1

Introduction

This book is about developing activity theory as an approach to the investigation of information technologies in the context of human practice. *Acting with technology* is a phrase to position our relationship to technology as one in which people act intentionally in specific ways with technology—ways that we can study and for which we can produce effective designs.

Activity theory was introduced to an international audience in the late 1970s and early 1980s through two publications: the English translation of Leontiev's *Activity, Consciousness, and Personality* (1978), and a collection of papers by Leontiev and other activity theorists edited by James Wertsch with an excellent introduction by Wertsch (1981).

But until the 1990s, activity theory was effectively standing in Vygotsky's shadow. Vygotsky's approach had become popular in the West, having a substantial impact on a wide range of research in psychology and cognitive science (Cole and Scribner 1974; Wertsch 1985; Hutchins 1995), education (Lave and Wenger 1991), and computer support for collaborative learning (O'Malley 1995; Koschmann 1996a). International interest in activity theory increased dramatically during the 1990s, judging from the frequency of citation of key works in activity theory (Roth 2004). A number of papers and books published during that time (e.g., Engeström 1990; Bødker 1991; Nardi 1996a; Wertsch 1998; Engeström, Miettinen, and Punamäki 1999) contributed to the increased awareness of the ideas and potential of the approach. According to Roth (2004), part of the credit for the uptake of activity theory should be given to Yrjö Engeström, who "through his publications and presentations in a variety of disciplines spread the word. . . ."

The aim of *Acting with Technology* is to provide a thorough understanding of activity theory through a systematic presentation of its principles, history, relationship to other approaches, and application in interaction design. A decade ago, *Context and Consciousness: Activity Theory and Human–Computer Interaction*, a volume edited by one of us, and to which both of us contributed chapters, was published by the MIT Press (Nardi 1996a). *Context and Consciousness* presented a variety of positions and arguments unified by the common objective of making the case for activity theory as a potential theoretical foundation for human–computer interaction. *Context and Consciousness* contributed to the turn to contextual approaches in HCI, foregrounding an understanding of activity as central to the concerns of specialists in human–computer interaction.

The present book has different ambitions. *Acting with Technology* addresses three questions:

1. What impact has activity theory had on interaction design? We present and discuss key results of interaction design research based on activity theory.
2. How does activity theory relate to other theoretical approaches in the field? We contextualize activity theory in the ever-changing theoretical landscape of interaction design by way of a comparative analysis of current approaches.
3. What does "activity theory" really mean? Activity theory is sometimes considered an "esoteric" approach (Engeström 1999a) because systematic introductions to its main principles, intended for general audiences, rather than enthusiasts, are nonexistent. In this book we make an attempt to put together a primer in activity theory, to deliver activity theory "in a nutshell."

The domain of the book is *interaction design*, understood in a broad sense. The term has been used in the human–computer interaction (HCI) and computer-supported collaborative work (CSCW) communities (Winograd 1996; Preece, Rogers, and Sharp 2002; Bannon 2005; Pirhonen et al. 2005), and by those in the field of digital design who see their work as related to but distinct from human–computer interaction (Wroblewski 1991; Gaver, Beaver, and Benford 2003; Löwgren and Stolterman 2004). Löwgren and Stolterman (2004) defined interaction design as "the process that is arranged within existing resource constraints to create, shape, and decide all use-oriented qualities (structural, functional, ethical, and aesthetic) of a digital artifact for one or many clients." This definition reveals some reasons for the shift to the term "interaction design": it is not only computers, but digital artifacts of all kinds that interest us, and not only the computational abilities of such artifacts, but the totality of their potentials.

Winograd (1996) defined interaction design as "the design of spaces for human communication and interaction." This definition is similar in spirit to that of Löwgren and Stolterman, but more general. While Löwgren and Stolterman suggested a context of design in workaday settings, invoking clients and resource constraints, Winograd's definition can be construed as covering a wide range of issues, from empirical studies with design implications to work in hands-on design settings.

Interaction design is a broad term inflected in different ways in different communities. To us, interaction design comprises all efforts to understand human engagement with digital technology and all efforts to use that knowledge to design more useful and pleasing artifacts. Within this arena, the main audiences for this book are those who conduct work in the fields of human–computer interaction, computer-supported collaborative work, computer-supported collaborative learning, digital design, cognitive ergonomics, informatics, information systems, and human factors.[1]

Activity theory fits the general trend in interaction design toward moving out from the computer as the focus of interest to understanding technology as part of the larger scope of human activities. HCI began with the notion of a "user." Researchers developed a set of core concepts that advanced the field, such as "user-centered design," "the user experience," "usability," "usefulness," and "user empowerment" (Norman and Draper 1986; Thomas and Kellogg 1989; Cooper and Bowers 1995). Expanding these notions, Bannon (1991) coined the memorable phrase "from human factors to human actors" to emphasize actors in social contexts, consonant with the concerns of CSCW. More recently, attempts to incorporate human activity in interaction design have led to ideas of "activity-based," "activity-centered," or "activity-centric" computing (Norman 1998; Christensen and Bardram 2002; Geyer, Cheng and Muller 2003; Harrison 2004; Muller et al. 2004; Millen et al. 2005) and "activity management" (Moran 2003). These efforts seek to provide a richer framing for interaction design that more closely matches how people actually use technology at work and play.

While it is helpful that such notions of activity-based computing acknowledge the general importance of the meaningful context of interaction between subjects and the world, it is crucial to move to concrete understanding of what activities are. Activity theory can help bridge the gap between insights about the need for broader perspectives and the need for specific tools for thought. As we attempt to study human activities "in the world" (Bannon 2005), we will encounter issues long of interest to activity theory. We believe that activity theory fits a niche opened by the emerging sensibility that studying interaction and activity is essential to the development of interaction design. The basic principles of activity theory underwrite the emphasis in interaction design on the social, emotional, cultural, and creative dimensions of human actors in shared contexts.

Today activity theory is an approach that has transcended both international and disciplinary borders. It is used not only in Russia, where it originated, but also in Australia, Belgium, Brazil, Canada, Denmark, Finland, France, Germany, Italy, Japan, Norway, South Africa, Sweden, Switzerland, the UK, the United States, and other countries. It is applied in psychology, education, work research, and other fields. In this book, we discuss activity theory in the context of interaction design, but in appendix B the interested reader can find information and web links to international conferences, journals, and discussion forums devoted to research based on activity theory from a variety of perspectives.

The book consists of three parts. In part I we give an overview of the basic concepts of activity theory and how they have been used in interaction design research. We discuss the need for theory in interaction design in chapter 2. We explicate the fundamentals of activity theory in chapter 3. We describe applications of activity theory to practical problems of interaction design in chapter 4. We provide a detailed example of an application developed with activity theory in chapter 5.

In part II we turn to more advanced issues. We discuss the notion of the object of activity in chapter 6, describe the use of this notion in an empirical study in chapter 7, and review the history of activity theory, with a focus on key debates that shaped the development of the approach, in chapter 8.

In part III we draw on the discussions in parts I and II to outline current issues and future theoretical development in activity theory. In chapter 9, we compare activity theory with its leading contenders in interaction design—distributed cognition, actor-network theory, and phenomenologically inspired approaches. In chapter 10, we delve more deeply into issues regarding agency and asymmetry raised in chapter 9. We conclude in chapter 11 with some reflections on the future of activity theory.

If we have any advice to our readers, it is to be alert to the coherent whole that is activity theory. As we have explored other theories and empirical research, we sometimes have the sense of seeing a piece of activity theory developed independently. For example, early in his career, Herb Simon discussed the way people conserve "mental effort by withdrawing from the area of conscious thought those aspects of the situation that are repetitive" (Simon 1945). This sounds very much like the operational level of the activity hierarchy in activity theory. Without in any way critiquing Simon (who was not developing a psychological theory but rather describing organizational behavior), we can point to the way such insights crop up as "one-offs" across the theoretical landscape. In activity theory, the operational level is one of three linked levels in the activity hierarchy, not an isolated insight. Another example closer to home is that of GOMS models, which resemble the activity hierarchy but lack an activity level and the possibility of dynamic changes between levels that are part of activity theory. We hope to encourage a holistic reading of activity theory and a cognizance of the way concepts weave together into a patterned whole. Some of the power of activity theory lies in the way it ties insights into larger wholes to provide a clarifying framework for the bigger picture.

In this book we advocate and evaluate the continued development of activity theory as a basis for understanding how people act with technology. We hope to use theory to stimulate great design—the design of digital technologies that address the needs and desires of specific individuals and groups. We also want to understand the fundamentals of our human relationship with technology. These designs and understandings will include the usual activities that we know as the practice of interaction design, but may also stretch to less familiar projects involving how we act with technology, such as analyzing the impact of technologies on the environment or understanding the role of technology in viewing our spiritual relation to the cosmos. Though such projects may appear beyond the scope of interaction design, the technologies we design inevitably have major impacts in these arenas. If we are to continue to deepen our understanding of what it means to act with technology, such concerns will impinge on, and sometimes become central to, our labors.

Activity theory seeks to understand the unity of consciousness and activity. It is a social theory of human consciousness, construing consciousness as the product of an individual's interactions with *people* and *artifacts* in the context of everyday practical activity. Consciousness is constituted as the enactment of our capacity for attention, intention, memory, learning, reasoning, speech, reflection, and imagination. It is through the exercise of these capacities in everyday activities that we develop; indeed this is the basis of our very existence.

This social approach rooted in practical activity contrasts with, for example, biological explanations of consciousness that focus on genetically coded capabilities, or neuroscientific views that situate explanation at the level of nerve tissue, or the Jungian view positing universal archetypes accessible through dreams. Traditional cognitive science attends to representations, casting them as entities that can be modeled equally well for computers as humans. Freudian explanations focus on a small set of early social relations with parents and family. Activity theory proposes that consciousness is realized by *what we do* in everyday practical activity.

To take a simple example, let's consider how an activity theorist might analyze a young child learning arithmetic. Activity theory looks for key people in the child's universe and useful artifacts. In many cultures, children learn math from their teacher who explains numbers and arithmetic operations to them, and encourages and motivates them. The children may also consult more experienced peers. Children initially perform calculations on their own bodies, counting on their fingers silently until they have internalized addition and subtraction. The fingers come into play as a useful "artifact," appropriated by the child as a marking device to aid in counting. Once the child has mastered the facts of arithmetic, the calculation shifts to what activity theorists call the internal plane of actions, and the math is done in the head.

Part of what is distinctive about this formulation is that it goes beyond the representation of the arithmetic problem, beyond the bare bones of the arithmetical processes, out to the environment where the teacher, the friends, and the fingers are. These aspects of the child's universe are essential to our understanding of how the child learns arithmetic. Most theories miss these aspects, or see only one—perhaps the teacher, or the way the problem is represented on paper. In activity theory it is the *doing* of the activity in a rich social matrix of *people and artifacts* that grounds analysis.

This insight was expressed thousands of years ago in Eastern thought. In speaking to Vasettha, Buddha described the primacy of activity in human life:

> One is not a brahmin by birth,
> Nor by birth a non-brahmin.
> By action is one a brahmin,
> By action is one a non-brahmin.
> So that is how the truly wise
> See action as it really is.
> Seers of dependent origination,
> Skilled in actions and its results.
> Action makes the world go round
> Action makes this generation turn.
> Living beings are bound by action
> Like the chariot wheel by the pin.[2]

It is striking that the central image of this poem is a technical one, the chariot wheel with its pin. Here the poet intimates the close link between human action and the technologies that support it. Activity theory has developed the insights of the poets in a scientific idiom, delineating a set of core principles that frame the study of all human activity (see Zinchenko 1996).

We have found the principles of activity theory to be of help as we consider our own chariot wheels and how we design and use them. For several years we have advocated activity theory as a framework for thinking about human activity as it is expressed in the use of technology (Nardi 1992, 1993, 1996a; Kaptelinin 1992; Kaptelinin, Nardi, and Macaulay 1999; Bannon and Kaptelinin 2002). We have observed a steady and growing uptake in the adoption of activity theory among those who find a theoretical framework useful for negotiating the thickets of users and their needs, and technologies and their possibilities. We have been drawn to activity theory because of certain of its tenets that are encapsulated in the notion of people acting with technology. These tenets are:

- an emphasis on human intentionality;
- the asymmetry of people and things;
- the importance of human development; and
- the idea of culture and society as shaping human activity.

Let us first consider intentionality. We live in an ever increasingly designed world, furnished with technologies at every turn. Despite the clearly intentional nature of the act of design—behind every design there is an intention—many of our theories lack a concept of intentionality. In acting with technology, people deliberately commit certain acts with certain technologies. Such a mild statement, seemingly devoid of theoretical freight, is in fact at odds with theories such as actor-network theory and distributed cognition. These approaches posit a sociotechnical network whose generalized nodes are actors that can be either human or artifact. Such actors represent states that move through a system—whether the actor be a pencil or a person. Intentionality is not a property of these generalized nodes. Activity theory distinguishes between people and things, allowing for a discussion of human intentionality.

More broadly speaking, activity theory posits an asymmetry between humans and things—our special abilities to cognize through interactions with people and artifacts are distinctive from any sort of agency we could sensibly ascribe to artifacts. In activity theory, it is essential to be able to theorize intention, imagination, and reflection as core human cognitive processes. Accounts in which people and artifacts are the same deflect such theorizing.

In activity theory *people* act *with* technology; technologies are both designed and used in the context of people with intentions and desires. People act as *subjects* in the world, constructing and instantiating their intentions and desires as *objects*. Activity theory casts the relationship between people and tools as one of *mediation*; tools mediate between people and the world.

Another principle of activity theory is the notion of *development*. Activity theory shares the commitment of the cultural-historical school of psychology because of its commitment to understanding how human activity unfolds over time in a historical frame. Activity theory takes the long view: we cannot understand activity if we do not watch it cycle, grow, change. It would be desirable to establish a practice of design in which the development of users—their ability to grow and change with technology—is of paramount importance. In activity theory, development is a sociocultural process, but the individual is not reduced to society or culture. The dialogical nature of processes of internalization–externalization makes it possible for individuals to transform culture through their activity. As a psychological theory, activity theory has always had a strong notion of the individual, while at the same time understanding and emphasizing the importance of the sociocultural matrix within which individuals develop. As we will discuss in chapter 9, the individual is an important theoretical concept because of the need to account for the interrelated processes of creativity, resistance, and reflexivity. These processes take place in part within individuals as people have the capacity to radically restructure cultural conceptions, transcending culture in unpredictable ways.

Technological creativity is rooted in our primate past. Nonhuman primates can "think out of the box," developing and sharing simple tools to transform their activity. For example, capuchin monkeys have been observed using sticks to reach food (Beck 1980). The great apes, especially chimpanzees, have more sophisticated tool capabilities. In the wild, chimps may use assemblages of anvils and hammers to crack tough nuts (Mercader, Panger, and Boesch 2002). An individual animal in its own well-known environment can suddenly recognize a solution to a problem, and come to see an object as a tool for some useful purpose. As with humans, nonhuman primate development is cultural; tool use among higher primates is specific to distinct animal locales, with local tools and cultural practices providing knowledge of how to use the tools.

How does grounding our theory in a concept of intentionality and the asymmetry of people and things, as well as a strong notion of development, help us as interaction designers? We believe there are several benefits. First, such a theory can provide a matrix in which to reflect on our own practice, to arrange what seem to be disparate threads into a coherent framework. For example, the adoption of approaches such as participatory design and contextual design are responses to the larger problem of addressing the gap between the intentions of designers and the intentions of users. The continuing search for techniques of end user programming (Lieberman 2000) speaks to an unfilled need to increase end users' abilities to realize their own intentions so they can grow and develop over time, becoming increasingly adept with their technologies. The design of agent-based user interfaces, which seek to enact high-level intentions while sparing users the details, is one approach to bringing intentions into the user interface. The current state of designing and using information technologies in education also clearly indicates the importance of taking intentionality into consideration. There has been a growing realization that to have a positive impact on education, technologies should be designed to support purposeful actions of the human actors involved in everyday educational practices (Gifford and Enyedy 1999).

A second benefit of a theory grounded in intentionality, asymmetry, and development is that it can frame discussions of users' continuing frustrations. We do not have to go far to find users who are stymied in realizing their intentions because the technologies offered them are neither usable nor useful. And users often feel daunted by the rapid pace of technological change, which makes it ever more difficult to become skilled with a given technology. Only a decade ago, it was possible to write optimistically about "gardeners and gurus" (Gantt and Nardi 1992), those office experts who became especially proficient with the technologies in use in their local settings and could help their less technically inclined colleagues. Today, because technologies change more rapidly and work groups are less stable, we cannot be as sanguine about the role of local experts in the ecology of a given work setting. Activity theory's attention to issues of development commits us to taking such issues seriously.

The third benefit is that of reckoning with the long-term impact of the technologies we design. If a historical developmental perspective frames our view, we cannot merely hope for the adoption of the technologies we intentionally design; we must consider wider impacts. For example, the batteries and components of wireless devices contain arsenic, antimony, beryllium, cadmium, copper, zinc, nickel, lead, and brominated flame retardants—all toxic. Wireless devices, including cell phones, pagers, PDAs, pocket PCs, portable email readers, and mp3 music players, are being manufactured by the billions. Yet we have not designed or implemented adequate means of handling the wastes they release. Toxins leach into groundwater when wireless devices are discarded in landfills, and dioxins are created when they are incinerated. Used cell phones (and computers) are often donated to Third World countries, so the waste reaches its final resting place in the air and water of the poorest countries (see *Waste in the Wireless World: The Challenges of Cell Phones*, 2002). As designers, how do we respond to these realities?

Activity theory is self-reflexive, and we are encouraged to find ways to inform our own development. To mitigate the harmful effects of, say, the wireless devices we design, we might look to the fields of architecture and manufacturing which are working with techniques of "green design," "lifetime design," and life cycle assessment. While such a move might seem an unmanageable increase in the scope of our efforts, other disciplines have adopted these concerns as part of their practice. When our theories reveal intentionality and historical development as visible theoretical constructs, we are more likely to entertain conversations about long-term effects than if our theories conceal them. Miettinen (1999) noted that understanding the historical development of human consciousness is needed to make sense of the relations between humans and their environment. Such an understanding is critical when the aim is to analyze the work of constructing associations between heterogeneous entities and the work of creating "new assemblies of materials and humans" (Miettinen 1999).

Activity theory opens up avenues of discussion concerning human interaction with technology and potentially can be fruitful in encouraging participation in conversations about the larger global concerns that the deployment of our technologies unquestionably affects. If we are *acting with technology*, both possibilities and responsibilities expand. The object of this book is to stimulate further discussion of the theoretical basis for understanding how people act with technology.