1

Activity Theory and Human-Computer Interaction

Bonnie A. Nardi

What is activity theory, and how will it benefit studies of human-computer interaction? This book addresses these questions. Many HCI researchers are eager to move beyond the confines of traditional cognitive science, but it is not clear exactly which direction to move in. This book explores one alternative for HCI research: activity theory, a research framework and set of perspectives originating in Soviet psychology in the 1920s. Just as HCI research is concerned with practical problems of design and evaluation, activity theorists from the outset have addressed practical needs, applying their research efforts to the problems of mentally and physically handicapped children, educational testing, ergonomics, and other areas. Following the lead of dialectical materialism, activity theory focuses on *practice*, which obviates the need to distinguish ``applied'' from ``pure'' science—understanding everyday practice in the real world is the very objective of scientific practice.

Activity theory is a powerful and clarifying descriptive tool rather than a strongly predictive theory. The object of activity theory is to understand the unity of consciousness and activity. Activity theory incorporates strong notions of intentionality, history, mediation, collaboration and development in constructing consciousness (see Kaptelinin, chapter 5; Kuutti, this volume). Activity theorists argue that consciousness is not a set of discrete disembodied cognitive acts (decision making, classification, remembering), and certainly it is not the brain; rather, consciousness is located in everyday practice: you are what you do. And what you do is firmly and inextricably embedded in the social matrix of which every person is an organic part. This social matrix is composed of people and artifacts. Artifacts may be physical tools or sign systems such as human language. Understanding the interpenetration of the individual, other people, and artifacts in everyday activity is the challenge activity theory has set for itself.

Unlike anthropology, which is also preoccupied with everyday activity, activity theory is concerned with the development and function of individual consciousness. Activity theory was developed by psychologists, so this is not surprising, but it is a very different flavor of psychology from what the West has been accustomed to, as activity theory emphasizes naturalistic study, culture, and history.

The chapters in part I explain what activity theory is. They, along with the seminal article, ``The Problem of Activity in Psychology'' by the Russian psychologist Leont'ev (1974) (widely available in English in university libraries), form a primer of activity theory.

Activity theory offers a set of perspectives on human activity and a set of concepts for describing that activity. This, it seems to me, is exactly what HCI research needs as we struggle to understand and describe ``context,'' ``situation,'' ``practice.'' We have recognized that technology use is not a mechanical input-output relation between a person and a machine; a much richer depiction of the user's situation is needed for design and evaluation. However, it is unclear how to formulate that depiction in a way that is not purely ad hoc. Here is where activity theory helps: by providing orienting concepts and perspectives. As Engeström (1993) has noted, activity theory does not offer ``ready-made techniques and procedures'' for research; rather, its conceptual tools must be ``concretized according to the specific nature of the object under scrutiny.''

As we expand our horizons to think not only about *usable* systems but now *useful* systems, it is imperative that we have ways of finding out what would be useful. How can we begin to understand the best ways to undertake major design projects, such as providing universal access to the Internet, effectively using computers in the classroom, supporting distributed work teams, and even promoting international understanding in ways both small (e.g., international video/e-mail pen-pals for schoolchildren) and large (e.g., using technology to find new means of conflict resolution)? Laboratory-based usability studies are part of the solution, but they are best preceded in a phased design process by careful field studies to ascertain how technology can fit into users' actual social and material environments, the problems users

have that technology can remedy, the applications that will promote creativity and enlightenment, and how we can design humane technology that ensures privacy and dignity.

Recently a major American journal of HCI rejected a set of papers that would have formed a special issue on activity theory. The concern was that activity theory is hard to learn, and because we have not seen its actual benefits realized in specific empirical studies, the time spent learning it would be of dubious benefit. The chapters in parts II and III of this book speak to this concern by providing empirical studies of human-computer interaction developed from an activity theory perspective. In these pages you will meet Danish homicide detectives, a beleaguered U.S. Post Office robot and its human creators, disgruntled slide makers, absent-minded professors, enthusiastic elementary school students, sly college students, and others. These people and artifacts, and the situations in which they are embedded, are analyzed with concepts from activity theory. Several interesting ways to structure an activity theory analysis are provided in these chapters, so readers are offered substantial methodological tools to support practice.

Throughout the book we have tried to ``compare and contrast'' activity theory with other techniques and theories to make it ``easier'' to learn (if indeed it is truly difficult). Thus readers will find that as they read the chapters, they may think about activity theory in relation to cognitive science, GOMS, Gibson's work on affordances, Norman's cognitive artifacts, situated action models, distributed cognition, actor-network theory, and other social scientific artifacts. Bannon and Bødker (1991) have compared activity theory to task analysis and user modeling elsewhere, so we have not undertaken that task here. Briefly, they argued that these approaches are very limited in that (1) task analysis provides a set of procedural steps by which a task supposedly proceeds, with little attention to ``the tacit knowledge that is required in many skilled activities, or the fluent action in the actual work process,'' and (2) user modeling considers user characteristics (e.g., is the user an expert or a novice?) but says little about the situation in which the user works or the nature of the work itself.

Activity theory proposes a strong notion of *mediation*—all human experience is shaped by the tools and sign systems we use. Mediators connect us organically and intimately to the world; they are not merely filters or channels through which experience is carried, like water in a pipe (see Zinchenko, this volume). Activity theorists are the first to note that activity theory itself is but one mediating tool for research (as are all theories!) and that like any tool, its design evolves over time (see Kaptelinin, chapter 3, this volume). Activity theory is certainly evolving and growing; it is not by any means a static end point.

Activity theory has a tremendous capacity for growth and change, an intellectual energy that is being realized in research efforts in Russia, Europe, North America, and Australia. I think perhaps this is because of activity theory's rich philosophical and scientific heritage and because it permits such wide scope of analysis. Activity theory provides ample room in the intellectual sandbox for adventure and discovery and leads to the work of philosophers, psychologists, anthropologists, linguists, educators, and others whose thoughts have influenced activity theory. The chapters in part III of this book push on the frontiers of activity theory, expanding its conceptual base.

Let's talk for a moment about the most concrete practical benefit we could expect from activity theory in the near term. The most immediate benefit I hope for is the dissemination of a common vocabulary for describing activity that all HCI researchers would share. Activity theory has a simple but powerful hierarchy for describing activity that could be common coin for all HCI researchers. This hierarchy (described in several of the chapters in this book) has a superficial resemblance to GOMS but goes beyond GOMS in essential ways, especially in describing dynamic movement between levels of activity rather than assuming stasis.

The development of a common vocabulary is crucial for HCI. As we move toward ethnographic and participatory design methods to discover and describe real everyday activity, we run into the problem that has bedeviled anthropology for so long: every account is an ad hoc description cast in situationally specific terms. Abstraction, generalization and comparison become problematic. An ethnographic description, although it may contain much information of direct value for design and evaluation, remains a narrative account structured according to the author's own personal vocabulary, largely unconstrained and arbitrary. Ethnography—literally, ``writing culture''—assumes no a priori framework that orders the data, that contributes to the coherence and generalizability of the descriptive account. This leads to a disappointing lack of cumulative research results. One would like to be able to develop a comparative framework, perhaps a taxonomy as suggested by Brooks (1991), that would help us as we pursue design and evaluation activities. It would be desirable to be able to go back to previous work and find a structured

set of problems and solutions. Activity theory will help us to achieve this goal but not until its concepts become part of a shared vocabulary.

Let us look briefly at a few of the main concerns of activity theory: consciousness, the asymmetrical relation between people and things, and the role of artifacts in everyday life. Each of these concerns (and others) will be considered at length in this book, and I introduce them briefly here to anticipate some of what the reader will encounter.

A basic tenet of activity theory is that a notion of consciousness is central to a depiction of activity. Vygotsky described consciousness as a phenomenon that unifies attention, intention, memory, reasoning, and speech (Vygotsky 1925/1982; see Bakhurst 1991). Does HCI really need to worry about consciousness? The answer would seem to be yes, as we have been worrying about it all along. A notion of consciousness, especially one that focuses on attention and access to cognitive resources, permeates HCI discourse. When we speak of ``direct manipulation,'' ``intelligent agents,'' ``expert behavior,'' and ``novice behavior,'' we are really positing concepts in which consciousness is central. The notion of consciousness has continually snuck in the back door of HCI studies, as Draper (1993) has pointed out. We use the word ``transparent,'' to describe a good user interface—that is, one that is supportive and unobtrusive, but which the user need pay little, if any, attention to. We have borrowed the concept ``affordances'' from Gibson, which practically dispenses with the notion of consciousness but still implies a particular stance toward it. We speak of ``skilled performance,'' implying a kind of mental ease and access to certain cognitive resources peculiar to experts who have become very good at something. ``Novices,'' on the other hand, consciously labor to perform actions that will later become automatic, requiring little conscious awareness. Their less able performance is attributable to their need to focus deliberate attention on task actions while at the same time working with fewer cognitive resources than they will have available later as they gain expertise and experience in their tasks.

Even in the earliest HCI work we find concern with the user's consciousness. In 1972 Bobrow wrote that a programming technique ``can greatly facilitate construction of complex programs because it allows the user to remain thinking about his program operation at a relatively high level without having to descend into manipulation of details.'' This is a succinct statement of the interdependence of the ``how'' and the ``what'' of consciousness: the user's attention is at stake, and at the very same time, so is the content of what he thinks about as he programs.

Consciousness is still with us: Carey and Rusli (1995) argue that simply observing users does not tell the researcher enough; it must be discovered what the user is thinking. They give an example, asking, ``Was a switch in search tactics the result of abandoning an unproductive attempt, or the result of gaining knowledge from the last few actions?'' There are very different implications for technology design depending on the reason for the switch. Looking back more than a decade at Malone's (1983) classic paper on office organization, we find that Malone noted that users' behavior cannot be understood without reference to intentionality: is a user organizing her office so that she can find something later, or so that she will be reminded of something? The observer sees the same behavior but cannot know what it means without asking the user. Malone observed that finding and reminding are quite different functions, equally important for users, and that we cannot understand them if we do not take account of the user's intentions.

The unstudied use of a notion of consciousness will continue to crop up in HCI research, and rather than dealing with each new instance piecemeal, in a new vocabulary, as though we had never heard of it before, an overarching framework prepared to deal with the phenomenon of consciousness will be useful. Draper (1993) talks about ``designing for consciousness,'' and it seems that this is exactly what we should be doing when we discuss the possibility of, for example, ``intelligent agents.'' The notion of agents suggests that the user direct conscious awareness toward the user interface rather than that the user interface disappear ``transparently.'' In a direct manipulation interface, on the other hand, cognitive content concerns the nitty-gritty of one's task, with the interface ideally fading from awareness.

Thus we see from this brief excursion into the difficult subject of consciousness that already we have gained two insights: (1) we must know what the user is thinking to design properly, as Carey and Rusli (1995) argue, and (2) we have a larger conceptual space into which to place differing user interface paradigms such as intelligent agents and direct manipulation.

Of course, psychologists have studied attention and consciousness for a long time; this is not new to activity theory. Activity theory, however, embeds consciousness in a wider activity system and describes a dynamic by which changes in consciousness are directly related to the material and social conditions current in a person's situation (see Kaptelinin, chapters 3, 5; Nardi, chapter 4; Bødker; Raeithel and Velichkovsky, this volume). This extends the concept of consciousness past an idealistic, mentalistic

construct in which only cognitive resources and attention ``inside the head'' are at issue, to a situated phenomenon in which one's material and social context are crucial.

An important perspective contributed by activity theory is its insistence on the asymmetry between people and things (see Kaptelinin, chapter 5; Nardi, chapter 4; Zinchenko, this volume). Activity theory, with its emphasis on the importance of motive and consciousness—which belong only to humans—sees people and things as fundamentally different. People are not reduced to ``nodes'' or ``agents'' in a system; ``information processing'' is not seen as something to be modeled in the same way for people and machines. In activity theory, artifacts are mediators of human thought and behavior; they do not occupy the same ontological space. This results in a more humane view of the relationship of people and artifacts, as well as squarely confronting the many real differences between people and things.

Cognitive science has been the dominant theoretical voice in HCI studies since the inception of our young field. We are beginning to feel a theoretical pinch, however—a sense that cognitive science is too restrictive a paradigm for finding out what we would like to know (Bannon and Bødker, 1991; Kuutti, this volume). Activity theory is not a rejection of cognitive science (see Kaptelinin, chapter 5, this volume) but rather a radical expansion of it. One reason we need this expansion is that a key aspect of HCI studies must be to understand *things*; technology—physical objects that mediate activity—and cognitive science have pretty much ignored the study of artifacts, insisting on mental representations as the proper locus of study. Thus we have produced reams of studies on mentalistic phenomena such as ``plans'' and ``mental models'' and ``cognitive maps,'' with insufficient attention to the physical world of artifacts—their design and use in the world of real activity (Hutchins 1994). Norman (1988) has done much to alleviate this situation, turning our attention toward what Sylvia Plath called the ``thinginess of things'' (Plath 1982), but we still have a long way to go.

Activity theory proposes that activity cannot be understood without understanding the role of artifacts in everyday existence, especially the way artifacts are integrated into social practice (which thus contrasts with Gibson's notion of affordances). Cognitive science has concentrated on *information*, its representation and propagation; activity theory is concerned with *practice*, that is, *doing* and *activity,* which significantly involve ``the mastery of ... external devices and tools of labor activity'' (Zinchenko 1986). Kaptelinin (chapters 3, 5, this volume) and Zinchenko (this volume) describe the activity theory concept of ``functional organ,'' a fundamental notion pinpointing the way the mind and body are profoundly extended and transformed by artifacts (see also Vygotsky 1929, Leont'ev 1981). There are echoes of Haraway's (1990) cyborg here but in a different (and much earlier) voice. The notion of the functional organ, rather than being a riveting poetic image like the cyborg, is a tenet of a larger system of theoretical thought and a tool for further scientific inquiry.

Some readers may be impatient with activity theory terminology. It can be inelegant in translation from the Russian and, worse, confusing. The notion of an ``object,'' in particular, becomes a point of confusion as activity theorists use terms such as ``object-oriented'' in an entirely different way than they are used in the programming community. A degree of forbearance is helpful when first confronting activity theory terminology.

Activity theory challenges much that we have held useful and important in HCI research. But this book is not mounted as an attack on previous work; rather, it is an inquiry into satisfying ways to extend, and where necessary to reformulate, the basis for the study of problems in human-computer interaction. This inquiry is intended to be ecumenical and inclusive yet probing and questioning. There is a new kind of post-postmodern voice struggling to speak clearly here; it is polyvocal and dialogical, to be sure, but also committed to social and scientific engagement. This voice has little use for the peevish debate and posturing that mark much current (and past) discourse; instead the aim is to acknowledge, learn from, and yet go beyond existing theory, to reach for what Bertelsen (1994) calls a ``radical pragmatic science of HCI.'' Many who have come to find activity theory useful for HCI acknowledge a debt to cognitive science, especially the pioneering work of Card, Moran, and Newell (1983), for the suggestion that HCI design can benefit from a rigorous scientific foundation, as well as a debt to participatory design work (Kyng 1991; Muller and Kuhn 1993), which urges a humane, socially responsible scientific practice. That activity theory fuses these two intellectual impulses into a unified approach perhaps explains why we are seeking its counsel at this particular time in the history of our field.

REFERENCES

Bakhurst, D. (1991). *Consciousness and Revolution in Soviet Philosophy*. Cambridge: Cambridge University Press.

Bannon, L., and 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.

Bertelsen, O. (1994). Fitts' law as a design artifact: A paradigm case of theory in software design. In *Proceedings East-West Human Computer Interaction Conference* (vol. 1, pp. 37–43). St. Petersburg, August 2–6.

Bobrow, D. (1972). Requirements for advanced programming systems for list processing. *Communications of the ACM* (July).

Brooks, R. (1991). Comparative task analysis: An alternative direction for human-computer interaction science. In J. Carroll, ed., *Designing Interaction: Psychology at the Human Computer Interface*. Cambridge: Cambridge University Press.

Card, S., Moran, T., and Newell, A. (1983). *The Psychology of Human-Computer Interaction*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Carey, T., and Rusli, M. (1995). Usage representations for re-use of design insights: A case study of access to on-line books. In J. Carroll, ed., *Scenario-based Design for Human Computer Interaction.* New York: Wiley.

Draper, S. (1993). Critical notice: Activity theory: The new direction for HCI? *International Journal of Man-Machine Studies* 37(6):812–821.

Engeström, Y. (1993). Developmental studies of work as a testbench of activity theory. In S. Chaiklin, S. and J. Lave, eds., *Understanding Practice: Perspectives on Activity and Context* (pp. 64–103). Cambridge: Cambridge University Press.

Haraway, D. (1990). *Simians, Cyborgs and Women: The Reinvention of Nature*. London: Routledge.

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

Kyng, M. (1991). Designing for cooperation—cooperating in design. *Communications of the ACM* 34(12):64–73.

Leont'ev, A. (1974). The problem of activity in psychology. *Soviet Psychology* 13(2):4–33.

Leont'ev, A. (1981). *Problems of the Development of Mind*. Moscow: Progress.

Malone, T. (1983). How do people organize their desks? *ACM Transactions on Office Information Systems* 1, 99–112.

Muller, M., and Kuhn, S. (1993). Introduction to special issue on participatory design. *Communications of the ACM* 36:24–28.

Norman, D. (1988). *The Psychology of Everyday Things*. New York: Basic Books.

Plath, S. (1982). *The Journals of Sylvia Plath*. New York: Ballantine Books.

Vygotsky, L. S. (1925/1982). Consciousness as a problem in the psychology of behaviour. In *Collected Works: Questions of the Theory and History of Psychology*. Moscow: Pedagogika.

Vygotsky, L. S. (1929). The problem of the cultural development of the child. *Journal of Genetic Psychology* 36:415–432.

Zinchenko, V. P. (1986). Ergonomics and informatics. *Problems in Philosophy* 7:53–64.