

Simulating activities: Relating motives, deliberation, and attentive coordination

Action editor: Tom Ziemke

William J. Clancey^{*,1}

Computational Sciences Division, M/S 269-3, NASA/Ames Research Center, Moffett Field, CA 94035, USA

Received 4 May 2001; accepted 3 December 2001

Abstract

Activities are located behaviors, taking time, conceived as socially meaningful, and usually involving interaction with tools and the environment. In modeling human cognition as a form of problem solving (goal-directed search and operator sequencing), cognitive science researchers have not adequately studied ‘off-task’ activities (e.g. waiting), non-intellectual motives (e.g. hunger), sustaining a goal state (e.g. playful interaction), and coupled perceptual–motor dynamics (e.g. following someone). These aspects of human behavior have been considered in bits and pieces in past research, identified as *scripts*, *human factors*, *behavior settings*, *ensemble*, *flow experience*, and *situated action*. More broadly, *activity theory* provides a comprehensive framework relating motives, goals, and operations. This paper ties these ideas together, using examples from work life in a Canadian High Arctic research station. The emphasis is on simulating human behavior as it naturally occurs, such that ‘working’ is understood as an aspect of living. The result is a synthesis of previously unrelated analytic perspectives and a broader appreciation of the nature of human cognition. Simulating activities in this comprehensive way is useful for understanding work practice, promoting learning, and designing better tools, including human–robot systems.

© 2002 Elsevier Science B.V. All rights reserved.

Keywords: Situated action; Human–robot systems; Activity theory; Behavior simulation

1. Introduction

All human activity is purposeful. But not every goal is a problem to be solved (cf. Newell & Simon,

1972), and not every action is motivated by a task (cf. Kantowitz & Sorkin, 1983). For example, listening to music while driving home is part of the practice of driving for many people, but it is not a subgoal for reaching the destination. Listening to music is part of the *activity* of driving, with an emotional motive (Leont’ev, 1979). The goal of reducing stress or making the drive more pleasurable is accomplished and sustained by listening itself, an activity that occurs simultaneously with driving and

*Tel.: +1-650-604-2526; fax: +1-650-604-4036.

E-mail address: bclancey@mail.arc.nasa.gov (W.J. Clancey).

¹On leave from The Institute for Human and Machine Cognition, UWF, Pensacola, FL, USA.

is not controlled moment-by-moment through deliberation.

When we attempt to simulate how people actually spend their time and how they interact with their environment, we find that the range of motives underlying human behavior—and consequently the formulation of goals—has been inadequately characterized by problem solving theory (e.g. Newell & Simon, 1972). Furthermore, we find that the coordination mechanism controlling moment-by-moment behavior goes beyond deliberation and learned sequences ('chunks', Newell, 1990), to include ritual forms of behavior and dynamically coupled perceptual–motor activity (e.g. playing a video game). Thus, not all goal-directed behavior is inferred or 'compiled', some actions simply reproduce cultural patterns ('scripts', Schank & Abelson, 1977) and some are attentively and adaptively coordinated without deliberation ('situated action', Suchman, 1987; Agre & Chapman, 1987). This framework provides a much broader understanding of cognition, relating problem-solving theory to two other levels of analysis: The variety of motives by which goals are formulated and how operations fulfilling goals are attentively coordinated.

My goal in this paper is to clarify the theoretical foundations and contributions to cognitive science of Brahms, a language for work practice modeling (Clancey, Sachs, Sierhuis, & van Hoof, 1998; Sierhuis, 2001). In particular, the key representational construct in Brahms, an *activity*, may be easily confused with a *task*, a representational construct that describes human behavior in terms of goals and operators (cf. Clancey, 1992). Two viewpoints need to be addressed: One side argues that there is no distinction to be made—activities are simply tasks by another name, and thus all human behaviors can be simulated within a problem-solving framework (e.g. the formalization of Soar, Newell, 1990). Another point of view claims there is a fundamental difference between tasks and activities, which has already been fully explored and explained by social scientists (e.g. Gasser, 1991; Lave, Murtaugh, & de la Rocha, 1984; Lave, 1988; Luff, Hindmarsh, & Heath, 2000) or in most respects is indistinguishable from activity theory (Leont'ev, 1979). Both viewpoints suggest there is nothing new in Brahms—but the two interpretations are contradictory, and I will show that

both are incorrect. By better articulating the distinction between tasks and activities, we can recognize psychological and social aspects of human behavior that extend both problem-solving and activity theory.

To make the presentation concrete, I draw from field studies of a scientific expedition (Clancey, 2001b). In the setting of the Canadian High-Arctic on Devon Island, geologists, biologists, and computer scientists have built a habitat for living and working called the Flashline Mars Arctic Research Station (the 'FMARS hab'). Operations and experiments in this analog setting will teach us how to live and work on Mars. My ethnographic study over four field seasons focuses now on developing a Brahms simulation of 'a day in the life of the hab'. An analysis of FMARS activities shows how the Brahms activity formalism shares properties of both scripts and activity theory. But constructing a simulation of interacting people, facilities, and tools also involves a new level of detail, an understanding of human *ensemble* behavior, which adds a new element to cognitive modeling.

2. Background: Brahms

We originally developed Brahms in 1992 as a tool for design teams of workers and researchers at NYNEX, a telecommunications company in New York and New England. Existing business process modeling tools (adapted directly from models of manufacturing processes) did not allow representing the informal communications and job sharing NYNEX's ethnographers observed in workplaces (Clancey et al., 1998). For example, it was not possible to represent that a person worked on multiple jobs at one time (just as machines in an assembly line must complete one job before the next arrives). Our task in developing Brahms was to 'make social processes visible', simulating how distributed teams interacted through workflow software and protocols for coordinating organizational roles. The first Brahms models in the 1990s focused on order processing at NYNEX; for example, we simulated how a 'turf coordinator' used teleconferencing to coordinate circuit testing simultaneously between two distant locations.

Previous work in business anthropology and social

studies of work showed that understanding complex human–system interactions requires going beyond formal procedures and information flows to analyze how people interact with each other (Gasser, 1991; Wynn, 1991). Such *work practices* include conversations, modes of communication, informal assistance, impromptu meetings, workarounds, and so on. Accordingly, Brahms models describe the *activities* of people belonging to multiple groups, situated in a physical environment (geographic regions, buildings, transport vehicles, etc.) consisting of tools, documents, and computer systems. The focus is on simulating human behavior, not cognitive processes in the usual sense of structures and processes in the brain. In particular, we interpret model constructs (e.g. situation–action rules) as behavior patterns, not ‘knowledge’.

A Brahms model of work practice reveals *circumstantial, interactional influences* on how work actually gets done, especially how people informally involve each other in their work, thus changing the quality of the result. Building a Brahms model leads human–computer system designers to question how tasks and information actually flow between people and machines, what work is required to synchronize individual contributions, and how tools hinder or help this process (Greenbaum & Kyng, 1991; Corbett et al., 1991; Bagnara et al., 1995; Hutchins, 1995). Accordingly, we have begun to see how information processing models are abstractions that omit essential interactions between people, tools, and facilities—relationships that determine whether a tool ‘fits’ in the workplace.

Our developing experience over almost a decade suggests that Brahms activity models will be especially useful for system requirements analysis, instruction, and implementing software agents. Brahms has been applied at NASA for simulating astronaut activities on the moon, as well as human–robot system design (Sierhuis, 2001).

From our earliest collaborations at NYNEX, involving anthropologists and artificial intelligence researchers, we found that creating Brahms models crystallizes differences between cognitive and social theories of human behavior, and forces us to articulate how different kinds of analyses relate. In particular, we have been struggling for almost a decade to bring together the perspectives of activity theory,

task analysis, situated action, and human factors. This paper revisits and integrates these theoretical backgrounds, using examples of activities in the Flashline Mars Arctic Research Station, which we are currently analyzing and modeling.

3. Background: FMARS

FMARS is a two-story structure, about 8 m in diameter, in which six people live and work. The hab includes private staterooms for sleeping, a common area for working and eating, a laboratory space, a room for storing clothing and preparing for excursions, plus toilet facilities. During the July–August 2001 field season, I was a member of the Phase 2 crew, which lasted 8 days. My observational methods include participant observation, extensive video and time lapse photography, logging events, and interviews (Clancey, 2001b).

After about a week in the hab, I had created a table of about 50 activities, ranging from mundane aspects of personal grooming to tasks such as cleaning the suits we used when making excursions outside the hab (called ‘extravehicular activities’ or EVAs). By grouping and segmenting activities according to the time of day, it became possible to describe the overall day in the life of the Phase 2 crew:

<i>Time</i>	<i>Activity</i>
07:00–09:00	<i>Breakfast</i>
09:00–10:30	<i>Briefing/planning</i>
10:30–15:00	<i>EVA</i>
15:00–15:30	<i>Eat and clean-up</i>
15:30–17:00	<i>Briefing and planning</i>
17:00–20:00	<i>Computer work (e-mail, photo download, software testing); data analysis in lab; napping; equipment cleaning</i>
20:00–21:00	<i>Dinner and clean-up</i>
21:00–24:00	<i>Movies, refreshments</i>
00:00 +	<i>Sleep, reading and writing</i>

This outline is a broad abstraction, averaged over the week, and not a schedule we followed. Nevertheless, the pattern of events can be striking. For example, on three sequential days, the EVA crew stepped into the FMARS ‘airlock’ at 11:05, 11:06, and 11:08. No

procedure required that we do this; it was an emergent product of our intentions, the constraints of getting into suits and fixing radios, and our other habits (such as when we awoke, how long it takes to eat, the duration of the planning meeting, and time to arrange personal gear). For other activities, the times listed varied each day; but relative times, such as when a debriefing occurs after an EVA, are regular. This *chaining of joint activities* dominated the activities of the hab. Individual activities occur mainly before dinner and in the late evening (e.g. someone may take a nap while someone else is preparing dinner and others are working on their laptops).

In attempting to model and simulate a day in the life of a scientific expedition, we find many different kinds of activities. An initial classification may be organized roughly as follows:

1. *Intellectual*: Any form of *inquiry* (Dewey, 1938), manipulating things or ideas. Includes work-oriented problem solving, but also less-directed forms of attention, such as exploration (reading an online newspaper), documentation-art (taking pictures of the hab), and scientific crafts (gluing sensors to rocks).
2. *Interactional*: Acting out a scene from a movie (e.g. the group sings ‘Ack-Ack-Ack’ in unison from *Mars Attacks*); restaging an event for the TV documentary; getting to know someone in a conversation; playing a game of ‘Martian chess’.
3. *Physical/body maintenance*: Grooming, eating, sleeping, etc.

The point of this initial list is to make obvious that not all activities in the hab are focused exclusively on work. After I discuss activity theory, I will present a more elaborate and principled classification of activities based on *motives*. The overall nature of the analysis will show how motives give rise to goals, which are related in different ways to operations (behaviors).

Orthogonal to the kind of activity, at the other end of the motive–goal spectrum, we can identify different ways in which behavior is explained or predicted by goals. We certainly find many examples of goal-directed problem solving in the hab. For example, during the first few days the generator failed often,

and troubleshooting was required to determine the cause. During Phase 2, problem solving was required to plan traverses, remove vehicles from the mud, clean suits, deploy instruments, etc. Such activities may be simulated within the conventional goal–plan framework of task analysis. But other kinds of coordinated behavior are possible.

Rather than being a form of search, as in troubleshooting or planning, in which goals and actions are changing moment-by-moment, an activity may be regulated to *maintain a state*, or perhaps better phrased, to *sustain a mode of interaction*.² Consider, for example, listening to music while taking a nap (Fig. 1). More generally, the goal of an activity may be to sustain a relation between the person and the environment (in the case of napping, the action includes eliminating unwanted stimuli).

The idea that sleeping is a cognitive activity might at first seem bizarre—agreed, sleeping is not a task, but what cognitive processing is involved? This response stems from viewing sleeping as a mental state, rather than an *activity*. When and where does sleep or rest occur? How is the person dressed? How long does the rest occur? Does the person tell others that he/she is going to take a rest? Sleeping is an activity like any other, socially constrained in how it is enacted or not. Thus, holding closed one’s eyes briefly is acceptable during a business meeting, while resting one’s head on the meeting table might suggest a sudden illness and would elicit comment. In this respect, placing ourselves into a sleep posture, doing so while listening to music (Fig. 1), and the duration of the rest are all *norms*, part of the content of a model of *practice*. Here, given that the rest occurs in the workstation area, while a colleague is working nearby, the activity falls in the realm of work practice, rather than being a strictly ‘personal’ matter. In conventional terms, a person has knowledge about when, where, and how to rest. In Brahm we model the napping behavior as a situation–action rule called a *workframe* (roughly put, a workframe for the activity of ‘computer work’ might state, “If

²I view the term ‘state’, imported from electromechanical engineering, as at best an abstraction when applied to human behavior.



Fig. 1. The activity of napping.

you are working on your computer in the workstation area and feel tired, then rest your head on your laptop for a few minutes”).

Another striking example of a goal involving stable self-regulation (Iran-Nejad, 1990) is waiting for someone or something to happen. Examples include waiting in the EVA Prep Room for assistance to remove the suit (Fig. 2) and waiting while sitting at the mess table for other people to arrive at the debriefing meeting.

Waiting is a form of attentive activity. Like other activities, a variety of socially acceptable behaviors may occur (e.g. joking, giving news, telling anecdotes), and other behaviors are unacceptable (e.g. moving around impatiently or angrily asking why other people are keeping you waiting). Waiting is especially interesting because like any activity it is a conceptually coordinated behavior (as reference to social norms indicates), and it is purposeful (usually with the goal of beginning another activity). Yet, moment-by-moment behaviors while waiting are not driven by a search process; there is no immediate problem to be solved (though a subgoal may be formulated and problem solving undertaken if the

process drags on too long). The activity of waiting highlights how the notion of cognition has been biased to focus on a certain type of attention, namely actively seeking a path of action or carrying out a practiced procedure with structured steps. When we view waiting as an activity, we realize it has content, quite unlike the ‘WAIT’ instruction in a computer program.

The activity of waiting shows that the goal–plan view of action is inadequate for describing what people actually do in practice—and this analysis will carry over to activities that accomplish tasks. This point requires some elaboration. To begin, a task analysis suggests that actions are always determined by the relation of a desired end state and the current state (e.g. making replacement liners for the incinerator toilet, checking email to get the weather forecast). The implication is that every action is a *move* from one state to another, in order to achieve a desired state.

Now, strictly speaking, waiting is an action that satisfies a goal. One might say that ‘just waiting’ is a kind of move (a non-operation or ‘no-op’), with the goal to start the next activity (e.g. to get out of the



Fig. 2. The activity of waiting in the EVA Prep Room for assistance to remove the suits.

suit, to have the debriefing meeting). Not doing anything, or perhaps better stated, continuing to sit or stand but not moving elsewhere, is itself an action. So the claim that behavior is driven by a goal remains intact.

This waiting-as-action formulation is fine, but it ignores that people do not put their minds into stasis, they still do something while waiting, whether it is humming a song, talking to a colleague, making plans for later in the day, or whatever. Waiting is an activity subject to norms like any other, not a ‘no-op’. And without work-specific goals to accomplish, other motives will come into play. Behaviors that are contextually constrained do occur—and that’s what we want to capture in the Brahm’s model of hab activities. How do people wait? What do they say to each other? What gets accomplished? For example, waiting for others to arrive, one crew member seizes the opportunity to take a picture that will later appear on the CNN web site (Fig. 2). The motive here clearly transcends the immediate task of the EVA and scientific work. As the author’s smile indicates (Fig.

2, far right), taking photos is part of establishing interpersonal relations and reconfirms a sense of identity of being a member of a team. My relation to the picture-taker changes.

In summary, on the one hand, I point out that actions like waiting are left out of cognitive models because they do not involve composite behaviors that accomplish a work goal. On the other hand, I show that where waiting might be conventionally included, as in business process models of work flow,³ the activity is still not viewed as composite, but just as something that ties up a human resource and takes time. By definition, waiting is ‘off-task’ (the person is waiting *for* work to be done).

In contrast, an activity perspective—describing what people actually do—reveals that people do not wait like horses standing in a field, but engage in a

³For references and detailed comparison of workflow, Brahm’s, and other multiagent systems, see (Clancey et al., 1998).

variety of behaviors that fit within the context. And having described these behaviors, we might no longer view waiting as lost time or inefficient. Rather, we can better recognize the variety of motives that drive human behavior (e.g. hunger, interpersonal relationships, expressing stress and other emotions), and inquire whether waiting has broader implications within the purposes of the setting (e.g. maybe it's good that the tired and sometimes frustrated EVA crew is left alone a few minutes to unwind).

Actions that fit the task-analysis framework, in which every behavior is a move *towards* a work-oriented goal, are defined as 'rational'; everything else has no purposeful explanation, so is ignored by the conventional cognitive modeler. Yet, as I have illustrated, cognition plays a role in how 'non-work' activities such as resting, relaxing, partying, reading, browsing, touring, etc. are carried out. These examples show that the notion of 'goal-driven behavior' needs to be unpacked and differentiated, so that we appreciate not only the motives driving human activity, but different kinds of attentive behavior.

To make this point clearer, I will examine another FMARS activity, and contrast *what people do* with a functional, task-oriented, normative analysis (Vincente, 1999) of what *should* happen.

4. Example: filling a water tank

So far I have illustrated some hab activities that a task analysis would not include or would not represent as structured and meaningful. But the implications are more far reaching. When we study activities that are plainly task-oriented, we find behaviors that are non-optimal from a purely functional view of work. Here we must avoid the temptation to assume workers are at fault (Wynn, 1991). As in the case of off-task behaviors, a broader analysis of activities and human motives suggests that how people accomplish tasks reflects a variety of motives, whose satisfaction makes work tolerable, interesting, learnable, shareable, etc. (Wenger, 1998).

For example, consider the following situation in the FMARS hab. A water tank located eight feet high above the staterooms must be filled every day; a water pipe comes into the hab from a pump located a half-mile away. A radio is used to communicate with base camp to indicate that filling is necessary and then to communicate with the person at the pump site to indicate when to open and close the water valve leading to the hab (see Fig. 3).

A functional analysis of the setting suggests that, assuming the radio remains fixed on a shelf above the workstation area, two people are required within



Fig. 3. The activity of filling the water tank. The communication specialist is standing by the mobile radio; the biologist is above the staterooms, gesturing to represent the amount of water in the tank. Two other crew members are visible, observing and looking out the portal.

the hab to accomplish the task of filling the tank: One person looks into the tank, holds the pipe in place, and reports when the tank is (almost) full. A second person reports when the first person is ready and when the tank is full.⁴ Furthermore, we could automate the entire process within the hab by including a sensor in the water tank that transmits the water level information to base camp and to the pump area (perhaps the base camp would be radioed by a software agent when water is required, and a light at the pump would indicate a full tank). And apparently, everything could be fully automated with appropriate sensors and pumps.

An alternative, complementary approach—before seeking to automate this process or make it more efficient—is to study the practice and ask what is being accomplished by having three or four people involved. Filling the tank might be automated or have no analog on Mars; but what might we learn more broadly about cognition and collaboration from this example, which might in fact generalize and be useful in the design of a Mars mission?

One possibility is that the functional view omitted some constraints, some aspect of the work to be done. This is evidently not an issue here. More likely, a functional analysis examining only the single task of filling the tank, and not modeling the environment, might not take into consideration the other co-located activities occurring in the hab and the value of redundancy in the case of emergencies. For example, if the generator fails, the communications specialist, who is on the radio, will need to drop what he is doing and attend to the power problem. Other people in the setting, who are peripherally aware of what is going on, will be able to take over to complete the water-filling task (a serious matter, for if the water is not stopped in time, it will flood the hab). Thus, an efficiency analysis might take into account a range of scenarios, and not over-commit the crew so that it lacks flexibility in

unusual, high-risk situations. More generally, *peripheral participation* by onlookers promotes learning (Lave & Wenger, 1981), so the group as a whole is more able to reconfigure how the task is performed (this is called ‘cross training’ in office settings). Furthermore, it is not necessary that everyone be doing something essential at every moment. A chance to reflect, to watch and participate in interesting activities (e.g. filling the tank), provides variety to the day and promotes a feeling of being a member of a group that works well together and is competent.

To simulate the filling-the-tank activity in Brahms we need to represent the peripheral crew members as engaging in the same activity, but not having any routine role. These people are not multitasking, in the sense of doing multiple jobs in an interleaved way, but rather attentively *blending* parallel concerns. The onlookers detect the progress of the work and are ready to step in, but they may be doing other things at the same time (e.g. the crew member looking out the window reported later that she was gathering food in the galley, just to her left). The person below the ladder, looking up, is the crew commander. His role may be viewed partly a supervisor, though the practice of filling the tank (as observed on other occasions) does not require his presence.

In Brahms, peripheral participation is simulated through a *subsumption architecture* (Brooks, 1991; Clancey, 1997a, chapter 5) by which a higher-order activity does not ‘call’ or invoke a subactivity like a procedure. Behavior (simulated by *workframes*), inference (*thoughtframes*) and activity-specific perception (*detectables*)—constituting the higher-order activity—may be activated while the lower-order ‘subactivity’ continues.⁵ In particular, we can model ‘filling the tank’ as a *group activity* (inherited by all crew members) that becomes activated and provides a context for individual behavior. Thus individuals in the vicinity who become aware of the filling activity carry out other activities *in the context of* filling the tank. This represents a conceptualization that crew

⁴Fig. 3 provides a fine example of using gestures and artifacts to represent a quantity. By simply positioning his hand, the first person represents the amount of water in the tank that he sees by looking over the edge; no verbal communication is necessary. Everyone else in the room can see the amount and how it changes over time, without anything having to be numerically measured or described.

⁵See (Sierhuis, 2001) for a detailed description of the Brahms language and simulation engine, with detailed examples of models.

members share—‘we are filling the tank now’—even though individually they may be pursuing other goals (e.g. preparing lunch). The Brahms simulation of this blending of conceptualizations and identities is primitive, and its nature is only vaguely developed in neuropsychological theory (e.g. see Clancey, 1999a, 2001a; Sacks, 1987). But awareness that human attention operates in this way plays a major role in the analysis of practice (Clancey, 1997b; Wenger, 1998).

5. Overview of related work

With the FMARS examples in mind, it is now useful to review other research relevant to simulating activities. Recall that the broad goal of this paper is to show that simulating human behavior in terms of activities is different from task analysis. Used in isolation, for designing or understanding work systems, the task analysis perspective fails to observe and formalize certain motives, activities, and interactions:

1. *Motives not related to work proper*, and their resulting goals and activities, are not studied (e.g. activities done for pleasure, such as watching a movie).
2. When problem-solving behavior is observed and modeled, *actions that are not functional* with respect to the goal are not included (e.g. listening to music) or viewed as ‘no-ops’ in which nothing happens (e.g. waiting).
3. Highly interactive, *ensemble behaviors* (e.g. singing as a group), with which deliberation interferes, have not been extensively studied.

The examples of napping and waiting, reading for pleasure, playing a musical instrument, etc. reveal that the concern of cognitive science has been biased by the notion that intelligence is concerned with solving problems (cf. Gardner, 1985). As has been said in different ways since the mid-1980s, problem solving behavior is left hanging in mid-air, not properly related to human motives (especially how goals and solution methods are formulated) or to human action (how behavior is coordinated moment-by-moment).

What makes this critique so large and complex is the range of concerns, from social (focusing on motivation) to neuropsychology (focusing on conceptual coordination). There is no clear consensus of the meaning of ‘situatedness’ or ‘embodiment’ in the cognitive science community because the implications are all pervasive, including neuroscience (Edelman, 1992), robotics (Steels & Brooks, 1995; Clancey, 1997a, Part II), education and instructional design (Eckert, 1989; Bamberger, 1991; Schön, 1987; Clancey, 1998, 2001c), work systems design (Greenbaum & Kyng, 1991; Clancey et al., 1998), and organizational theory (Clancey, 1995; Wenger, 1998; Brown & Duguid, 2000). Research in this arena is marked by multidisciplinary analyses, most notably involving anthropology (Sacks, 1987; Hutchins, 1995; Clancey, 2001b).

Simulating activities in natural settings, such as the FMARS hab, provides a special opportunity to relate social and physiological concerns to problem solving, as I have begun to do in the examples above. Furthermore, as I will show, simulating behaviors such as gathering for a meeting, following people, representing by gesturing, collaboratively planning a traverse, etc. enables us to make concrete and formalize arguments about *situated action* (Suchman, 1987), which may have seemed vague or counterintuitive without formal modeling to explicate the challenges and insights (essentially Vera & Simon’s (1993) complaint).

Bits and pieces of my claims about activities have surfaced in various ways over several decades:

- Schank and Abelson’s (1997) representation of activities as *scripts* emphasizes that some behaviors are cultural rituals, not deliberately planned or even ‘compiled’ from previously constructed plans (e.g. having a waiter show you to a table versus finding an empty table yourself).
- Agre and Chapman’s (1987) Pengi simulation demonstrated how interactive behaviors in video games occur as tight coupling of perception and action, using physically grounded representations. A variety of situated robots (Steels & Brooks, 1995; Clancey, 1997a, Part II) demonstrate architectures for perception-motor coupling (which builds on original cybernetic insights about feedback control).

- Leont'ev's *activity theory* (1979) emphasizes that behavior is comprehensively organized by motives, which are not only intellectual or labor-oriented, but include hunger and fatigue. Human factors research (Kantowitz & Sorkin, 1983) has emphasized the influence of these motives on human performance in work situations.
- Lave et al.'s (1984) analysis of arithmetic calculations in natural settings (e.g. grocery shopping) shows how properties of artifacts provide resources for making comparisons without using school-taught formal methods.
- Barker's (1968) analysis of *behavior settings* reveals a broad variety of places and action patterns, with a variety of motives including recreation, religious worship, and aesthetic pleasure.
- Maue's (1979) formalization of ensemble behavior (inspired by the music of John Cage), shows a remarkable, taken-for-granted social understanding of timing and development in everyday interpersonal dynamics.

Remarkably, putting this work together to simulate work practices, we find a wealth of human behavior that has not been viewed as cognitive. Perhaps more surprisingly, we find mechanisms of self-regulation and attention at play (pun intended) that are not included in most cognitive models—because most focus has been placed on modeling *reasoning* rather than *behavior*. Modeling practices in Brahms challenges us to model *reasoning as a behavior* (partly illustrated by gesturing in the water-in-the-tank example), and we are far from doing this in realistic detail. As the FMARS examples show, by modeling behavior comprehensively, over the course of a full day, we discover phenomena that we would almost certainly have omitted from a 'work flow' model of the hab, and thus not recognized different ways in which cognition operates.

Most centrally, as situated cognition research has emphasized from the start (e.g. Lave, 1988; Schön, 1987), we find that how goals are formulated and problems resolved are inadequately described and explained by problem-solving models that do not represent the broader activities in which problems arise and are pursued. Brahms research goes beyond including the physical context to include other

activities occurring nearby in parallel, as well as the effect of multiple conceptions of activity within a single person (e.g. the activities of 'being a scientist on a NASA expedition'; 'being a member of the FMARS crew'; 'being the lead biologist on the Phase 2'; 'doing an EVA that will play well on TV'). In particular, in modeling activities we are made aware that we are modeling *identities*—a social-psychological mixture of goals, responsibilities, and norms (e.g. see my analysis of Harold Cohen as a computer artist, Clancey, 1997a, chapter 1).

Furthermore, in simulating practice, we are pressed to move from an *idealized inventory* of the knowledge required to solving problems (the knowledge engineering approach) to understanding *whose knowledge is brought to bear* in actual situations and what cultural and physical constraints affect who participates in activities during which problem solving occurs (cf. Fig. 3). Thus, a proper explanation of the *quality of work* and how to improve it requires moving beyond idealized goal–plan formulations of reasoning to models of settings, activities, and group interactions that affect what (whose) knowledge is engaged in given circumstances. In many respects, this is a rewarding, productive reformulation, for it respects both the social scientist's concern with settings and identity, as well as the cognitive psychologist's concern with domain models and inferences.

In describing 'how work actually gets done' we are not just giving details or showing how procedures are not followed or worked around, etc. We are led to explain how joint behavior is possible at all. A restaurant script (explicated below) may say that the waiter takes patrons to their table, but a restaurant simulation must model the patron agents as recognizing that the person standing before them is a waiter, break off whatever else the group was saying when it entered the restaurant before the waiter approached, and engage in the odd *passive, but goal-directed behavior* of following someone to an unknown location (categorized abstractly as 'our table'). Thus, a multiagent simulation of activities requires consideration of matters of attention, agreement to engage in a mutual behavior, and coupled perception–action sequences that 'in the head' models of cognition do not explore. Nevertheless, a happy realization in simulating work practice is that problems arise all the time in the everyday course of

events, and the formulation of goal-directed reasoning and planning may be adapted to activity models.

In subsequent sections, I will briefly review a variety of related work, elaborating the notion of activities as I go along.

6. Activity theory

Leont'ev's (1979) formulation of activity theory explicitly relates problem solving research to a broader framework of human motives. Rather than giving an independent exposition of activity theory, which may be readily found elsewhere (Wertsch, 1979; Lave et al., 1984; Nardi, 1996), I will present the idea by drawing similarities and contrasts with the activity construct in Brahms. However, I emphasize that my theoretical understanding of *activities* and how they are simulated in Brahms are not the same.⁶

Activity theory, as described by Leont'ev and others, has been applied more to work ('labor') settings than to groups of people living together, as in FMARS. Furthermore, activities described by social scientists have not been simulated heretofore in a comprehensive model of attention, agent movements, communication, and an active environment of objects (e.g. computer systems, microscopes, stoves). The careful reader of Leont'ev will notice that activity theorists were writing with a particular background, different from my own. In particular, Leont'ev uses terms, such as 'task', that I believe he would take pains to define if writing for a cognitive science audience today. Similarly, he takes pains to avoid terms that I do not, such as 'behavior', because he was writing against the background of behaviorism, with which he didn't want to be identified. In particular, when Leont'ev (1979) says "we are dealing with activities—not behavior" (p. 45), I believe he is identifying 'behavior' with stimulus–

response theory, which focused only on immediate organism actions, and ignored or denied the nature of goals, motives, mental representations, and tool mediation.⁷

Activity theory was formalized by Leont'ev, as a consolidation and interpretation of concepts introduced by Vygotsky in the 1930s (Wertsch, 1985). Activity theory may be viewed as a response to two competing theoretical extremes, namely the passive, innate response of the organism without subjectivity in *behaviorism*, and the disembodied, contemplative, ego-centric response of the organism in *mentalism*. Activity theory emphasizes that the organism is doing something all the time (the essence of the word 'activity') and that subjectivity is realized within and constructed by *interaction*. Hence, the focus is on describing and understanding practice—action in naturally occurring, social contexts. By the 1980s, the notion developed in cognitive science as *constructivism* (Papert, 1980) and *interactivism* (Bickhard & Terveen, 1995).

Activity theory may be viewed as an early version of what we call situated cognition today:

(The Stimulus → Response scheme) excludes the process that active subjects use to form real connections with the world of objects. It excludes their objective activity Such an abstraction from the subject's activity is justified only within the narrow confines of the laboratory experiment that tries to clarify elementary psychophysiological mechanisms. As soon as one goes beyond these narrow confines, however, its groundlessness becomes evident. This compelled earlier investigators to explain psychological facts on the basis of special powers such as active apperception, inner intentions, etc.—that is, they appealed to the subject's activity, but only in its mystical, idealized form (Leont'ev, 1979, p. 42).

Ironically, the situation–action rule formalism of knowledge-based models of expertise directly pat-

⁶A Brahms model of activity is only an abstraction for some purposes, such as work design and training. A typical Brahms simulation will describe aspects of human behavior, rather than generatively explain why they happen; this becomes obvious when emotional explanations come into play. For example, does a crew member in FMARS turn on music during dinner by being reminded of a restaurant? Or does the silence of the hab give an eerie, isolated and uncomfortable feeling that music relieves?

⁷Not recognizing that there are different ways of attentively coordinating behavior, early reactions to situated cognition (e.g. Vera and Simon, 1993; Clancey, 1993) viewed it as a form of behaviorism, too.

terns the S→R formalism of behaviorism. Knowledge-based research attempted to make the behavior of the agent non-immediate and deliberated. Although it made the S→R chains more complex, introducing mental constructs, it never explained how goals are formulated and achieved through motives, in and through activity (e.g. see Schön's (1979, 1987) analysis of *problem framing*). Perceptual and motor aspects are left hanging as inputs and outputs, so feedback dynamics are under-appreciated. Thus, the conventional task analysis of work is an 'idealized form' of activity, abstracting away what happens in practice, never observing or formalizing key aspects of *coordination* (Agre & Chapman, 1987; Clancey, 1999b). Sadly, for those who like to believe science is rapid and monotonic, Dewey made precisely the same point in 1896 when he criticized S-R theory (Clancey, 1997a, pp. 92–96).

In summarizing activity theory, Wertsch (1985) writes: "Instead of focusing on the study of psychological entities such as skills, concepts, information-processing units, reflexes, or mental functions, it assumes that we must begin with a unit of activity" (p. 199). Wertsch relates this unit of analysis to the idea of a 'frame':

The level of analysis concerned with activities is seldom included in Western approaches to cognitive psychology. Perhaps the construct in contemporary Western social science that is most similar to this level of analysis in the theory of activity is the notion of 'frame' as outlined by Goffman (1974). As in Goffman's analysis, the notion of an activity focuses on socioculturally defined contexts in which human functioning occurs. Among the activities mentioned by Vygotsky's students . . . are play, instructional (formal educational) activity, and labor or work. One of the most important characteristics of an activity is that it is not determined or even strongly circumscribed by the physical or perceptual context in which humans function. Rather, it is a sociocultural interpretation or creation that is imposed on the context by the participant(s) (Wertsch, 1985, p. 203).

A frame is an actor's conceptual organization that

makes experience and behavior coherent. In artificial intelligence research, frames are associated especially with language interpretation and viewed as 'commonsense knowledge', the objects and events of everyday settings.⁸

A fundamental notion of activity theory is that behavior may be described by different levels of analysis: "To say that an individual is engaged in a particular activity says nothing about the specific means–ends relationships that are involved" (Wertsch, 1985, p. 203). Thus, the notion of goals and goal-directedness is a second, *different* level of analysis. An action is associated with a goal, and is general. For example, "the goal-directed action of moving from one point to another could be executed while participating in different activities, such as play, work, or instruction" (Maida, 1992, p. 204).

At the third level of analysis, "an operation is associated with the concrete conditions under which the action is carried out . . ." For example, how a person moves may depend on the distance involved, obstacles, etc. Leont'ev gives an example of the action of 'dismembering an object' via alternative operations of sawing or slicing, depending on conditions.

In summary, the three levels of analysis in Leont'ev's theory of activity are:

- Activity–Motive
- Action–Goal
- Operation–Conditions

Motives, goals, and conditions enable and constrain behavior, but they are fundamentally different in character: Motives are non-conceptual 'forces'; goals are conceptual constructs, and conditions are interactively perceived dynamics in the situation ('affordances', see Clancey, 1997a, p. 252).

An action may be viewed as 'what must be done', while the 'operational aspect' is 'how it can be done'. From the perspective of task analysis, the relation of the action and operation levels is represented as a top-level goal and a plan decomposed

⁸For an excellent discussion of the many historical threads relating the term 'frame', see (Maida, 1992). See also my subsequent remarks in the discussion of scripts about the distinction between a knowledge model and a behavior simulation.

into subgoals and primitive actions (operations). The apparent emphasis in Leont'ev's analysis is not so much that an action requires a specific procedure, but that the action has a *modality* sensitive to the conditions under which the action is done. For example, we may walk gracefully from our seat to the exit in a restaurant, but move with great urgency when there is a fire. Thus the emphasis of operational analysis is not just on the steps of the plan, but how the steps are carried out.⁹

Of special concern for understanding activities is understanding the kinds of motives that drive them: "The main feature that distinguishes one activity from another is its object. After all, it is precisely an activity's object that gives it a specific direction The motive can be either material or ideal Some need always stands behind it." (Leont'ev, 1979, p. 59).

Zinchenko and Gordon (1979) give related definitions (p. 74):

- Motive = "the object that impels the activity, that toward which the goal is directed"
- Goal = "representation of the result of an action"
- Objective conditions = "resources for reaching the goal"
- Task = an action required to resolve the difference between the result of an action and the goal (p. 101).

When an action (e.g. shifting gears in a car) becomes one of the methods for carrying out a higher-order action (e.g. changing the speed of an automobile), "it becomes an operation . . . no longer carried out as a special goal-directed process" (Leont'ev, 1979, p. 64). Leont'ev says that such an operation "'drops out' of the driver's activity entirely", meaning that it is no longer a consciously organized action. This emphasizes that for Leont'ev activities are consciously undertaken; 'goal-directed' refers to conscious, deliberate organization of operations to accomplish an action. Thus, an agent's having a goal ('distinguishing a goal') depends on the agent's experience (knowledge); in contrast, a

task-analytic view says that a behavior is 'goal-directed' even when the action involves a chunked sequence of behavior (as in Soar; Rosenbloom, Newell, & Laird, 1991).

Finally, it should be emphasized that as a conceptual construct, the notion that activity is a unit means that it is 'not additive' and not composed of parts (the actions) (Leont'ev, 1979, p. 61). The levels of analysis have 'different realities', actions are not an 'implementation' of an activity, as in a reductionist model. In particular, taken out of the motive–activity context, goals and actions have no purpose, no meaning. To make this explicit in Brahm's would at least involve making claims about consciousness (a motive "is converted into a *motive–goal* precisely because it is conscious" p. 62), which is beyond our present purposes and understanding.

7. An activity–motive classification

Leont'ev's notion of an 'object' towards which the motive is constructively oriented is useful for sorting out different kinds of motives. By 'object', Leont'ev apparently meant a thing ('material or imagined'). Yet we find some activities are not oriented towards the properties of things, but the properties of *an interactional process*. Furthermore, I suggest that we separate out the self (personal body and identity) as a special kind of object, giving the following three broad classes of motives:

Thing/Idea ↔ Intellectual motive
 Process ↔ Interactional motive
 Self ↔ Physiological/psychological motive

These categories may be illustrated as follows:

- (1) Thing/Idea ↔ Intellectual motive (inquiry)
 - (a) *Attendee* (observer, listener), e.g. classroom lecture
 - (b) *Explorer* (museum, reading, touring, photographing, drawing, exploring, shopping)
 - (c) *Maker* (cook, carpenter, gardener, weaver, farmer)
 - (d) *Problem solver* (researching, modeling (designing, auditing, controlling), inferring, calculating)

⁹Leont'ev says, "The performed action is in response to a task. The task is the goal under certain conditions" (Leont'ev, 1979, p. 63). This is a notion of a task as a *situated action*.

- (2) Process \leftrightarrow Interactional motive (ensemble)
 - (a) Professional Performance: Actor, musician, sports player
 - (b) Informal: Games, an impromptu face-to-face conversation, playing in general, TV channel-hopping, sex
- (3) Self \leftrightarrow Physiological/psychological motive
 - (a) Eating from hunger
 - (b) Exercising in order to feel stronger
 - (c) Resting from fatigue
 - (d) Grooming to avoid disease
 - (e) Relaxing, praying, exercising, etc. to relieve stress

By ‘intellectual’ I mean the usual focus in cognitive science on reasoning about the properties of things (more broadly, ‘systems’) in order to design them, diagnose their behavior, repair faults, predict and control behavior, etc.¹⁰ Within this category, I include crafts (e.g. landscaping), aesthetic invention (e.g. landscape photography as an art form), and forms of learning that vary from very active (visiting a museum) to passive (attending a lecture). The corresponding things and ideas—the objects of the activities—include: a garden, scenes being photographed, museum exhibits, and the subject of the lecture (e.g. Mars).

In contrast, an interactional motive focuses on the qualities of the interactive behavior itself, as a dynamic process; dancing and playing music are paradigmatic examples. In sports, the purpose of course is to win, but one overarching motive during a team game is how the players interact as a whole, forming an ensemble. Attention is focused on smoothness of interaction, responsiveness, handoffs, joint coverage of opponents’ moves, etc. (of course, interactive intensity varies greatly from hockey to baseball). Obviously, in ball games the ball is an object of attention. My point is that players are also conceiving *the quality of how they are relating* to each other and the overall quality of the team, too. Thus, they focus on the process and not just a thing. (No important distinction is intended in separating

professional from informal interactions; it is only a heuristic for calling examples to mind.)

In the third category, the person is the object of attention. The body is obviously an object, but the motive is primarily physiological, a biological need. I’ve qualified the listed activities to emphasize that the motive relates to survival. But as we all know, in many cultures people eat, exercise, groom, etc. for other motives that may be psychological (e.g. stress), treating the body as a thing (e.g. aesthetic adornment) or oriented toward a dynamic interpersonal activity (e.g. sexual attraction).

Where should we place interpersonal relations? In evaluating the quality of interactional processes, the quality of an interpersonal relation may become a focus of attention (e.g. am I a good parent, a good manager, a good neighbor?). Thus, the object of the motive is the self—what kind of person am I? One approach might be to view the three motive categories from the perspective of Freudian psychiatry (1923): (1) ego-dominated interaction, (2) egoless interaction, and (3) reconciling the id, ego, and superego. Not surprisingly, attempting to classify human motives brings into play psychiatric analyses—generally ignored in problem solving analyses—and demonstrates the value of activity theory for relating different branches of psychology.

For the point of this paper, nothing hinges on having included all human motives in this table or even having adequately analyzed the examples (which I have not attempted). Instead, I have a few broader points to make:

- Not all motives are oriented towards things and ideas in the manner emphasized by studies of expertise (e.g. Chi, Glaser, & Farr, 1988). It’s not just that knowledge goes beyond professional expertise to include commonsense, everyday skills and facts (as Lave emphasized). Rather, the motives for learning and doing are not exclusively oriented towards objects-attributes-values in the manner that knowledge representation languages suggest. In particular, in some behaviors attention focuses on the temporal, dynamic quality of the interaction. Models of cognition rooted in reasoning (deliberation) are insufficient for explaining how attention, ongoing conceptualization, and behavior are coordinated (Clancey, 1999b).

¹⁰See (Clancey, 1992) for a systematic analysis of such operations and the broader notion of ‘system’ applied to medicine, electronic troubleshooting, etc.

- Activity theory claims that goals originate in a variety of motives; understanding how people pursue goals (solve problems) depends on understanding their broader motives (purposes for formulating a goal). For example, puzzles like cryptarithmic were often studied to formalize human behavior (Newell & Simon, 1972). However, in general, the motive in doing puzzles is not only to solve the puzzle, but to have a good time, to pass the time in a pleasurable way. The intellectual activity is a stimulus, not the objective (e.g. the solution of a crossword puzzle has no purpose). This point is entirely ignored in studies of cognition that view intelligence as only serving functional, work-related objectives. Psychologists giving puzzles to subjects in laboratories are so caught up in their own sense of what is being measured, they ignore the meaning of puzzles in real life and never study how they are taken up or solved. Attempts to ‘situate’ puzzle solving in physical interaction (e.g. Zhang, 1997), still ignore the motive for puzzle solving and how this may influence puzzle solving behavior (Clancey, 1999b, pp. 186–191).
- A key kind of activity involves *exploration*, without a particular object in mind. Many activities combine both intellectual stimulation with pleasure (e.g. visiting a museum, touring a city, browsing the newspaper). In these activities, the motive orienting the overall behavior is to have a good time, to be intellectually stimulated, to be refreshed by learning new perspectives. Along the way, a subgoal may be formulated (a question about something, a desire to learn about something in particular). But the overall regulation of the activity is to sustain a certain kind of interaction with the environment, a certain state of mind, not to search for something (whether an object or idea). In general, the notion of goal-directed behavior, when viewed from a problem solving framework, does not fit exploration activities, which constitute a very broad class of human behavior (including scientific exploration of the crater on Devon Island). Accordingly, the nature and role of curiosity is ignored in cognitive models—theories of scientific discovery (Darden, 1997) start with data, not looking and manipulating things (Dewey, 1938; Schön, 1987).

To summarize, motives are oriented towards desirable qualities of things/ideas, processes, and the self. Motives are causative; they are how an activity originates, what Leont’ev called the ‘energizing force’. A motive is not another goal, but something not yet conceptual that affects thought:

The origin of an action is to be found in the relationships among activities, whereas every operation is the result of the transformation of an action The formation of the operations occurs in a way that is completely different from that in the formation of the goal (Leont’ev, 1979, pp. 63–64).

I believe here lies the essential reality cognitive psychology is slowly grasping: All aspects of cognition do not fit into the problem solving (goal–plan–operator) framework; other mechanisms are required to understand the formulation of goals. In particular, a motive is a *pre-conceptual* causal influence, not mystical but neurophysiological in character. Bartlett’s theory of reminding makes this explicit in making emotion the organizing influence for re-collecting details in a story (Clancey, 1997a, pp. 47–50; Clancey, 1999b, chapter 8).

A motive orients a person to formulate goals for action, whose manifestation as operations will be regulated by the nature of the motive. Thus, one usually stops eating when no longer hungry, but a fashion model may leave a plateful of food, ignoring her hunger, so that she may make her body an object of desire. Leont’ev’s analysis emphasizes that cultural values and norms create new objects and qualities in our consciousness, such that new motives are socially constructed. One powerful approach is to transform a physiological or psychological need into a need oriented around a commercial object (the ‘objects of desire’ are themselves produced, and not inherently biological; Leont’ev, 1979, p. 50).

8. How operations are dynamically selected and ordered

The third subtopic of this paper’s title is ‘attentive coordination’. I have several times referred to physical behavior that is not directed moment-by-moment

by deliberation and doesn't fit the problem-solving framework of being stepwise-planned to fit a goal. A paradigmatic example is following someone (e.g. the commander on FMARS says, "Come with me, I need your help on the lower deck"). Improvisation in dance, games, sports, conversations, etc. also illustrate how a non-verbal conceptualization can organize behavior.

So what are the ways in which behavior may be organized moment-by-moment? The examples so far suggest that operations may be organized in four ways:

- Conscious deliberation ('problem solving' proper)
- Chunked sequences ('compilation' of consciously selected and sequenced operations)
- Ritualized sequences (culturally reproduced patterns and personal habits)
- Consciously sustained dynamic relations (invariant processes)

In short, consideration of how attention is regulated outside of conscious deliberation or rote practiced behavior reveals a broadly applicable kind of cognition that has been poorly studied. This form of attention regulation sustains an invariant dynamic relation: Rather than achieving a goal state, the motive is to sustain a state. Attention is regulated this way in musical performances, conversations between friends, brainstorming discussions during meetings, and reading for pleasure. Thus, again, many intellectual activities (e.g. exploring a museum, browsing on the web) are not pervasively organized by plan-goal deliberation (though it often plays some role). Instead, attention is regulated by coupled, dynamic relations of perceiving, reflecting, and moving—a interactional 'state' is sustained, rather than pursued by search.

Any given activity commonly involves a mixture of episodes with different forms of attentive coordination. A well-practiced breakdown may occur and be responded to automatically (Schön's reflection-in-action); a transient motive and associated goal may develop and be set aside or handled by brief deliberation; a serendipitous discovery may occur during problem solving, leading to a phase of browsing

exploration. Accordingly, prompted by the situated action analysis, AI researchers in the late 1980s began to investigate 'reactive planning' (Lyons & Hendriks, 1992).¹¹

Crucially, the overall notion of 'pursuing a goal' holds for both problem solving search and dynamic interactions with feedback. (Leont'ev: "All activity has a looplike structure" p. 49.) The distinction is the degree of difference conceived during the interaction: In dynamic interactions, usually only small variances occur, so coordination is not broken, and behavior develops on a non-problematic path. In problem solving, large differences require description of the situation and reasoning, such that a path can be found (involving backing up from false branches). Schön has articulated this spectrum very well as: knowing-in-action, reflection-in-action, conversations with the situation, etc. (Schön, 1979, 1987). I reformulated this from a problem-solving perspective as: doing, adapting, framing, history-telling, and designing (Clancey, 1997a, p. 216).

Notice how the idea of motives or motivation is rarely mentioned in the problem-solving literature. This is because intellectual motivations are taken for granted in the setting (e.g. a study of expertise in work settings). How thing/idea-oriented motives arise *within* activities driven by other motives (e.g. learning in order to become like someone else; Eckert, 1989) is not considered. In effect, tasks are viewed as objectively given, inherent in a situation, rather than the goals of what Lave calls 'interested persons'. That is, goals are adopted, interpreted, and resolved within broader contexts of subjective motives. Social scientists emphasize that these motives are often socially constructed (e.g. matters of identity, as in the example of the fashion model); even survival motives are recognized and resolved as norms within social contexts (e.g. how to behave when exiting a burning building).

¹¹See (Clancey, 1997b) for a discussion of how all plan following is reactive because learning is always occurring (subconsciously). All action is situated. Contrast with: "The most important property of such a plan is that it minimizes the number of occasions when an emergency calling for situated action will arise." (Vera and Simon, 1993, p. 41).

9. Critique of problem solving framework

Lave draws on activity theory to show how knowledge and problem solving in natural contexts draws upon the social settings. In particular, problem solving should be analyzed as an *activity*, instead of disembodied mental manipulations. For example, Lave et al. (1984) provide an activity-setting analysis of grocery shopping, showing how shoppers do not always use school arithmetic calculations. Instead, readily available information, such as the size of packages, enables making value comparisons directly. Lave et al. (1984) wish to emphasize that mathematical competence should not be equated with methods taught in school (and indeed, different methods are advantageous in everyday life).

Although Lave et al. (1984) draw on and cite Leont'ev's activity theory, they say that the notion of motive "appears too abstract to relate to the person-acting-in-context approach" (p. 72). I believe this stems from their focus on the *physical context*, in an attempt to show the influence of the setting, rather than the *conceptual context* (the shopper's ongoing life activities) in which the shopping occurs. For example, is the shopping event they studied a regular weekly visit, a stop coming home from work, or an urgent trip to secure a missing item in a recipe? Is the person hungry, stressed by children hanging on, feeling overweight, etc.? Is this an upscale market, where price comparisons are perhaps irrelevant (everything is expensive)? Or is the shopper relatively wealthy, and buying items in a discount store (where everything is relatively cheap)? Lave et al. (1984) were not making a systematic study of shopping price comparisons, but rather only intended to show alternative, non-formal methods in one illustrative situation.

Lave et al. (1984) remark on the "necessity of analyzing any segment of activity in relation to the flow of activity of which it is a part" (p. 73), but make no mention of the shopper's outside-the-store activity that affords the time—or financial need—to make the price comparisons they observed. I do not take issue with the emphasis on arithmetic activity as an opportunistic practice in which "people and settings together create problems and solution shapes, and moreover, they do so simultaneously"

(p. 94). However, broader social considerations inevitably come into play when studying an entire day in the life of a person, and not just a half-hour inside a grocery store. Because the design of FMARS hab must reflect all motives that will arise during a 500-day mission, we must model at least an entire day, and not just isolated tasks or work periods. And obviously, this is just a starting point for understanding what might happen to six people in such a small space over the course of the entire mission.

In several important respects, my analysis goes beyond Leont'ev's notion of activity and hence extends Lave's criticism of problem solving in a different way. First, Leont'ev's use of the terms 'task' and 'production' (Leont'ev, 1979, p. 46) and his examples suggest an emphasis on *labor*, rather than recreation, art, and exploration in general. In studying work practice today, we are drawn to consider how work settings interweave 'off-task' motives with formal settings and schedules (e.g. a company gym, an art exhibit in the lobby, a school intern program). As Leont'ev would no doubt say, these broader concerns reflect new realities in the workplace, such that our way of analyzing work reflects our cultural formulation of what a workplace should be. It is probably different from a communist factory.

Furthermore, Leont'ev's and Lave et al.'s analyses were not formalized in simulation models. To create a Brahms simulation, we must articulate the *transactional relations* of mental representations, materials in the environment, and the work settings (Clancey, 1997a, chapters 8 and 9). This interaction involves psychological processes, including the motive of the activity, which Lave et al. tend to ignore.

On the other hand, Leont'ev's analysis resembles a problem-solving view of the world, and doesn't examine the variety of human motives and attentive processes either. For example, he says "(the agent) moves toward the object of this need, and it terminates when it satisfies it" (p. 59). Reading a paper for pleasure satisfies a need, but the object of the activity is only conceptual and one doesn't 'move toward it' so much as engage in a dynamic process that *inherently* satisfies the need (reading produces relaxation). Also, in practice, one doesn't necessarily

end an activity when the need is satisfied, so much as other needs become more pressing or external conditions interrupt what one is doing (e.g. the telephone rings). A mechanism of competing activation is involved. The idea of termination of an *activity*, like the process of beginning an activity, is much more subtle than a goal–operation analysis describes. For again, the motive of many activities is not to *move* to a goal, but to *engage in a process* whose sustained enactment is the goal itself (e.g. playing a game of catch, talking to a friend).

10. Behavior settings

Lave et al. (1984) refer to “Barker’s functionalist brand of setting determinism” (p. 73) as contrary to their activity-setting analysis. Nevertheless, from the perspective of other cognitive science research, this is splitting hairs. The notion of behavior settings—a study of *places* where activity occurs—is an important aspect of how *social worlds* are constituted by and influence human behavior, and is well worth our consideration.

Barker’s (1968) study was a survey and inventory in a particular city of hundreds of ‘behavior settings’ (e.g. lumber yards, hallways, carnivals, hayrides) and ‘action patterns’ (eleven kinds of activities that may occur). Barker’s concern was to *study the environment*, in order to understand the ecology of behavior in ‘bounded, physical–temporal locales.’ Like Leont’ev, Lave, and many others, Barker was concerned that a description of behavior that didn’t study the environment provided only an idealized notion of plan–goal relationships—“the structure of the context is dismantled and rearranged; the structure is destroyed” (p. 9). He wanted to understand how settings are physically nested (e.g. work areas inside a pharmacy) and the temporal relations of periodicity and duration of activities. From another perspective, Barker’s study reminds us that behavior is a conceptually organized relation of perceived physical boundaries (places), functional roles, and episodes. In all respects, his view of psychology moves data gathering and analysis from outside the laboratory, to a full view of human enterprises.

Barker’s list of *action patterns* is especially useful

for considering the variety of motives (pp. 55–66) that drive activities. His eleven categories are:

1. Aesthetic (artistic work, e.g. landscaping, singing; personal appearance is a special case)
2. Business (exchanges where payment is obligatory)
3. Education (formal, not incidental)
4. Government (e.g. printing ballots)
5. Nutrition (e.g. eating or serving food)
6. Personal appearance (grooming, getting well-dressed; excludes purely functional dressing)
7. Physical health (e.g. exercise, physical examinations, learning medical skills)
8. Professionalism (paid, not voluntary performance)
9. Recreation (“playing, having fun, reading for enjoyment, relaxing, being entertained . . . guided by the person’s own needs; there is little compulsion . . .” “Immediate gratification must be the primary function”. p. 64)
10. Religious (includes worship, as well as preparing religious materials)
11. Social contact (“interpersonal relations of any kind”)

Activities in a setting often involve a combination of action patterns listed above. Also, a given setting may involve different action patterns at different times (e.g. a school auditorium used for off-hours movies as well as lectures).

In contrast with an analysis of motives, deliberation, attentive coordination, etc., behavior settings are not cognitive distinctions. Behavior settings are *places with purposes*, not activities (e.g. Barker labels ‘nutrition’ as an action pattern). Barker describes settings with illustrative activities (e.g. ‘eating’). Strikingly, the places are all public (e.g. nursing homes are included, but not places in a typical house). His interest is not on modeling individuals (how people spend a typical day or their path through settings), but modeling *socially constructed places* (what happens here?).

Barker’s analysis revealed activities I had not originally considered in the motive classification, such as landscaping and religious worship. He also highlights the importance of aesthetic concerns in

many settings. He emphasizes how simple activities such as dressing combine aesthetic and survival motives. Similarly, going to church often combines religious ideas with sustaining interpersonal relationships (being a member of the community). However, his analysis has limited direct relevance for building an FMARS model because he omits survival motives (e.g. dressing to keep warm) and personal spaces (e.g. places to read privately).

In summary, consideration of settings, motives, and activities shows that there is no simple mapping between what people do, where they do it, and why they do it. This is why I believe the more powerful levels of analysis are to study kinds of motives and ways in which attention is coordinated:

- Contrast motives according to thing/idea, process, and self orientations.
- Distinguish between goals that are invariant (interactional) relations and those that are actual gaps (differences from the current state) requiring changes to the world and/or mental models.

11. Scripts

In the cognitive science literature, the notion of a script developed by Schank and Abelson for natural language processing in the 1970s is probably the closest to the activity construct in Brahms:

Scripts encode culturally shared knowledge of stereotypes actions that occur in socially ritualized activities, such as: going to stores, restaurants, and museums; riding trains and subways; attending plays and banquets; and playing games or driving cars. Scripts are intended to capture situations in which the behavior is so stylized that the need for complex goal and planning analysis rarely arises. People probably acquire scripts through repeated exposure to stereotypic situations. Scripts are knowledge constructs that tell people what can happen in a situation, what events follow, and what roles various people are expected to play in a given social setting.

An example of a prototypical script is \$RESTAURANT, which captures the activities involved in eating in a restaurant The ordinary course of affairs is that the patron enters, is seated, and orders a meal. The meal is then prepared and served, and the patron eats it. Finally, the patron pays the bill and leaves. Each of these activities is described by a stereotyped sequence of events, which prescribes the order in which things happen and the people and objects participating (sic) in the action. Entering the restaurant, looking for a table, walking over to one, and sitting down comprise one such event chain. Each event has resulting states that in turn become the enabling conditions for further events to occur. For example, one must be physically inside the restaurant before one can look for a table. Seeing an empty table enables walking over to it. As a result of walking to a table, one can sit down at it. (Dyer, Cullingford, & Alvarado, 1992, p. 1444)

In view of the misunderstandings about situated action (e.g. Vera & Simon, 1993), it is ironic to find these earlier claims about the insufficiency of the problem solving framework for explaining certain forms of human behavior. In particular, many actions are not chunked by individuals (the result of deliberate plans), but learned and reproduced as cultural patterns (what I called ‘ritualized sequences’ in the typology of Section 8: ‘How operations are dynamically selected and ordered’). Thus few people eat lunch on weekdays at noon specifically because they are hungry. Rather the time of day culturally determines when the activity occurs. Individuals differ in their adherence to schedules for organizing their daily life, but everyone must relate to broad norms (e.g. planning meetings, avoiding traffic). Importantly, these ritualized patterns are widely reproduced despite individual preferences, and not logically planned from anyone’s work goals or other motives.

Scripts were invented primarily as a means of processing ‘text involving stereotypical situations’ (p. 1444). A script ‘provides the context necessary for inferring events not explicitly mentioned in the story’. In Brahms, activities are like scripts, but are

used for *generating* agent behavior in simulations.¹² The locations, movements, and perceptions of activities are modeled independently. Thus simulated behaviors are *interactions*, conditional on circumstances that are liable to change; behaviors are not rote enacted as if following a script in a play. Furthermore, behaviors are blended through higher-order activities that are currently active and modulating how actions are carried out. Thus, for example, how one greets someone already seated in a restaurant depends on whether one is engaged in the activity of a business lunch or celebrating a family event, such as a wedding. In effect, in human understanding there is not a single \$RESTAURANT script, rather a variety of experiences occurring in different contexts. Schank (1982) and his colleagues recognized this in later developing the MOP memory model.

A distinction is drawn in script theory between necessary events, called *scenes* (e.g. entering the station in order to ride a subway train) and optional actions, included as *episodes* (e.g. sitting down on the platform), which do not contribute directly to the accomplishment of the scene (p. 1447). In this respect, scripts represent practices, what people actually do, rather than an idealized, functional view of what is essential to accomplish a task.

Because of the focus on simulation in Brahms, in contrast with text processing, the modeling of activities emphasizes aspects and details of behavior not found in script theory:

- Association of scripts with groups (to allow agents from different groups to have different behaviors in a single setting, e.g. children behave differently from adults in a restaurant).
- Representing duration associated with operations (*primitive activities* in Brahms occur for fixed or probabilistic durations).

¹²Wilensky (1983) applies scripts to planning ‘mundane activities’ within a problem-solving framework (e.g. the goal of sending a birthday card) rather than simulating *life* (e.g. an errand activity or the birthday party itself). ‘Themes’, as in script theory, ‘give rise to goals’ in a Goal Detector, but they are functional (e.g. fulfilling an organizational role). Rather than explaining how goals are formulated (they are ‘noticed’ and reasoned about), the framework emphasizes ‘meta-planning’, e.g. ‘achieve as many goals as possible’.

- Communication, agent movement, and carrying things are operations built into the Brahms language.
- Perception is modeled as *detectables* (patterns associated with actions, which match environment facts and generate agent beliefs; e.g. on opening a box an agent perceives what is inside).
- Representation of sounds broadcast in an area (e.g. a radio call).
- Simulation of object behavior (e.g. how fax machines work, how telephones ring).
- Agents notice when other agents enter or leave a setting (e.g. crew members on the upper deck notice when someone arrives at the top of the ladder from the lower deck).
- Distinction between *workframes* (conditional actions) and *thoughtframes* (activity-dependent inferences).
- Termination and impasse conditions for activities, plus an interrupt-and-resume capability in the simulation engine.
- An activation/priority mechanism for dynamically organizing behaviors hierarchically and sequentially in a subsumption architecture.

Research on interpretation and learning in script theory (and related work on memory modeling) might be useful to incorporate in Brahms. For example, a Brahms model could be made available to a simulated agent, and used as a script for understanding what real people are doing (‘user modeling’). However, an epistemological distinction is drawn between Brahms’ activities and scripts. Scripts are presented by Schank et al. as psychological entities, which explain human ability to understand text; thus they call scripts *knowledge* constructs. Activities in Brahms are a modeler’s descriptions of stereotypic behavior, which are not necessarily consciously articulated by people themselves. Put another way, Brahms activities are a representation of a person’s *conceptualization* of activity, but there is no claim that a person could or would talk about activities in the manner they are articulated in a Brahms model. One might say that a Brahms model represents an agent’s knowledge of how to behave in a socially acceptable manner, but not *the agent’s* (descriptive) *theory* of behavior. As my analysis of Figs. 1–3 should make clear, and indeed

the study of motives reveals, there are many recurrent, regular aspects of behavior that people do not consciously describe or have theories about (though of course any theory, e.g. Freudian psychology, may become part of the culture and begin to play a causal role in consciously organizing behavior).

Another important connection can be made between scripts and activity theory. It was recognized in formulating scripts that there must be some ‘top-level goals’ that cause behavior to occur, otherwise it would not be purposeful:

- *Schank and Abelson (1977) suggest five general categories of top-level goals.*
- *Satisfaction Goals. Basic physical needs, such as hunger, thirst, and fatigue, that arise periodically.*
- *Preservation Goals. The desire to preserve certain key personal states, such as the preservation of life, health, and possessions.*
- *Achievement Goals. Large-scale ambitions accomplished over a long term, such as raising a family or success in a career.*
- *Entertainment Goals. The short-term enjoyment of some activity, such as seeing a movie.*
- *Delta Goals. The acquisition of certain goods, particularly wealth and knowledge (Davis, 1992, p. 1292–1293).*

These top-level ‘goals’ correspond to motives in Leont’ev’s activity theory. Satisfaction and Preservation correspond to what I have called the self orientation. Delta goals relate to the thing/idea orientation. Entertainment goals relate to a process orientation. Achievement goals, as I have discussed, appear to combine self and process orientations.

We can call these top-level conceptualizations ‘goals’, but we still need a theory of motives to explain what causes these goals to be conceived from non-conceptual processes such as feelings and emotions. Some attempt has been made to model emotions (Ortony, 1992), usually as traits influenced by events—again shoehorned into the problem-solving framework (e.g. “representing knowledge about a wide range of emotions”, p. 447). A recent, more comprehensive analysis given by Damasio (1994, 1999) makes emotion not just an evaluation of ideas but a causal part of idea formation.

In developing the FMARS simulation, we are starting to include physiological needs and emotions (e.g. boredom/excitement, flow/frustration) in our agent models. In this respect, we believe we can reduce the need for fixed workframe priorities and predesigned sequences, and develop a more realistic simulation of daily life. Consequently, as we understand better how practices and motives are related, we are coming to better understand the intent of other psychological research, especially the discipline of human factors, and perhaps will subsequently be able to integrate these different modeling perspectives and methods.

12. Ensemble

Once we move beyond modeling reasoning to simulating situated action in everyday contexts, especially as we move beyond ordinary work settings to study how people live, new research problems arise in understanding how coordinated, group behavior develops and is sustained. For example, suppose that that FMARS commander has announced during breakfast that a planning meeting will occur at 09:00 h. What is the process by which people gather for this meeting? How do they wrap up what they are doing? What tolerance is allowed in being late? What will happen if someone does not appear ‘on time’? What will other individuals do while they are waiting at the meeting table? These are all aspects of practice that we may observe and model in Brahms. Importantly, these behaviors are dynamic interactions. Individuals are attending to what others are doing (will the meeting start late today? do I have more time?), and adjusting their conception of the group’s intent and what behavior is acceptable. Thus, norms are not fixed rules and universal patterns, but locally adapted and improvised. For example, as a joke everyone in a restaurant may hide while someone is out of the room—conventions provide opportunities for creative violation of norms in actions that confirm the group’s awareness and endorsement of these very patterns. Modeling such adaptive, interactive effects generally goes beyond what we can handle in Brahms’ current language, as we repeatedly find in ethnographic

observations aspects of human cognition that cannot be modeled today.

One person who studied the dynamics of human behavior is Kenneth Maue, a musician interested in “how whole structures of consciousness are generated” (Maue, 1979, p. 2). Maue claimed that there were “very abstract patterns, which we ordinarily take for granted without noticing them directly, that lend shape to all our ideas and experiences”. He focused on the process of an *ensemble*, “the sharing of cooperative awareness” (p. 13). The idea originates in music ensembles (a trio or quartet presenting a paradigmatic example). But as my example of how the FMARS crew arrives at a meeting illustrates and as Maue (p. 5) says, the idea is very general:

Well-developed ensemble means that the players are interacting with sensitive and supportive teamwork. It means a balancing of the intentions of all the individual players with each other, and with the needs of the group as a whole.

Ensemble can be more or less interactive, demanding more or less internalized skill of the players. We can also notice what kinds of roles are being set up. Does everyone have the same role, or are there differentiated roles? Do different roles fall into a hierarchy? Are people equally involved, or are some more central while others are more peripheral? Are there provisions for players entering or leaving the piece?

I extend the sense of ensemble to mean the interrelationships of all the parts within the whole. That means everything: the time and the place, the circumstances and surroundings.

Maue provides a general description of ensemble, as a kind of procedure, which may be hypothesized as a general conceptual structure that organizes a variety of joint behavior (quoting from p. 6):

1. *There is a beginning.*
2. *The player proceeds in turn.*
3. *That which is done becomes precedent.*
4. *Some acts are unacceptable.*
5. *There is an ending.*

Maue began with the idea of a musical ‘piece’, but

then generalized: “a job is a piece; a single day at work is a piece; a relationship is a piece; a walk after supper is a piece” (p. 7). Thus, a piece is an activity that is conceived as a unit, and whose participants constrain their behavior according to their understanding of how the process is developing.

Maue provides games, which are exercises for creating and exploring pieces. Playing these games gives the player “a sense of beginning; of parts and cycles and transitions; of ending; of sensitive moments such as the beginning of the ending” (p. 7). In particular, “part of doing a piece is finding its own proper length, as a natural expression of the activity itself”. For example, in ending FMARS debriefings an observer is aware of bids to end the activity and of certain kinds of summary statements that remark on the quality of the team’s performance: Should we be proud or humbled? What are the implications for next time? Usually people are finely tuned to the moment when such a meeting ends, and signals such as retrieving personal items from the table or starting to stand may be used to indicate strong desires to reach closure.

Maue studied “the way things occur in time” and was interested in how processes develop in our consciousness. He identified five dynamic stages of pieces:

1. *Gathering: A number of elements gather together, out of some common identity or purpose; a piece of constituent elements is formed.*
2. *Interacting: These elements begin to play with each other, participating in each other’s movements.*
3. *Synthesizing: Now the elements influence each other, alter each other, become each other.*
4. *Transforming: The whole experience begins to become something else; the original idea is becoming something that was not known at the outset.*
5. *Yielding: Both meanings apply; the piece itself ends, and gives birth to an offspring.*

Prototypical examples in FMARS are the top-level stages in a day, especially the meetings, but also improvised joint activities, such as refilling the water tank.

The essence of Maue’s analysis is that activities

are not merely logically derived action sequences, but dynamically conceived as meaningful interactions. Through playful interaction, participants influence each other's understanding and emotions about what is happening and what the overall experience means. For example, an FMARS debriefing is not just a review of events, but an airing of grievances, an exchange of viewpoints, a reinterpretation of the past, and a transformation of future intentions and plans. How the group feels about their overall stay in the hab, about their leader, and about their relations to each other is expressed and changed through the debriefing interaction. Accordingly, the end of an activity is not merely marked by the attainment of a goal, but in the case of meetings and conversations especially, is marked by a lowering of the need to interact (brought on by stress, boredom, concern, lack of information, etc.). In the case of regular meetings without time limits, as in FMARS, the end of the activity is emotionally determined, often marked by an explicit statement about yielding to other interests or planned actions (e.g. "Let's do that together; but first I need to use the toilet"; "Okay, I think I'll get started on the generator").

In summary, an activity cannot be fully described or understood as a cognitive process from a goal-plan analysis alone. One must take into account participants' conception of their relationships, their appraisal of the interaction with respect to group coherence, and their weighing of the meaning of the interaction relative to other activities. These notions are consistent with activity theory, but not emphasized in Leont'ev's (1979) presentation. Attempting to simulate activities in a computer model leads us to focus on dynamic characteristics, highlighted by Maue, such as the gathering and the ending. These are viewed only intellectually in a problem-solving analysis as 'having a goal' and 'accomplishing the goal', not as located experiences of persons and not as negotiations of meaning. Indeed, as I have emphasized, a strict notion of 'the goal of a meeting' is inadequate, for an activity like a debriefing occurs for its own sake, as a ritualized interaction, and is not planned de novo every day or terminated by some single goal-specific condition. Instead, the typical duration of a meeting is a more reliable predictor of when it ends.

13. Flow experiences

The idea of 'flow experiences' (Csikszentmihalyi & Csikszentmihalyi, 1988) is related to the kind of attentive coordination in which perception and action are coupled, so that behavior occurs for long periods without a sense of having a problem. Flow experience tends to involve physically interactive and well-practiced behaviors. Activities in which flow experiences appear to occur spontaneously (40% of reported experiences) include swimming, reading, studying, and praying. Other examples of flow experiences include crafts (knitting, carving, book binding, silversmithing, working in fields), playing musical instruments, and religious ceremonies. I think we can safely add 'editing publications'.

The state of flow is induced when a good fit results from the interaction between two lists of instructions: those contained in the rules of a cultural 'game' (e.g. a tennis match, a religious ritual, a professional activity) and the list of intrasomatic instructions—based on biological predispositions—which constitute the actor's skills A person in flow wishes to do what he or she is doing for the sake of the activity itself, independently of external consequences. (Massimini, Csikszentmihalyi, & Delle Fave, 1988, pp. 65–66)

Of course I don't endorse the term 'instruction' here, which suggests a descriptive statement, as in a computer program. Rather the neurological process involves conceptual schemas, the norms of actions and operations I have been describing in this paper. It is the dynamics of a flow experience that are remarkable: "an equilibrium of challenges and skills, clear goals, and immediate feedback" (p. 32). Notice in particular how the activity becomes its own reward. Thus, engaging in the activity is itself a goal state, which is sustained for minutes or hours on end. Because of the pleasurable nature of flow experiences, people and society tend to replicate them and seek them out. Consequently, one approach for designing the FMARS habitat is to examine when and where flow experiences occur and to what extent the habitat, technologies, roles, and practices facili-

tate them. The hours people spent at their laptops and that the lead biologist devoted to creating microscope images suggest leading candidates for such a study.

14. Related multiagent systems

Other researchers in the past decade have discovered that problem solving architectures must be extended or complemented by additional mechanisms to simulate multiagent activities. Examples include the Phoenix fire-fighting simulation (Cohen, Greenberg, Hart, & Howe, 1989), Soar agents (Tambe, Johnson, Jones, Laird, Rosenbloom, & Schwamb, 1995), and team games (Nakashima, Noda, & Handa, 1996). A pioneering example is the simulation of a video-action game by Agre and Chapman (1987).

In general, this research has emphasized that action in a geographically complex space involves tight coupling of perception and action. Actions are not infrequent transformations of the environment resulting from long and detailed deliberation. Instead, actions are quick, often small transformations whose effects are immediately perceived and interpreted, redirecting the next behavior over the course of at most a few seconds. The time scale and degree of deliberation may be placed on a spectrum: A painter or jazz pianist may make many changes for several minutes without stopping to think about the overall effect and redirecting the effort (a flow experience). A writer may similarly pour out many ideas for an hour or more without having to reflect on the structure of what has been written or to reconceptualize an argument. A manager or team may go away from a meeting with a problem that will take many days or weeks to consider and formalize.

The situated activity simulations of video games and firefighting, like Lave's analysis of grocery shopping, begin with a circumscribed world, in which goals are predefined by the researcher, whether it be 'minimizing purchase cost', 'piloting a fighter jet', or 'playing a board game'. The broader activities in which these goals are carried out are generally not analyzed, and hence the effect of conceptual context (matters of identity, self-interest

and periodic physiological motives (hunger, fatigue) are not taken into account. So we see the incremental progress in research: The 'situated' research of the early 1990s, which criticized cognitive modeling for being disembodied, did not model physiological influences on behavior.

15. Conclusions: making meaning

In developing Brahms to simulate work practice since the early 1990s, our goal has been to describe chronological behavior, especially aspects of human behavior not represented in task models. In analyzing simple activities, such as napping or following someone, we became aware that the notion of goals and actions did not always fit a task analysis, problem solving framework.

So what is the relation of tasks to activities? This is like asking how a functional specification of a circuit relates to how a device behaves under different conditions. In engineering, everyone recognizes that these are different kinds of analysis. To be specific: A task is a specification of work (a function) to be performed. People accomplish tasks in activities. A task is not a naturally occurring phenomenon, but a theorist's formalization of a process. Applied to human behavior, task analysis abstracts and summarizes real world events, usually focusing on the transformation of products ('work flow'). An activity model is also an abstraction, but activities, as a unit of analysis of human behavior, are naturally occurring. Activity theory makes the psychological claim that people conceptualize the temporal, located aspects of their lives in terms of activities (e.g. reading mail, relaxing, going to work, waiting). People may of course abstract their own activities in terms of tasks, especially within the rubric of family and organizational responsibility (a job, an order, a chore). A problem is also a naturally occurring phenomenon, experienced as a conceptual *breakdown* (Winograd & Flores, 1986); a goal is conceived subsequently, *within activity* (Dewey, 1938; Schön, 1987).

Task analysis reduces practice, what people actually do, to descriptions of goals, conditional operators, and problem states. With this focus, task analysis omits 'off task' behaviors relating to emo-

tional motives (e.g. resting, listening to music), as well as the circumstantial detail of how goals are accomplished as located, mediated, inherently social interactions. For example, MYCIN's model of medical diagnosis did not describe a typical physician's work day, how a nurse gathers initial information and does a physical exam, how the physician calls the pharmacy to get drug dosage information, how the clinic negotiates with another hospital to get copies of X-rays, etc. Instead, medical diagnosis was reduced in MYCIN to intellectual goals (determine the cause of the patient's complaint) and reasoning. (Or put another way, medicine as a human practice is reduced to mentally relating symptoms, diseases, and therapies.)

The logical, inference-based view of action breaks down when we attempt to simulate everyday human behavior: Practice includes behaviors not derivable or determined by the avowed goal, such as conventions and habits. Practices also mix different motives. This notion was originally captured in scripts; but that representational framework was used primarily for text comprehension, rather than generation (simulation) of situated action. Memory was referred to as 'dynamic' (Schank, 1982); but behavior was not. And even a recent treatment of 'cognitive dynamics' (Dietrich & Markman, 2000) says nearly nothing about behavior in the world.

The dynamic nature of human behavior becomes clearer when we ask why a particular activity stops. Some activities are simply scheduled for a certain time (e.g. a talk during a conference); in other cases we stop from fatigue, a sense that we have 'just had enough', or because of the developing conception of another motive/activity. In contrast, the goal-driven view of behavior cannot explain why we stop browsing a newspaper (when the goal wasn't to find anything in particular), why a phone conversation ends (when the information desired by the caller was received ten minutes earlier), and so on. Some kind of motive-based activation model, related to emotional and physiological factors, is required to explain why many human behaviors stop and others begin.

Once we move out of the laboratory, leaving behind the experimentalist's preconception of what constitutes a goal, we find that behavior in the world is not controlled by goals but by motives. The idea

that goal satisfaction is the determining force in human behavior is incomplete; a higher-order model of motives and their origins, affecting conceptualization of a 'situation' prior to goal formulation (Dewey, 1938), is required. In reducing motives (and emotions) to conceptual structures (e.g. 'top-level goals'), cognitive science has failed to explain how an individual's conceptual systems develop and interact. As a step in this direction, extending Leont'ev's activity theory, I suggested that motives are oriented towards things (material or imagined) and ideas, interactive processes, and the self. A deeper neuropsychological explanation would refer to the dynamics of conceptual systems (e.g. curiosity, humor, mental fatigue).

In this paper, I have shown that a description of work practice is not just 'how work is done' but a comprehensive story of everything that occurs. The task view of work is especially prone to omit (and in any event is not designed to explain) behaviors caused by emotional-physiological motives. For example, having a snack break after an EVA has motives of fatigue and hunger, with corresponding goals of eating and/or drinking something and attending to something other than the previous activity (the idea of a break). Ironically, in not describing and attempting to explain these behaviors, studies of cognition have implicitly viewed fatigue as a purely physical phenomenon (like sore muscles) rather than a cognitive process (how is it that one is too tired to attend another meeting, but quick to pick up the newspaper and read?). An activation model based on physiological feelings and emotions is more comprehensive than the goal-plan mechanism for explaining moment-by-moment human behavior. This is true even in an office setting, where people apparently vary their activities by the time of day and the day of the week according to their energy and enthusiasm.

Furthermore, modeling the variety of human activities reveals a spectrum of *forms of attention*, including passive waiting, daydreaming, resting, browsing, wandering, exploring, listening, watching, highly interactive perceptual-motor coordination, and causal reasoning. How a conscious goal affects moment-by-moment behavior varies a great deal. A goal may be a sustained 'state' (e.g. waiting), or a *desired state* that drives rapidly changing motions

(e.g. looking for information on the world-wide web) or drives changes within a mental model (e.g. imagining an EVA plan). The moment-by-moment organization of behavior is not always deliberated or packaged up from previously inferred action paths. Other forms of conceptual coordination (Clancey, 1999b) are possible and indeed common. FMARS observations suggest that *exploration behavior* in general is poorly understood. More generally, *ensemble interaction* is especially challenging to simulate because it is based on tacit, non-verbal cues and concepts.

An activity model, unlike a task model, is explicitly oriented towards how people conceive of their behavior with respect to norms. Norms are socially constructed patterns, including aspects of when, where, and how actions may be performed. Other, broader activities in which the person is simultaneously engaged provide a *conceptual context* for regulating how a subactivity's actions are operationalized. This idea—that the context for behavior is a social conceptualization (a conceptualization of acceptability with respect to norms)—was originally misunderstood in early debates about situated cognition. ‘Context’ was viewed as either input information or something exclusively physical. And consequently ‘social’ was viewed as meaning ‘physically with other people’. Instead, a psychological theory of activity was required, something social scientists were not likely to posit, and which psychologists, viewing work in terms of functional (task) models and reasoning behavior, were not able to understand.¹³

In Resnick, Levine and Teasley's (1991) introduction to *Socially shared cognition*, in a section subtitled “The social character of motives and task representations”, activity theory is aptly described as “an antifunctionalist point of view in which intentionality and affect are components of cognitive activity” (p. 4). Intentionality (what Lave calls ‘interest’) involves a goal arising from an authentic need of a *person* (e.g. to be contrasted with a

scientist's conception of a *subject* in a laboratory task). A person is a human being. A person has a life with day-to-day activities, with a variety of motivations including professional standing, family relations, physiological concerns, and temperamental biases (e.g. degree of inhibition, Kagan, 1994). These interests only rarely appear in laboratory analyses of cognition. For example, Burton and Brown (1992, p. 96) detected that children would play a game called WEST, not primarily to win, but also to enjoy the side-effects of certain moves and play configurations. Brown and Burton called this a ‘strategy’, locating it within a goal–plan formulation, and hence still begged the issue of how such goals develop from *competing motivations*. Nevertheless, just mentioning that the subjects broke the rules of the experiment is remarkable.

Part of the difficulty in relating the problem solving framework to activity theory is that social science studies are often couched in terminology that other researchers (including myself) find difficult to unpack. For example, Lave (1991) refers to the “inherently socially negotiated quality of meaning and the interested, concerned character of the thought and action of persons engaged in activity” (p. 67). The shift in perspective is that action is not only (or even primarily) oriented towards things (the products of work), but always has an overarching implication, a meaning, as *a move in a socially conceived world*. From this perspective, human behavior inevitably produces and changes meaning. It operates through ‘non-work’ conceptualizations, visible all around us in facial expression, body posture, and choices of who gets to participate in activities. Meaning making is evident in dynamic social interactions (e.g. in ending a meeting), but also is implicit in every individual act. Whether a behavior is deliberately sensitive to norms or deliberately violates them, it is still conceived with respect to a norm. There is no escaping the relationship to what other people desire or expect, whether it is affirmed or denied. Even in the psychopath, norms are not ignored.

As our understanding of motives and attentive coordination develops, we realize that the Brahms language is still primitive. Because a Brahms model is descriptive (i.e. all constructs are statements in a language), conceptualizations are modeled as

¹³Then the whole debate degenerated into the cognitivist's preferred formulation: “Are there representations in the brain?” with the absurd conclusion, “Clancey... wants to deny any internal representations at all,” (Clancey, 1994). See also (Greeno, 1997).

'beliefs', which may be inferred (thoughtframes), communicated (through different explicitly modeled media, such as radios), or form from perceptions (detectables). Simulated meetings, such as impromptu conversations between individuals, require an explicit handshake protocol. Blending of activities is only hierarchical, through an inheritance mechanism, and does not represent the value judgments individuals make when weighing different identities (and norms for behavior) in new settings. An obvious improvement is to model learning in Brahms, capturing how it occurs in natural settings (e.g. by repeated informal assistance as in the water tank example, Fig. 3). But more broadly, the development of the group's practices, as might occur in a 500 day mission might be considered: "Cognition and communication, in and with the social world, are situated in the historical development of ongoing activity" (Lave, 1991, p. 67).

Coming from my original knowledge engineering perspective, I think the most important idea associated with situated cognition is that conceptual systems (in the brain) must be distinguished from verbal descriptions and diagrams (Clancey, 1997a, 2002 in preparation). This idea applies most directly to the interpretation of Brahms models: They are descriptions of human behavior, not simulations of mental processes. But the idea applies as well to the development of Brahms and situated cognition research itself: We may have a conceptual understanding that is inarticulate and even incoherent, but which guides our research and model development. Accordingly, Brahms is a reification of our understanding. It is not the theory itself, but a research tool for helping us improve and communicate our understanding. Papers like this are required for articulating what the theory is and how it relates to a simulation model and the Brahms language more generally. Strikingly, the theories of activity theory and attention presented here go well beyond what is represented explicitly in Brahms (e.g. we have added hunger to the physical properties of agents, but how shall we represent curiosity?). Whether this gap signifies only a transition (as we improve the Brahms architecture) or is a paradigmatic change in the relation between theory and model in cognitive science is unclear.

In conclusion, by attempting to comprehensively

simulate human behavior as it occurs over the course of a day or more in natural settings, we can now locate conventional problem-solving research between two other levels of analysis—*activities conceived with respect to motives* (intellectual, interactional, and physiological–psychological) and *attentively coordinated action* (e.g. sustaining, rather than pursuing, a goal state). By including behaviors that task analysis leaves out (e.g. resting) or does not view as having composite structure worthy of description (e.g. waiting), we reveal a variety of human experience with origins, means of regulation, and effects not considered in conventional cognitive or business process models. On this basis, we have extended cognitive theory, bridging social and psychological analyses, and thus provided a better basis for designing and evaluating work systems.

Acknowledgements

I thank my colleagues at the Institute for Research on Learning (1987–1997) who introduced me to activity theory and its relevance for work system design—especially Jean Lave, Pat Sachs, and Gitti Jordan. My ongoing interactions with Maarten Sierhuis (RIACS) have been helpful to move the theory forward. *Cognitive Systems Research* reviewers provided essential guidance for formulating this paper. The Haughton–Mars Project is a NASA/SETI collaboration; Pascal Lee is principal investigator. The Mars Society owns and operates the FMARS facility on Devon Island. Boris Brodsky (QSS) is implementing the FMARS model; Bruce Damer and associates (Digital Space) are developing Brahms' 3-D representations of facilities, objects, people, and gestures. Mike Shafto has provided essential support and encouragement at NASA/Ames. Funding for this research has been provided by NASA's Intelligent Systems and Space Human Factors Engineering Programs, and the University of West Florida.

References

- Agre, P. E., & Chapman, D. (1987). Pengi: an implementation of a theory of activity. In *Sixth national conference on artificial intelligence*. San Mateo, CA: Morgan Kaufmann, pp. 268–272.

- Bagnara, S., Zuccermaglio, C., & Stucky, S. (Eds.). (1995). *Organizational learning and technological change, Papers from the NATO Workshop, Siena, Italy, September 22–26, 1992*. Berlin: Springer.
- Bamberger, J. (1991). *The mind behind the musical ear*. Cambridge, MA: Harvard University Press.
- Barker, R. G. (1968). *Ecological psychology: concepts and methods for studying the environment of human behavior*. Stanford, CT: Stanford University Press.
- Bartlett, F. C. (1932). *Remembering: a study in experimental and social psychology*. Cambridge: Cambridge University Press, Reprinted 1977.
- Bickhard, M. H., & Terveen, L. (1995). *Foundational issues in artificial intelligence and cognitive science: impasse and solution*. Amsterdam: Elsevier.
- Brooks, R. (1991). How to build complete creatures rather than isolated cognitive simulators. In Van Lehn, K. (Ed.), *Architectures for intelligence*. Hillsdale, NJ: Erlbaum, pp. 225–239.
- Brown, J. S., & Duguid, P. (2000). *The social life of information*. Boston: Harvard Business School Press.
- Burton, R. R., & Brown, J. S. (1992). An investigation of computer coaching for informal learning activities. In Sleeman, D., & Brown, J. S. (Eds.), *Intelligent tutoring systems*. London: Academic Press, pp. 79–98.
- Chi, M. T. H., Glaser, R., & Farr, M. J. (1988). *The nature of expertise*. Hillsdale, NJ: Erlbaum.
- Clancey, W. J. (1992). Model construction operators. *Artificial Intelligence*, 53(1), 1–124.
- Clancey, W. J. (1993). Situated action: A neuropsychological interpretation (Response to Vera & Simon). *Cognitive Science*, 17(1), 87–107.
- Clancey, W. J. (1994). Comment on di Sessa. *Cognition and Instruction*, 12(2), 97–102.
- Clancey, W. J. (1995). Practice cannot be reduced to theory: knowledge, representations, and change in the workplace. In Bagnara, S., Zuccermaglio, C., & Stucky, S. (Eds.), *Organizational learning and technological change, Papers from the NATO Workshop, Siena, Italy, September 22–26, 1992*. Berlin: Springer, pp. 16–46.
- Clancey, W. J. (1997a). *Situated cognition: on human knowledge and computer representations*. Cambridge: Cambridge University Press.
- Clancey, W. J. (1997b). The conceptual nature of knowledge, situations, and activity. In Feltoovich, P., Hoffman, R., & Ford, K. (Eds.), *Expertise in context*. Menlo Park, CA: AAAI Press, pp. 247–291.
- Clancey, W. J. (1998). Developing learning technology in practice. In Bloom, C., & Loftin, R. B. (Eds.), *Facilitating the development and use of interactive learning environments*. Hillsdale, NJ: Lawrence Erlbaum, pp. 3–20.
- Clancey, W. J. (1999a). Studying the varieties of consciousness: Stories about zombies or the life about us? *Journal of the Learning Sciences*, 8(3–4), 525–540.
- Clancey, W. J. (1999b). *Conceptual coordination: how the mind orders experience in time*. Mahwah, NJ: Erlbaum Associates.
- Clancey, W. J. (2001a). Conceptual coordination bridges information processing and neurophysiology. *Behavioral and Brain Science*, 23(6), 919–922, Special issue on Sleep and Dreaming.
- Clancey, W. J. (2001b). Field science ethnography: Methods for systematic observation on an Arctic expedition. *Field Methods*, 13(3), 223–243.
- Clancey, W. J. (2001c). Is abstraction a kind of idea or how conceptualization works? *Cognitive Science Quarterly*, 1(3–4), 389–421, Special issue on abstraction.
- Clancey, W. J. (2002 in preparation). Modeling the perceptual component of conceptual learning—a coordination perspective. In Johansson, P. (Ed.), *Cognition, education, and communication technology*. Mahwah, NJ: Erlbaum Associates.
- Clancey, W. J., Sachs, P., Sierhuis, M., & van Hoof, R. (1998). Brahms: Simulating practice for work systems design. *International Journal of Human–Computer Studies*, 49, 831–865.
- Cohen, P. R., Greenberg, M. L., Hart, D. M., & Howe, A. E. (1989). Trial by fire: Understanding the design requirements for agents in complex environments. *AI Magazine*, 10(3), 34–48.
- Corbett, J. M., Rasmussen, L. B., & Rauner, F. (1991). *Crossing the border: the social and engineering design of computer integrated manufacturing systems*, Artificial Intelligence and Society Series. London: Springer.
- Csikszentmihalyi, M., & Csikszentmihalyi, I. S. (Eds.), (1988). *Optimal experience: psychological studies in flow in consciousness*. Cambridge: Cambridge University Press.
- Damasio, A. (1994). *Descartes' error: emotion, reason, and the human brain*. New York: Putnam.
- Damasio, A. (1999). *The feeling of what happens: Body and emotion in the making of consciousness*. New York: Harcourt Brace.
- Darden, L. (1997). Recent work in computational scientific discovery. In Shafto, M., & Langley, P. (Eds.), *Proceedings of the nineteenth annual conference of the cognitive science society*. Mahwah, NJ: Lawrence Erlbaum, pp. 161–166.
- Davis, E. (1992). Plans and goals. In Shapiro, S. C. (Ed.), *Encyclopedia of artificial intelligence*, 2nd ed. New York: Wiley, pp. 1292–1294.
- Dewey, J. (1938). *Logic: the theory of inquiry*. New York: Henry Holt.
- Dietrich, E., & Markman, A. B. (2000). *Cognitive dynamics: conceptual and representational change in humans and machines*. Mahwah, NJ: Lawrence Erlbaum.
- Dyer, M. G., Cullingford, R. E., & Alvarado, S. J. (1992). Scripts. In Shapiro, S. C. (Ed.), *Encyclopedia of artificial intelligence*, 2nd ed. New York: Wiley, pp. 1443–1460.
- Eckert, P. (1989). *Jocks and burnouts: social categories and identity in the high school*. New York: Teachers College Press.
- Edelman, G. (1992). *Bright air, brilliant fire: on the matter of the mind*. New York: Basic Books.
- Freud, S. (1923). *The ego and the id*. New York: W.H. Norton.
- Gardner, H. (1985). *Frames of mind: the theory of multiple intelligences*. New York: Basic Books.
- Gasser, L. (1991). Social conceptions of knowledge and action. *Artificial Intelligence*, 47(1–3), 107–138.
- Greenbaum, J., & Kyng, M. (Eds.), (1991). *Design at work: cooperative design of computer systems*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Goffman, E. (1974). *Frame analysis*. New York: Harper and Row.
- Greeno, J. G. (1997). On claims that answer the wrong questions. *Educational Researcher*, 26, 5–17.

- Hutchins, E. (1995). *Cognition in the wild*. Cambridge, MA: MIT Press.
- Iran-Nejad, A. (1990). Active and dynamic self-regulation of learning processes. *Review of Educational Research*, 60(4(Winter)), 573–602.
- Kagan, J. (1994). *Galen's prophecy: Temperament in human nature*. New York: Basic Books.
- Kantowitz, B. H., & Sorkin, R. D. (1983). *Human factors: understanding people–system relationships*. New York: Wiley.
- Lave, J. (1988). *Cognition in practice*. Cambridge: Cambridge University Press.
- Lave, J., Murtaugh, M., & de la Rocha, O. (1984). The dialectic of arithmetic in grocery shopping. In Rogoff, B., & Lave, J. (Eds.), *Everyday cognition*. Cambridge, MA: Harvard University Press, pp. 67–94.
- Lave, J. (1991). Situated learning in communities of practice. In Resnick, L. B., Levine, J. M., & Teasley, S. D. (Eds.), *Perspectives on socially shared cognition*. Washington, DC: American Psychological Association, pp. 63–82.
- Lave, J., & Wenger, E. (1981). *Situated learning: legitimate peripheral participation*. Cambridge: Cambridge University Press.
- Leont'ev, A. N. (1979). The problem of activity in psychology. In Wertsch, J. V. (Ed.), *The concept of activity in soviet psychology*. Armonk, NY: M.E. Sharpe, pp. 37–71.
- Luff, P., Hindmarsh, J., & Heath, C. (2000). *Workplace studies: recovering work practice and informing system design*. Cambridge: Cambridge University Press.
- Lyons, D. M., & Hendriks, A. J. (1992). Reactive planning. In Shapiro, S. C. (Ed.), *Encyclopedia of artificial intelligence*, 2nd ed. New York: Wiley, pp. 1171–1181.
- Maida, A. S. (1992). Frame theory. In Shapiro, S. C. (Ed.), *Encyclopedia of artificial intelligence*, 2nd ed. New York: Wiley, pp. 493–507.
- Massimini, F. C., Csikszentmihalyi, M., & Delle Fave, A. (1988). Flow and biocultural evolution. In Csikszentmihalyi, M., & Csikszentmihalyi, I. S. (Eds.), *Optimal experience: psychological studies in flow in consciousness*. Cambridge: Cambridge University Press, pp. 60–81.
- Maue, K. (1979). *Water in the lake: real events for the imagination*. New York: Harper & Row.
- Nakashima, H., Noda, I., & Handa, K. (1996). Organic programming language GAEA for multi-agents. In *Proceedings of the second international conference on multi-agent systems, December 10–13, 1996, Kyoto, Japan*. Menlo Park, CA: AAAI Press, pp. 236–243.
- Nardi, B. (1996). *Context and consciousness: activity theory and human–computer interaction*. Cambridge, MA: MIT Press.
- Newell, A., & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.
- Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.
- Ortony, A. (1992). Emotion modeling. In Shapiro, S. C. (Ed.), *Encyclopedia of artificial intelligence*, 2nd ed. New York: Wiley, pp. 446–448.
- Papert, S. (1980). *Mindstorms*. New York: Basic Books.
- Resnick, L. B., Levine, J. M., & Teasley, S. D. (Eds.), (1991). *Perspectives on socially shared cognition*. Washington, DC: American Psychological Association.
- Rosenbloom, P. S., Newell, A., & Laird, J. E. (1991). Toward the knowledge level in Soar: The role of the architecture in the use of knowledge. In *Architectures for intelligence: The twenty-second Carnegie Mellon symposium on cognition*. Hillsdale, NJ: Lawrence Erlbaum, pp. 75–111.
- Schank, R., & Abelson, R. P. (1977). *Scripts, plans, goals and understanding*. Hillsdale, NJ: Erlbaum.
- Sacks, O. (1987). *The man who mistook his wife for a hat*. New York: Harper & Row.
- Schank, R. C. (1982). *Dynamic memory: a theory of reminding and learning in computers and people*. Hillsdale, NJ: Lawrence Erlbaum.
- Schön, D. A. (1979). Generative metaphor: a perspective on problem-setting in social polity. In Ortony, A. (Ed.), *Metaphor and thought*. Cambridge: Cambridge University Press, pp. 254–283.
- Schön, D. (1987). *Educating the reflective practitioner*. San Francisco: Jossey-Bass Publishers.
- Sierhuis, M. (2001). *Modeling and simulating work practice*. PhD thesis, Social Science and Informatics (SWI), University of Amsterdam, SIKS Dissertation Series No. 2001-10, Amsterdam, The Netherlands, ISBN 90-6464-849-2.
- Steels, L., & Brooks, R. (Eds.), (1995). *The 'artificial life' route to 'artificial intelligence': Building situated embodied agents*. Hillsdale, NJ: Lawrence Erlbaum.
- Suchman, L. A. (1987). *Plans and situated actions: the problem of human–machine communication*. Cambridge: Cambridge University Press.
- Tambe, M., Johnson, W. L., Jones, R. M., Laird, J. E., Rosenbloom, P. S., & Schwamb, K. (1995). Intelligent agents for interactive simulation environments. *AI Magazine*, 16(1 Spring), 15–39.
- Vera, A., & Simon, H. (1993). Response to comments by Clancey on 'Situated action: A symbolic interpretation'. *Cognitive Science*, 17(1), 117–133.
- Vincente, K. J. (1999). *Cognitive work analysis: toward safe, productive, and healthy computer-based work*. Mahwah, NJ: Erlbaum.
- Wenger, E. (1998). *Communities of practice: learning, meaning, and identity*. New York: Cambridge University Press.
- Wertsch, J. V. (Ed.), (1979). *The concept of activity in soviet psychology*. Armonk, NY: M.E. Sharpe.
- Wertsch, J. V. (1985). *Vygotsky and the social formation of mind*. Cambridge, MA: Harvard University Press.
- Wilensky, R. (1983). *Planning and understanding: a computational approach to human reasoning*. Reading, MA: Addison-Wesley.
- Winograd, T., & Flores, F. (1986). *Understanding computers and cognition: a new foundation for design*. Norwood: Ablex.
- Wynn, E. (1991). Taking practice seriously. In Greenbaum, J., & Kyng, M. (Eds.), *Design at work: cooperative design of computer systems*. Hillsdale, NJ: Lawrence Erlbaum, pp. 45–64.
- Zhang, J. (1997). The nature of external representations in problem solving. *Cognitive Science*, 21(2), 179–217.
- Zinchenko, V. P., & Gordon, V. M. (1979). Methodological problems in the psychological analysis of activity. In Wertsch, J. V. (Ed.), *The concept of activity in soviet psychology*. Armonk, NY: M.E. Sharpe, pp. 72–133.