic equipment such as the ability to pay attention, activity theory says that consciousness is the result of development. What you spend your time doing is what shapes your consciousness. Because you have some say over how you spend your time, you are in charge of your consciousness. If you design mediating tools for others (such as computer hardware or software), you are also responsible, in part, for the consciousness of others. Our tools make us who we are, says activity theory. As designers this gives us tremendous power and tremendous responsibility.

COGNITION, PART II

Activity theory has provided a social theory of consciousness. A much newer paradigm, "distributed cognition" has to some extent revisited much of what the activity theorists discovered and explicated. But there are some crucial differences in the two approaches that I will explore here. Distributed cognition has taken a sharp turn away from Vygotsky's notion of the individual as a being with unbounded potential, the notion of consciousness as capable of limitless development. Distributed cognition is much less optimistic concerning human development and sees tools as a means of reducing the complexity of human tasks.

In my opinion, Edwin Hutchins is the most articulate proponent of the distributed cognition approach, so I will consider his work in this brief paper rather than provide a full overview (but see also Halverson, 1993; Salomon, 1993; Sandberg and Wielinga, 1993; Rogalski, 1994; Rogers and Ellis, 1994; and Lave, 1988 though there are major differences in Lave's approach which I discussed in Nardi, 1996b). Hutchins' book *Cognition in the Wild* (1995) presents the most carefully worked out theory of distributed cognition with which I am familiar. While I provide a critique of distributed cognition, I acknowledge that it is a relatively new body of work that has not yet had time to ferment and mature. But its potential impact merits a close critical look.

The points of similarity between distributed cognition and activity theory are (1) cognition/consciousness are mediated by tools, so the artificial is profoundly important to any theory of what it means to be human; (2) cognition/consciousness are social in nature and (3) the "functional system" of a person or group with a set of tools is an important unit of analysis. Each approach differs in (1) what is meant by "cognition" (2) the attitude toward the potential for human development (3) the notion of objects (4) concepts of awareness.

Hutchins (1995) details the functioning of navigation aboard an amphibious helicopter transport in the U.S. Navy. He studied how the sailors on board the ship used tools and human communication to navigate the ship. *Cognition in the Wild* is a tour de force ethnography, with beautifully detailed descriptions of the minutiae of navigation and the best discussion of how people use tools that I have ever seen.

I will now look in more detail at the points of similarity between activity theory and distributed cognition. I will have to use the narrower "cognition" rather than "consciousness" here as that is what the distributed cognitivists deal with.

DISTRIBUTED COGNITION AND ACTIVITY THEORY: CONNECTIONS

Cognition is Mediated by Tools

Hutchins has a real feel for tools (technical, rather than psychological tools, to use Vygotsky's terms) and his descriptions of the use of nautical tools such as hoeys, alidades, fathometers, compasses, charts, and so forth are brilliant. Through these rich descriptions, Hutchins makes clear that cognition is mediated by tools. This finding is in complete concert with what Vygotsky proposed in 1930. Hutchins has done the ethnographic work to demonstrate—in the most vivid manner—the importance of tool mediation. Like the activity theorists, Hutchins understands that the crucial role of tool mediation in cognition means that cognition is embedded in the artificial. As Hutchins says, "the environments of human thinking...are artificial through and through" (p. xvi). (See also Leont'ev, 1981.)

Cognition is Social in Nature, Embedded in Practical Human Activity

Hutchins expends lavish ethnographic attention on the social arrangements by which the cognition involved in navigation gets done. He makes clear that cognition is a social affair involving nuanced communication, learning, interpersonal interaction. Again, this insight is not new to Hutchins but his ethnographic descriptions bring the point alive.

But I would like to point out that Vygotsky proposed a social theory of consciousness more than 60 years ago. Vygotsky theorized, for example, that learning takes place in what he called the "zone of proximal development"; that is, the "space" of what a learner can accomplish with the help of a more experienced peer or teacher vs. what she can accomplish by herself. We learn much more rapidly under the tutelage of a more experienced person than solo. (It isn't merely a matter of rich paths between neurons. Nor can we just open up the head and pour in the right representations as a cognitive scientist would have it.) Activity theorists (and distributed cognitivists) understand that there is a whole process by which the teacher assesses what it is the learner knows or doesn't know; decides how to shape representations so the learner can grasp them; uses the physical world to embody the task (perhaps guiding the hand to trace a letter of the alphabet, for example); gives the learner exercises to practice the new task, and so forth. Cognition is not fundamentally a matter of clever representation, or forming those neural pathways (though that certainly is a physical result of activity)

but of a special kind of social interaction embedded in a practical activity, whether attending school or navigating a large ship.

After Vygotsky, activity theorists continued to stress the point that cognition is rooted in the social and practical. As Luria (1971) noted, "Cognitive processes are not independent and unchanging 'abilities'...; they are processes occurring in concrete, practical activity and are formed within the limit of this activity." Leont'ev (1974) observed, "...a person's mental processes acquire a structure necessarily linked to sociohistorically formed means and modes, which are transmitted to him by other people through teamwork and social intercourse." This could have been the epigram for Hutchins' Chapter 6. Hutchins reprises (without attribution) these basic activity theory points throughout the book as he describes how quartermasters learn as they work on a team, how tools are used in navigation, how the practical activity of navigation shapes cognition. Hutchins' points resonate clearly with activity theory when he talks about "locating cognitive activity in context," (p. xiii), language as a mediating technology (p.300), a cognition that is "not entirely internal to individual persons" (p. 118), and "learning that happens in the doing" (p. 373).

The Functional System

A key point Hutchins makes in *Cognition in the Wild* is the importance of functional systems which he defines as "systems composed of a person in interaction with a tool" (p. xvi). His point is that these systems mean that what a person can do with a tool is profoundly different than what a person can do without the tool. To talk about the person without the tool—cognition without tools—is to make a huge mistake. This is the most telling point Hutchins makes against the traditional cognitive science approach. While cognitive scientists would not deny the importance of tools, in practice their analyses are confined to abstract representations. Simulating on a computer, or even taking seriously, how to count on your fingers, or the hug the teacher gives you, or the way you close your eyes when you are trying to retrieve a memory are not stunts a cognitivist would attempt. These actions are just not amenable to being reduced down to abstract representations. But such actions count among the tools of human cognition. (Activity theory forces you to look much more broadly at the notion of "tools.") I don't know why squeezing your eyes shut helps when you're trying to remember something, but sometimes it does. The human need for encouragement means that you may keep at a task when you get the hug and give up in frustration when you don't. The hug is a real tool in getting the task done. Counting on your fingers is a great way to learn simple arithmetic (even if you have to do it under the desk because the teacher is not a proper activity theorist). Cognition depends on these tools in a very material and everyday way. Hutchins gives many wonderful examples of functional systems on board the helicopter transport. He points out that such systems may be very temporary, coming together at a particular moment to solve a particular problem.

Again, this insight was provided earlier by activity theorists. And it has been elaborated beyond Hutchins' discussion. The notion of a "functional organ" as the activity theorists refer to it, goes back to Marx (and Hegel), as Zinchenko (1996) explains:

From Marx's point of view, the major property of any living system—whether individual or social—is its capability to create needed organs in the course of its growth and development. Similar ideas were being developed by the Russian physiologist Ukhtomsky, who defined a dynamic, integral functional organ: "Usually we associate the name 'organ' with the notion of something that has already formed, something static and constant. It is not necessarily so. Any temporary combination of forces which is capable of attaining a definite end can be called an organ" (Ukhtomsky, 1978, cited in Zinchenko, 1996).

Zinchenko (1996) elaborates the notion of a functional organ building on the work of the Russian physiologist Bernshtein (1924) and the psychologist Zaporozhets (1986). (See also Leont'ev, 1981 and Engeström, 1991 for discussion of functional organs.)

DISTRIBUTED COGNITION AND ACTIVITY THEORY: DISCONNECTS

As we have just seen, in many ways, distributed cognition and activity theory are singing the same song. But the human that lies beneath the surface of a distributed cognition account is not the human that activity theory speaks of. Let's look at the ways in which distributed cognition veers radically away from activity theory. What leads to this divide is the insistence in distributed cognition that people and things are fundamentally the same, that the same language can be used to describe how people and things behave, that both are similar parts in a larger computational/cognitive system. This notion that people and things are the same plays out in many ways throughout Hutchins' book. It is to these manifestations of the notion of the symmetry of people and things in the distributed cognition approach that we now turn.

Cognition, Computation, Task Performance and the Division of Labor

Although an anthropologist, Hutchins has been very influenced by cognitive science; indeed he states that he has written his book for cognitive scientists. While Hutchins has much new to say to that community, at the same time he wishes to preserve their definition of cognition as computation. Computation, as Hutchins defines it, is "the propagation of representational state across a variety of media" (p. xvi). Cognitive science has not paid special attention to different kinds of media, but the idea of propagating changing state across representations is squarely within the cognitive science view of the world.

For Hutchins, a person is a "medium" just as surely as a fathometer or chart is a medium. Cognition takes place, at the highest level of abstraction, indiscriminately across media of all types. That Hutchins believes this is evident in his statement that he seeks a "concept of computation that does not require a change of theory to cross the skin" (p. 117). In other words, what's cognitive for the hoey is cognitive for the human. To understand cognition we must understand how state is propagated across various media, be they people or tools.

Hutchins' theory of propagating representations has been a point of great confusion for me in reading his book. I cannot make out why he applies the term "cognition" to inanimate tools as well as to people. In activity theory, a tool mediates a relationship between a cognizing person and reality, but the tool does not in itself exhibit any cognition. The tool bears information but it does not think. The neuroscientists and philosophers speak of cognition and awareness as intimately related; for example, Searle (a philosopher), defines consciousness as materially involving awareness, as we discussed earlier in the section on neuroscience (Searle, 1996). (Common sense concepts of cognition also posit awareness, e.g., my dictionary defines cognition as "the act or process of knowing including both awareness and judgment" (Webster's Ninth New Collegiate Dictionary)). Flattening cognition to obliterate awareness and judgment seems to me to profoundly miscast what cognition is all about.

The problem may be that Hutchins confuses the cognitive and the functional. A tool functions; it does not think. Further evidence for this confusion lies in Hutchins' notion of distribution. He goes back and forth between descriptions of what individual people can do and the tasks that can be accomplished by the group, wanting all the while to call both "cognitive." Group properties are different than individual properties he says. No one would argue with such a statement. Whether you look at classical systems theory (Wiener, 1948; Ashby, 1956; Bertalanffy, 1968) or work on the division of labor which makes clear that a society can accomplish things that an individual cannot (and Hutchins himself cites this work), it is clear that a group functions differently—in terms of the work that can be accomplished—than a single individual.

Hutchins however wants to call this group function "cognitive." That seems obscurantist. Where is the awareness and judgment in a fathometer? Or in the system as a whole? Hutchins presents no data that a system has awareness or judgment beyond that of the individuals in the system. The system can be organized to take maximal advantage of human awareness and judgment, and set things up so that complex tasks get done by leveraging that awareness and judgment, but it is obscure to assert that a social system, or the media within it (apart from people, if you insist on calling people media), are cognizing entities. A system can evolve over time such that human cognition is effectively leveraged, but that is not to say that the system itself cognizes anything.

I think Hutchins gets tangled up in the notion of assigning the term "cognitive" to a system because he so profoundly recognizes that large scale tasks like navigation involve the transmission and transformation of a huge amount of information across people and their tools. Information is part of cognition, but it is not cognition. The practice of navigation itself does not "know" anything or cognize in any way, but it is ordered to take advantage of the cognitive contributions of individuals as they use their many tools.

Hutchins is aware, at some level, that his description of "cognition" is problematic. In Chapter 8, "Organizational Learning," he attempts to demonstrate how the shipboard organization "learned" something that no individual team member learned. The learning was not achieved by "conscious reflection about the work" (p. 317) as he says; rather, a good solution to a problem on the ship was achieved by small adjustments to local conditions, rather than being designed by any one person.

But this just tells us that sometimes such local adjustments work in certain situations. It does not tell us that the system exhibited cognition. What if there had been a disaster (as in other complex environments such as nuclear power plants) and the local adjustments had not worked? Would we say the organization had failed to learn? We could, but I prefer the explanations of those such as Perrow (1984) who point to specific instances of human error, badly designed tools, problematic power relations, and so forth in analyzing system failure. (See also Linde, 1988 for a fascinating account of pilot error caused by communication breakdowns attributable to the pilots' status hierarchy.) It seems much less mystical to me to talk about real people making real mistakes with bad tools (and bad lines of communication, etc.) than to assert a cryptic level of system cognition.

A great deal of Hutchins' book is devoted to how individuals learn and how they use information-bearing tools. When Hutchins talks about individual "task performers" as he sometimes refers to the sailors on the ship, he's on solid ground. It is here that the rich descriptions of tool use shine. When he attempts to talk about the "cognition" in the other media, stretched across a socio-material system, or the "cognitive properties of human groups" (p. 176) the discussion is less lucid. The "cognitive properties of human groups" are nothing more than a division of labor and the functional interdependencies of the components of any system. Hutchins again senses his own confusion and he begins, well into the book, (p. 176) to talk about "cognitive labor." (The plethora of terminology in the book is an indication of muddles in the models. Hutchins keeps trying to invent new terms to describe what he's talking about, without saying what the distinctions between the terms might be.) What Hutchins really means is that knowledge is distributed across the different people on the ship as is the information contained in the tools of the ship. Knowledge and information are then used by various people in their own cognitive tasks.

Cognitive tasks are distributed in a social system—as are manual, social and emotional tasks. At times Hutchins himself make this crucial distinction, e.g., when he speaks of "computational tasks [that] are socially distributed" (p. 185). Here the book is fluent and coherent. That the system is organized to bring labor and information into coordination is undeniably true. And Hutchins' wonderful contribution is that he details this process with great care, precision and clarity.

But such coordination is not cognition. (In a review of *Cognition in the Wild*, Latour (1996) calls coordination "one of the fetish words of the book"). This sentence from the book sums up the confusion: "These systems are simultaneously cognitive systems in their own rights and contexts for the cognition of the people who participate in them" (p. 288). But only the latter makes any sense. The *system* is not cognitive any more than it is manual or emotional; it is rather an organizing framework in which various kinds of labor—done by real humans—are accomplished in an organized way. The system itself does not exhibit cognition just as it does not perform manual labor or feel emotion or behave sociably.

A fetish of long-standing in cognitive science has been that cognition can be wrenched away from body and soul and reified as something apart from the total person (such that it can be simulated on computers). And that cognition is different in kind than other human capabilities such as physical skill (remember skiing in deep powder?) or spirituality. Hutchins makes this very mistake of reification in thinking that a system can be a cognitive entity because it can instantiate a cognition apart from body and soul. I don't think Hutchins would suggest that a system performs manual labor or that it has a good time at parties or prays to its god(s).

While Hutchins yearns for one big flat system in which people do not have their pesky peculiarities and are really no different than any other "medium," he is nevertheless aware that this really isn't quite right. He smells the fishiness of a theory in which a person is not so different than a hoe (which, by the way, is a protractor like device). Hutchins thus declares people a "special medium": "The thinker...is a very special medium that can provide coordination among many structured media—some internal, some external, some embodied in artifacts, some in ideas, and some in social relationships (p. 316).

Of course the whole game is blown if the thinker must be "special." We then do not have a theory that does not have to change when it "crosses the skin."

At the end of the book Hutchins takes cognitive science to task for ignoring perception, motor behavior, emotion, the body, history, context and culture. So he sees the complexity of real people. But the theory of propagating representations across media (animate and inanimate) could not possibly account for any of these things either. It's not just that Hutchins was looking at something else, or that he'll get to emotion and so forth sooner or later; it's that a view of a person as a computing entity propagating representational state cannot find its way to a view of a person who knows the tango, or writes poetry or laughs at Beavis and Butthead. And it cannot find its way to a description of a cultural system in which such things are possible.

Scaling Back the Person

What most disturbs me about the distributed cognition point of view is the idea that people are not too bright and that our tools exist to prop up our pathetic intellects. (How we are smart enough to create these tools to make up for our stupidity has not yet been addressed by distributed cognition theorists.) Activity theory looks to the potential for development inherent in all of us and says that we start out needing a lot of help from other people to learn things, but then we acquire the psychological tools to do a lot of developing on our own. Other people and artifacts continue to be important but there is a very real sense in which the individual bootstraps his way to greater and greater development (at least under certain historical conditions). Thus activity theory has no problems with an Albert Einstein or a Martha Graham or a Martin Luther King.

The distributed cognition perspective, by contrast, is chary of expert performance (and here is a real reaction to cognitive science which loves those smart puzzle solvers and brilliant chess masters). Distributed cognition

sees the world as being structured so that most tasks are pretty easy (see also Lave, 1988; Norman, 1991). Hutchins details the simple tasks performed by the sailors and the way the tasks are broken down and organized so that nothing is terribly difficult. Of course terms like "difficult" are extremely relative; to me the literacy of the sailors is a major cognitive achievement as are the social skills necessary to interact with their fellow sailors. Hutchins does not agree with me on this point. He says, "...the cognitive abilities that navigation practitioners employ in their use of the forms and inscriptions are very mundane ones—abilities that are found in a thousand other task settings" (p. 133). He doesn't mention social skills a great deal, but after describing a fascinating non-verbal exchange between two sailors who negotiated a complex, novel alignment of tools, he merely remarks, "The social skills required to enter into shared task-performance relationships probably develop fairly early in life." This offhand comment implying that the social skills were not the result of a great deal of development dismisses the accomplishment of the two men. This kind of complex coordinated behavior is a threat to the distributed cognition approach because it is so clearly something that only humans can do. It means the theory has to change when it crosses the skin.

To underscore the point that navigation is organized to be easy, Hutchins describes how the use of algebra and arithmetic are avoided on board ship: "...tools and techniques permit the task performer to avoid algebraic reasoning and arithmetic...Rather than amplify the cognitive abilities of the task performers...these tools *transform* the task the person has to do by representing it in a domain where the answer or the path to the solution is apparent...the existence of such a wide variety of specialized tools and techniques is evidence of a good deal of cultural elaboration directed toward avoiding algebraic reasoning and arithmetic..[the tasks] are part of a cultural process that tends to collect representations that permit tasks to be performed by means of simple cognitive processes" (p. 155, emphasis in original).

I find this line of reasoning unsatisfying. Algebra and arithmetic are human artifacts and lots of people understand and use them. Ordinary people such as carpenters use the Pythagorean theorem when calculating dimensions for stairs. Insurance sales people use probabilities in figuring out what products their customers should buy. Civil engineers design bridges using algebra. There is nothing beyond human capabilities about algebra or arithmetic.

I would have been quite happy to see a *cultural analysis* in *Cognition in the Wild* in which the fact that many people who end up in the American armed forces do not have much formal education (or did badly in school) means that algebra is to be avoided. But there is no such analysis. Rather, it is suggested that culture tends, in a general sense, toward making things simple. The very title of the book—*Cognition in the Wild*—asserts that what is at stake is a "natural" depiction of cognition as it "really is." Not cognition in a *particular* culture, but cognition as it naturally occurs. Others (e.g., Brazerman, 1996) have pointed out that what happens in specific navigational tasks within the confines of an amphibious helicopter transport in the highly elaborated culture of the U.S. military is not necessarily what happens elsewhere. (See Ratner, 1997 on the propensity for cognitive anthropologists to take the culture out of cognition.)

While Hutchins avoids making a blanket claim that tools do not amplify cognition, it is worth looking at an opposing view since Hutchins does not dwell on expert performance. Kaptelinin (1996), an activity theorist, responds to Norman's (1991) argument about amplifying cognition vs transforming tasks (exactly the same argument Hutchins makes, as described above). Kaptelinin points out that, "tools not only change the task but often empower the individual even if the external tool is no longer used." He provides the example of a novice pool player who has to actually hit a ball with the cue to see where the ball will go, vs the expert who knows in advance where the ball will go and can plan a shot accordingly. The expert pool player has not substituted a simple task for a hard one; he has learned how to do a hard task through the practical activity of shooting a lot of pool. His cognition has truly been amplified (and so would say the marks of pool sharks who have learned the hard way, as memorialized in many American movies).

Citing Rumelhart et al. (1986) Hutchins goes on to enumerate the things he thinks "people are good at: recognizing patterns, modeling simple dynamics of the world and manipulating objects in the environment." This is entirely too simplistic, and says nothing about expert performance and the amazing and surprising things people do all the time. Lave (1988) has a similar perspective in her description of people who she labels "just plain folks." (Her book details how "just plain folks" avoid arithmetic when grocery shopping). Sandberg and Wielinga (1993) note the danger of reduction in the distributed cognition approach where "mind is a simple organism interacting with its environment and producing complex behavior through the application of simple behavioral rules." The inadvertent elitism in the distributed cognition stance is unfortunate; surely a Hutchins or a Lave would not regard their own work as "simple" or "mundane" (a word Lave also uses) or merely a matter of doing tasks where a "solution is apparent." A theory of cognition in which some people are smart and some just slog through the easy tasks seems wrong to me.

It is obvious that people have amply proved throughout human history that they are good at all kinds of things. Some people make beautiful gold jewelry and some people tell fabulous stories and some write terrific code.

None of these activities is possible without a great deal of development. And yet none of these activities is rare or bizarre. And people are always looking for new things to be good at. I recently saw a television show on "skysurfing" in which people train themselves to jump out of airplanes, and using a snowboard-like board, "surf" through the air (using the board as a wing) until the last possible second when they have to open a parachute. While I admit this isn't the stuff of everyday life, neither is it somehow not human. Hutchins offers a discussion of difficult cognitive tasks that are not "representative" that he thinks researchers have paid too much attention to. He paraphrases D'Andrade (an American anthropologist) who has noted that, "If we want to know about walking, studying people jumping as high as they can may not be the best approach" (p. 367). But an activity theorist would say you have to walk before you can jump and there is a developmental path that anyone might choose in going from walking to high jumping. It would be a matter of defining an "object" in the activity theory sense (wanting to be a high jumper), and then going for it, full bore. And if you couldn't walk, you couldn't high jump. They are utterly related to one another. High jumping does not have to be seen as "unrepresentative"; it is something a person might very well choose to pursue, if they already know how to walk and if they decide they want to do it.

The view of culture as tasks-made-simple reaches a nadir in Hutchins' analogy (borrowed from Simon, 1981) between people and ants. Extending Simon's original analogy, Hutchins suggests that what history means to an ant is that an ant occurring later in ant history can follow chemical trails left by its predecessors earlier in ant history and can then find food sources much more readily than the earlier ants. "Is this a smart ant?" asks Hutchins. "No, it is just the same dumb sort of ant, reacting to its environment in the same ways its ancestors did. But the environment is not the same...Generations of ants have left their marks on the beach, and now a dumb ant has been made to appear smart through its simple interaction with the residua of the history of its ancestors' actions" (p. 169).

The same dumb ant! Please. It pains me deeply to hear a fellow human referred to this way, even in an analogy. The notion of a person "reacting to its environment" was what behaviorism was all about; we seem to be revisiting that sad view of humanity again.

The analogy with the ant falls precisely into the trap Sandberg and Wielinga (1993) warn of: the person is made to seem a "simple organism interacting with its environment and producing complex behavior through the application of simple behavioral rules." But in the ant world, the chemical trails are automatic biological excreta of the ant; neither producing nor placing the trails involves any creativity or awareness or intelligence on the part of the ant. In the human world, tools need to be dreamed up, manufactured, distributed and explained to other people. There is no real analogy between biologically produced chemical trails and culturally produced and reproduced human tools, especially in the need people have to creatively design tools that will meet future needs.

The question remains: how can people be so clever as to create a world in which they can be so dumb? Hutchins does not ask how the tools that mediate task performance were invented. Tools appear full blown, like Venus on the half shell, as "crystallizations of knowledge and practice in the physical structure of artifacts" (p. 96). Latour notes that "Hutchins should have countered the objection that there must be a huge difference in applying routine knowledge and producing new knowledge" (Latour, 1996).

Hutchins asserts that the whole social organization of navigation on board ship is dedicated to bringing the ship in safely. But again, the procedures and arrangements in the social organization must be devised and enforced by someone. Even if a system did the right thing without any larger conscious agent designing a solution, someone must recognize a solution as useful and codify and enforce it in the future. Otherwise such innovations will be one-time occurrences, not cultural accomplishments. Hutchins uses the word "residua" to describe the accomplishments collected in a culture; it is an odd choice of words, suggesting a set of inadvertent by-products (as in the residua from a chemical process). By contrast, in talking about how a person encounters culture Leont'ev (1974) observed, on a more positive note, "a person's activity assimilates the experience of humanity."

The Object of the Game

Every game has an object. Every murderer has a motive. One of Leont'ev's principal achievements was pointing out that the very definition of an activity comes from its object—that to which the activity "always answers."

In distributed cognition while there is a systemic goal, such as bringing the ship safely into harbor, there are no human motives. To introduce human motive would be to introduce conflict, contradiction, tension, ambiguity, dislocation, discoordination. It would be to admit that individuals count, or can count, for a lot in the system. It would mean that the person has to become an even more "special medium" because having a motive is not

something a thing has (though Latour, 1993, who sees people and things as even more equivalent than Hutchins does, gives it the old college try with his assertion that tools have goals "built into them" by their designers). The whole thrust of the distributed cognition approach is that transitions between media happen in a smoothly coordinated way. As Halverson (1993) notes, distributed cognition accounts work best for "highly rationalized systems." Hutchins does describe a couple cases of breakdowns in ship function, but they are cast as a depiction of how quickly the system gets back on its feet, how rapidly it damps out any dysfunction—not how raw open conflicts can escalate and profoundly change a system.

But even if there were open conflict on the ship (Billy Budd or Captain Ahab, say), a distributed cognition account would have few tools for describing or analyzing what was happening. Activity theory has a strong notion of contradiction, straight out of Marx, and looks for movement and growth in systems as a partial result of conflict and tension (see Kuutti, 1996). Raiethel (1996) has expanded this notion to describe how organizations go through cycles of *coordination* (the level of smooth functioning Hutchins describes) to *cooperation* (small disruptions in the organization that do not require major realignment) to *co-construction* (major breakdowns that require extensive communication and realignment within the organization). This is a much more expansive picture of how a system functions than confining analysis to the coordination level and proposing that coordination is the major modus operandi of a system—and then calling this "cognitive." Raiethel (1996) does not call his levels "cognitive"; he is comfortable in describing system dynamics, at the level of the organization, and calling it just that.

Awareness Again

I have proposed that we follow philosophy and neuroscience and Webster's in asserting that cognition has to do with awareness. Hutchins has set out to do something else, but he does use the word cognition—a rather special word—to label the coordinated flow of media that he describes so richly. It would be more precise, in my opinion, to describe that phenomenon as information coordination (or some such), leaving "cognition" a larger space in which to range. Hutchins' contribution is in drawing attention to the intricacies of information flow between people and tools, rather than in providing a general framework in which to conceptualize cognition.

Hutchins does not talk much about awareness, but when he does, it is uneasily. He notes that "...human institutions can be quite complex because they are composed of subsystems (persons) that are 'aware' in the sense of having representations of themselves and their relationships with their surroundings" (p. 350). The rhetoric here is telling: people are "subsystems" (a term he continues with throughout the long paragraph)—not flesh-and-blood people with active intellects. The term "aware" is placed carefully in quotes by Hutchins. To make the propagating-representations theory work, people have to be "media" or "malleable media" or "subsystems" or "adaptive systems" (all terms Hutchins uses at different times in the book for *people*) and even sometimes ants. Hutchins also refers to the navigational team as a "computational machine" (p. 185). When we call people people, we are less likely to forget how messy and unpredictable and creative and surprising we are. Media and systems are things—not people. A vision of people as things just makes it harder to figure out who we are. It makes it harder to design and evaluate technology too, if we forget about whose needs we are really addressing.

The neuroscientists do not make the mistake of eliminating awareness or equating people and things. Edelman's (1992) language, for example, at some points is very close to that of activity theory.² He observes that to get started in the game of understanding consciousness, or mind, a few basic presuppositions are in order. First, things do not have minds. Second, normal people have minds. Third, beings with minds can refer to other beings or things. Edelman calls the third proposition intentionality (after the German philosopher Brentano). As Edelman says, there must always be an "awareness *of* something;" "it always has an object" (*italics in original*). This is object in the second activity theory sense: that toward which activity is directed, around which activity is coordinated, and which will be crystallized in a final form when the activity is complete.

² Edelman proposes a strongly biological theory of mind, but at the same time he is aware that "no amount of neuroscientific data alone can explain thinking" and that social and cultural interaction are crucial. But his focus on biology inhibits him from going far enough to grasp the basic tenet of activity theory—that consciousness is quite literally *manufactured* through our interactions with others and with tools. That is the essential basis of mind, though the brain and its wonderfully complex structures provide the material substrate on which such interactions are stored, used, organized and reorganized.

WHERE NEXT?

Activity theory provides, in my estimation, the best conceptual framework for a scientific study of consciousness. It has been fruitfully applied to many problems of human-computer interaction which materially involve consciousness (see Nardi, 1996) and has potential for further growth and development. There is a strong recognition of the importance of culture in activity theory, though it has not yet been realized in a large body of empirical work (Ratner, 1997). Modern activity theorists such as Zinchenko (1996) propose a spiritual dimension to humanity and describe in secular terms how we might talk about spirituality (seen as an aspect of human development). Activity theory does not fall into the Cartesian trap as cognitive science does. Coming out of a squarely materialist tradition, activity theory is quite prepared to go down to the level of the wetware, as the neuroscientists want to, as long as the importance of practical activity in the everyday world is a fundamental tenet of any theory of consciousness. Activity theory is optimistic concerning human development, celebrating achievement and the social arrangements that are necessary for achievement and development. Activity theory is set up to account for people engaged in the whole range of human activities, from those of great ordinariness to those of great genius.

While I believe the neuroscientists are off a bit on the wrong track in assuming that the brain *causes* consciousness, they are asking the right questions, the big questions. They are ready to ski in deep powder—to look at the messy problems of ambition, joy, sorrow, free will—and they deserve tremendous credit for that. They understand that people are not things, that a fundamental feature of people is their intentionality. Perhaps a marriage of neuroscience and activity theory will one day show that we can understand those 100 billion neurons if we come to understand the way everyday practical activity shapes human consciousness.

In the meantime, it is exciting to live in a time when science is tackling the big questions and, I think, making some progress. All of the four traditions we have visited—neuroscience, cognitive science, activity theory and the distributed cognition approach—make substantial contributions to the investigation of cognition and consciousness. Continuing the investigation will keep us busy for some time to come. A focus on consciousness, in particular, is especially likely to bear fruit as it is an expansive, grand, three-dimensional, somewhat vague concept, requiring analyses that include an account of culture, a sense of the individual as much more than an engine of representation, and notions of intentionality, spirituality and physicality. As Nabokov (1980) put it, "The brain only continues the spine: the wick really goes through the whole length of the candle."

ACKNOWLEDGMENTS

I am indebted to Yvonne Rogers for a careful, critical reading of an earlier draft of the paper. Thanks also to Carl Ratner for his helpful comments. Errors and omissions are my own.

REFERENCES

- Ashby, W. R. (1956) **Introduction to Cybernetics**, London: Chapman and Hall.
- Bannon, L. (1990) **A Pilgrim's Progress: From Cognitive Science to Cooperative Design**. *AI and Society*, 4, 259-275.
- Bertalanffy, L. (1968) **General System Theory**, New York: George Braziller.
- Bannon L. & Bødker S. (1991) Beyond the interface: Encountering artifacts in use. In: J. Carroll (ed.): **Designing Interaction: Psychology at the Human-Computer Interface**, Cambridge: Cambridge University Press, pp 227-253.
- Bellamy, R. (1996) **Designing Educational Technology: Computer-Mediated Change**. In Nardi, B., Ed. **Context and Consciousness: Activity Theory and Human-Computer Interaction**, Cambridge: MIT Press.
- Bertelsen, O. (1994) **Fitts' law as a design artifact: A paradigm case of theory in software design**. Proc. East-West Human computer Interaction Conference, 2--6 August. St. Petersburg, Russia. Vol. I, pp 37--43.
- Blum, B. I. (1995) **Beyond Programming. To a New Era of Design**, New York: Oxford University Press.
- Blumenthal, B. (1995) **Industrial Design and Activity Theory: A New Direction for Designing Computer-Based Artifacts**. In Blumenthal, B. et al., (Eds). **Selected Papers: Human-Computer Interaction**. 5th International Conference, EWHCI '95, Moscow, Russia, July 3-7. Lecture Notes In Computer Science. V. 1015. Springer.
- Bødker, S. (1989) **A human activity approach to user interfaces**. *Human Computer Interaction*, 4,3, 171-195.

- Bødker, S. (1991) **Through the Interface: A Human Activity Approach to User Interface Design**, Hillsdale, NJ: Lawrence Erlbaum.
- Bourke, I., Verenikina I., Gould, E. (1993) **Interacting with proprietary software users: An application for Activity Theory?** Proc. East-West International Conference on Human-Computer Interaction, Moscow, Russia, 3-7 August Moscow. v. 1, pp 219-225.
- Brazerman, C. (1996) **Review of Cognition in the Wild**. Mind, Culture and Activity 3, 1.
- Card, S., Moran, T. and Newell, A. (1983) **The Psychology of Human Computer Interaction**, Hillsdale, NJ: Lawrence Erlbaum Associates.
- Christiansen, E. (1996) Tamed by a Rose: Computers as Tools in Human Activity. In Nardi, B., Ed. **Context and Consciousness: Activity Theory and Human-Computer Interaction**. Cambridge: MIT Press.
- Cohen, A., Candland, K. & Lee, E. (1995) **The Effect of a Teacher-designed Assessment Tool on Instructors' Cognitive Activity**. Proceedings Interact'95, London: Chapman & Hall, pp 405-409.
- Crick, F. (1993) **The Astonishing Hypothesis: The Scientific Search for the Soul**. New York: Touchstone.
- Draper, S. (1993) **Activity Theory: The new direction for HCI?** International Journal of Man-Machine Studies 38.
- Edelman, G. (1992) **Bright Air, Brilliant Fire**, New York, Basic Books.
- Engeström, Y. (1987) **Learning by Expanding: An Activity-Theoretical Approach to Developmental Research**, Helsinki: Orienta-Konsultit Oy.
- Engeström, Y. (1991) Developmental work research: Reconstructing expertise through expansive learning. In: M. I. Nurminen and G. R. S. Weir (eds.): **Human Jobs and Computer Interfaces**, Amsterdam: North-Holland.
- Engeström, Y. & Escalante, V. (1996) Mundane Tool or Object of Affection?: The Rise and Fall of the Postal Buddy. In Nardi, B., Ed. **Context and Consciousness: Activity Theory and Human-Computer Interaction**, Cambridge: MIT Press.
- Gibson, J. J. (1979) **The Ecological Approach to Visual Perception**, Boston: Houghton Mifflin.
- Halverson, C. (1993) **Distributed Cognition as a Theoretical Framework for HCI**. COGSCI Tech Report # 94-03. San Diego: University of California at San Diego.
- Holland, D., & Reeves, J. (1996) Activity Theory and the View from Somewhere: Team Perspectives on the Intellectual Work of Programming. In Nardi, B., Ed. **Context and Consciousness: Activity Theory and Human-Computer Interaction**, Cambridge: MIT Press.
- Hutchins, E. (1994) **Cognition in the Wild**, Cambridge: MIT Press.
- Imaz, M. & Benyon, D. (1996) **Cognition in the Workplace: Integrating Experientialism into Activity Theory**. ECCE-8: Proceedings of the 8th European Conference on Cognitive Ergonomics, T. R. G. Green, J. J. Canas and C. Warren (Eds.) Granada, Spain: University of Granada Press.
- Kaptelinin, V. (1993) **Item recognition in menu selection: the effect of practice**. Adjunct Proc. INTERCHI'93 Conference on Human Factors in Computing Systems. Amsterdam 24-29 May, pp 183-184.
- Kaptelinin, V. (1996) Computer-mediated Activity: Functional Organs in Social and Developmental Contexts. In Nardi, B., Ed. **Context and Consciousness: Activity Theory and Human-Computer Interaction**, Cambridge: MIT Press.
- Kaptelinin, V., Kuutti, K., Bannon, L. Activity Theory: Basic Concepts and Applications. (1995) In Blumenthal, B. et al. (Eds.) **Human-Computer Interaction**. 5th International Conference, EWHCI '95, Moscow, Russia, July 3-7, Lecture Notes In Computer Science. V. 1015. Springer.
- Kuutti, K. (1991) **The Concept of Activity as a Basic Unit for CSCW Research**. Proceedings Second European Conference on CSCW (ECSCW'91), Amsterdam: Kluwer pp 249-264.
- Kuutti, K. (1996) Activity Theory as a Potential Framework for Human Computer Interaction Research. In Nardi, B., Ed. **Context and Consciousness: Activity Theory and Human-Computer Interaction**, Cambridge: MIT Press.
- Latour, B. (1993) On technical mediation. Three talks prepared for " **The Messenger Lectures On the Evolution of Civilization**" Ithaca, New York, Cornell University, April.
- Latour, B. (1996) **Review of Cognition in the Wild**. Mind, Culture and Activity 3, 1, 54-63.
- Lave, J. (1988) **Cognition in Practice**. Cambridge: Cambridge University Press.
- Leont'ev, A. (1974) **The problem of activity in psychology**. Soviet Psychology 13, 2, 4-33.
- Leont'ev, A. (1978) **Activity, Consciousness, and Personality**. Englewood Cliffs, NJ: Prentice Hall.
- Leont'ev, A. (1981) **Problems of the Development of Mind**, Moscow: Progress.
- Linde, C. (1988) **The quantitative study of communicative success: Politeness and accidents in aviation discourse**. Language and Society 17, pp 375-399.
- Luria, A.R. (1972). **The Man with a Shattered World**, New York: Basic Books.

- Nabokov, V. (1980) **Lectures on Literature**, New York: Harcourt, Brace, Jovanovich.
- Nardi, B., Ed. (1996a) **Context and Consciousness: Activity Theory and Human-Computer Interaction**. Cambridge: MIT Press.
- Nardi, B. (1996b) Studying context: A comparison of activity theory, situated action models, and distributed cognition. In Nardi, B., Ed. **Context and Consciousness: Activity Theory and Human-Computer Interaction**, Cambridge: MIT Press.
- Norman, D. (1991) Cognitive artifacts, in J. Carroll, ed. **Designing Interaction: Psychology at the Human computer Interface**, New York: Cambridge University Press.
- Perrow, C. (1984) **Normal Accidents**, New York: Basic Books.
- Raeithel, A. (1992) Activity Theory as a foundation for design. In: Floyd C. et al. (eds.) **Software Development and Reality Construction**, Berlin: Springer.
- Raeithel, A. (1996) **From Coordinatedness to Coordination via Cooperation and Co-construction**. San Diego: Workshop on Work and Learning in Transition, January.
- Raeithel, A. and Velichkovsky, B. (1996) Joint Attention and Co-Construction of Tasks: New Ways to Foster User-Designer Collaboration. In Nardi, B., Ed. **Context and Consciousness: Activity Theory and Human-Computer Interaction**, Cambridge: MIT Press.
- Ratner, C. (1997) **Cultural Psychology and Qualitative Methodology: Theoretical and Empirical Considerations**, New York: Plenum Press.
- Rogalski, J. (1994) **Analyzing distributed cooperation in dynamic environment management: the "distributed crew" in automatized cockpits**. Proceedings Seventh European Conference on Cognitive Ergonomics (ECCE 7), Bonn, Germany, September 5-8, pp 187-199.
- Rogers, Y. and Ellis, J. (1994) **Distributed cognition: An alternative framework for analysing and explaining collaborative working**. Journal of Information Technology 9, 119-128.
- Rubinshtein, S. L. (1940) **Osnovy obschei psichologii [Basics of General Psychology]**, Moscow: The State Teaching-Pedagogical Publishers (UchPedGiz). (in Russian)
- Rumelhart, D. (1986) Schemata and sequential thought processes in PDP models. In J. McClelland, ed. **Parallel Distributed Processing**, Cambridge: MIT Press.
- Saarelma, O. (1993) **New Work -- New Tools. Information system development in the context of development of work activity**. CEC ESPRIT project 6225 COMIC report Oulu-2-2. Oulu: University of Oulu.
- Salomon, G. (1993) **Distributed Cognitions: Psychological and Educational Considerations**. Cambridge: Cambridge University Press.
- Sandberg, J. & Wielinga, R. (1993) **Situated Cognition: A Paradigm Shift?** JAIE 3,2.
- Schön, D. (1983) **The Reflective Practitioner**, New York: Basic Books.
- Searle, J. (1995) **The Mystery of Consciousness**. New York Review of Books, November 2 and November 16.
- Simon, H. (1981) **The Sciences of the Artificial**, Cambridge, MIT Press.
- Sjöberg, C. (1996) **Activities, Voices, and Arenas: Participatory Design in Practice**. Linköping Studies in Science and Technology. Dissertation no. 439. Linköping, Sweden.
- Tikhomirov, O. K. (1991) The psychological consequences of computerization. In: J. Wertsch (ed.) **The Concept of Activity in Soviet Psychology**, Armonk, NY: M. E. Sharpe.
- Vygotsky, L. S. (1930/1978) **Mind and Society**, Cambridge: Harvard University Press.
- Vygotsky, L.S. (1934/1986) **Thought and Language**, Cambridge: MIT Press.
- Wiener, N. (1948) **Cybernetics**, New York: Wiley.
- Zaporozhets, A. V. (1986) **Zbrannye psichologicheskie trudy v 2-h t.t. [Selected Psychological Works in 2 Vols.]** In V. V. Davydov and V. P. Zinchenko, Eds. Vol. 2, Moscow: Pedagogika.
- Zinchenko, V. & Munipov V. (1989) **Fundamentals of Ergonomics**, Moscow: Progress Publishers.
- Zinchenko, V. (1996) Developing Activity Theory: The Zone of Proximal Development and Beyond. In Nardi, B., Ed. **Context and Consciousness: Activity Theory and Human-Computer Interaction**. Cambridge: MIT Press.