

3. THEORY OF MODELING WORK PRACTICE

This chapter describes the theory of modeling work practice. I start out, in the first section, with given credence to practice as a valid concept of human behavior, separate from cognitive problem-solving behavior. I discuss a number of views on practice from the research literature. This will give the reader some background for the next section, in which I define the elements of work practice at an epistemological level. These elements are what becomes the driving force in finding a modeling language to represent work practice. In the section before the conclusions, I describe a model-based approach to work practice modeling and the operationalization of a model into a computational form (Sierhuis and Clancey 1997).

3.1 HISTORY OF PRACTICE

In this thesis I take a strong stance by fully adhering to practice as a valid form of knowledge that drives the behavior and actions of people. I, therefore, am of the opinion that we can objectify this knowledge in a knowledge-level representation of practice, much in the same way as AI researchers have created an epistemology of problem-solving knowledge. Ironically, the principle of rationality¹⁴ in AI comes from the Technical Rationality model, and it is this model that I denounce as the only model for defining knowledge (Newell 1982) (Newell and Simon 1972) (Simon 1976). It is this view that has the field of AI ignore the value of practical knowledge.

In this section, I attempt to show that practice is a valid level of knowledge that can be represented, not independent, but complementary to the problem-solving knowledge of humans in organizations.

3.1.1 Practical knowledge as knowing-in-action

Donald Schön offers an approach to an epistemology of practice, based on close examination of what practitioners actually do (Schön 1982). When people talk about *practice*, they often mean the practice of professions that have great social importance, such as medical doctors, lawyers, engineers, architects, et cetera. We even go as far as calling the business of a medical doctor his “practice.” When people talk about the practice of such professionals, they mean the exercise of *professional activity*. People believe that the schools in which these professionals have been taught give them a level of practical knowledge and experience that can be applied to solve daily problems. This view of practice is embedded within the model of Technical Rationality.

According to the model of Technical Rationality, professional activity consists of the application of scientific theory and techniques in problem solving. The knowledge base of a profession is thought to have four essential properties: it is specialized, firmly bounded, scientific, and standardized. This view of professional knowledge forces people, still today, to view practical knowledge—what is known in practice—as the application of professional knowledge, while *practice* is viewed as minor knowledge. Practice is said to be the application of scientific theory. It is said that applied science “rests on the foundation of basic science, and the more basic and general the knowledge, the higher the status of its producer.” (Schön 1982)

Why is the application of scientific theory and techniques to problems in practice the dominant view of professional knowledge? Why do we not put practical knowledge at the same level as professional knowledge? Paraphrasing Schön, the answer lies in the history of Western ideas about knowledge over the last three hundred years. Technical Rationality is the heritage of Positivism¹⁵, and the Positivist’s epistemology of practice. In the history of Positivism, practice is an anomaly. Practical knowledge exists, but cannot be seen as a descriptive knowledge of the world, and therefore is not seen as knowledge whatsoever. By viewing practical knowledge as the knowledge of the relationship of means to an end, the question “How ought I to act?” became a scientific one and the best means could be selected by the use of

¹⁴ Principle of rationality: If an agent has knowledge that one of its actions will lead to one of its goals, then the agent will select that action.

¹⁵ Positivism is the philosophical doctrine that developed in the nineteenth century. It was a social movement aimed at applying the achievements of science and technology to the well being of mankind.

scientific-based techniques. From the perspective of Technical Rationality, professional practice is, therefore, a process of problem solving.

With this focus on problem solving, the problem of *setting* and *situation*—in AI this is referred to as *context*—is ignored. Nevertheless, problems do not present themselves to the practitioner as givens. As Schön writes (Schön 1982):

[Problems] must be constructed from the materials of problematic situations which are puzzling, troubling, and uncertain. In order to convert a problematic situation to a problem, a practitioner must do a certain kind of work. [...] It is this [...] that professionals are coming increasingly to see as central to their practice.

Thus, the model of Technical Rationality leaves out the context of work. The practical knowledge used in performing the work constraint by the context, in which it occurs, is not seen as knowledge. The definition of knowledge in the model of Technical Rationality is incomplete in the fact that it does not view practice as a real category of competence. As Schön says it profoundly (Schön 1982):

If the model of Technical Rationality is incomplete, in that it fails to account for practical competence in “divergent” situations, so much the worse for the model. Let us search, instead, for an epistemology of practice implicit in artistic, intuitive processes which some practitioners do bring to situations of uncertainty, instability, uniqueness, and value conflict.

If we are able to put the Technical Rationality model aside, we come to the realization that practical knowledge is a kind of knowing inherent in intelligent action. Common sense admits that the category of *know-how* is in the action. Meaning that the know-how of workers is revealed in the way they act in problematic situations. It is this know-how that constitutes the practice.

3.1.2 Hermeneutics and work practice

Here I touch upon Heidegger and Gadamer’s philosophy of being and understanding, as it relates to work practice. The main source has been the groundbreaking work of Winograd, and Flores (Winograd and Flores 1986). I do not claim to have a full and complete understanding of either Heidegger’s or Gadamer’s philosophy (Heidegger 1962) (Gadamer 1976), and want to stress that I mainly touch upon their work as it relates to my ideas of what constitutes work practice. Most, if not all, of the credit has to be given to Winograd and Flores, since they explained the importance of hermeneutics¹⁶ for artificial intelligence, and more broadly for system design. It is their thinking that made us, who initially worked on Brahms, realize that if we want to understand the way people work we need to understand how people interact with and interpret the world. Therefore, we need to go beyond a description of individual cognition to a more holistic and social view of cognition as it relates to the way people work.

As Winograd and Flores explain, it was Heidegger and Gadamer who placed the hermeneutic idea of interpretation as the foundation of human cognition. Just as we can ask how *interpretation* plays a role in understanding text, we can ask how it plays a role in understanding the world as a whole. Winograd and Flores put forward four assumptions that, simply put, explain the way humans interpret the world (Winograd and Flores 1986, p. 30-31). Here I relate this, more narrowly, to the way people work, and I postulate the following four worldviews:

1. We are the inhabitants of a ‘real world’ made up of objects bearing properties. Our actions take place in the world.

¹⁶ The science and methodology of *interpretation* of texts, particularly mythical and sacred texts, such as the bible.

This means that the way people work is constrained by the location in which this work takes place. Therefore, if we want to model work practice we need to model the “real world,” its locations and the objects it is made up of.

2. There are ‘objective facts’ about that world that do not depend on the interpretation (or even presence) of any person.

This means that we cannot model a world by just modeling the individual interpretation of that world. We need to separate the different individual interpretations from the “objective facts.” Here is where we get confronted with *solipsism*¹⁷, i.e. the modeler of the “objective facts” is also an individual in the world, and hence also interprets the facts of the world according to his or her subjectivity. However, it is important to make a distinction between modeling the interpretation of an individual in a world, and the interpretation of facts in the world. Both are subjective, but both are necessary if we want to take a holistic view of the way people work. However, we should never forget that this means that our model of work practice is *our* interpretation, and not reality.

3. Perception is a process by which facts about the world are (sometimes inaccurately) registered in our thoughts and feelings.

This seems a trivial point after having made the point that every interpretation is a subjective one. However, the important issue that needs to be emphasized is that people make *inaccurate* interpretations of what they perceive, and that they will *act* according to (inaccurate) interpretations. It is therefore important to not only model the facts about the world, but also each individual’s perception of those facts, since it is their perceptions that make people act independently from each other.

4. Thoughts and intentions about action can somehow cause physical (hence real-world) motion of our bodies.

This means that if we want to model work practice, we need to model physical motion of individuals. We can satisfy this assumption by simply modeling the causal relation between thoughts about action and physical motion, and we do not need to model how this happens in the human body (i.e. the neurophysiology).

These four worldviews are my starting point for talking about work practice as a knowledge-level concept. By defining what this level is about, we will be able to represent our practical knowledge in a computational model in a similar way as we are able to model our problem-solving knowledge at a knowledge-level (Newell 1982).

3.1.3 Understanding context

A broad range of work in psychology and anthropology has shown that to fully understand how people work we need to study context in order to understand the relation between individuals, artifacts and social groups (Leont’ev 1978) (Vygotsky 1978) (Suchman 1987) (Lave and Wenger 1991) (Rogoff and Lave 1984). This chapter describes three approaches to study context—situated action models, activity theory and distributed cognition—that have been fundamental in the development of my theory for modeling work practice. All these three approaches use the notion of *activity* as the central point in the way they analyze the context in looking at human behavior.

3.1.3.1 Situated action models

Situated action models emphasize the emergence of activities within the situation. The focus is therefore on *situated action* or, what I call practice, as opposed to problem solving, which means that it is an inquiry into the everyday activity of persons acting in a particular setting. The analysis of situated action is a moment-by-moment analysis of the interaction between people, and between people and the artifacts used in a particular situation (Suchman 1987). Lave identifies the basic unit of analysis for situated action as the

¹⁷ The theory or view that the self is the only reality (definition in the American Heritage Dictionary, 2nd college edition). Kant called it “a scandal of philosophy and of human reason in general” that no philosopher had been able to provide a sound argument against solipsism.

activity of people as it relates to the setting in which this activity takes place and is constructed at the same time; "The setting both is generated out of [the] activity and at the same time generates the activity" (Lave et al. 1984). A *setting* is the relation between acting people and the arena in which they act, almost like a theatrical play. The *arena* is the physical place, i.e. the geographical space, as well as the institution with its social, political and economical background, like the stage within the theatre.

An important aspect of the focus on the activity of persons acting within an arena is that it forces the analyst to pay attention to the flux of the ongoing activity, the minute-by-minute understanding of a real activity in a real setting (Nardi 1996). One of the interesting notions coming out of situated action studies, put forward by Suchman, is that plans are not the mechanism to action, but that plans are resources for action; a "retrospective reconstruction" of situated action (Suchman 1987). In that sense, I postulate that goals are generated within the activity, as an individual's rationalization of what the intention of the activity is; they are not the conditions of when activities are to take place.

3.1.3.2 Activity theory

Activity theory goes back to the 1920s, and developmental psychology work done in the former Soviet Union. The main developers of activity theory are Vygotsky and Leont'ev (Vygotsky 1978). In activity theory the unit of analysis is an activity. An *activity* is composed of a subject, the object, its actions and operations. A *subject* is the person or group of persons that is engaged in the activity. This makes the analysis of activities focus our attention on one or more people.

An *object* is the objective of the activity as it is held by the subject(s) and motivates them in the engagement. *Actions* are processes that must be undertaken to fulfill the object. Subjects are conscious about the actions to take to accomplish the object of an activity. Actions are more or less synonymous with tasks in cognitive science. The notion of an activity can span multiple actors being engaged together in coordinated actions. The actors engaged together might actually have different, even conflicting objects (Kuutti 1996). This is an important concept for the understanding of what collaboration between individuals is about.

Operations are routinized and unconscious actions. For example, when learning to drive a car with a standard gear, the shifting of the gear is at first a conscious action with an explicit goal. Later on, when we are well versed in driving with a stick shift, shifting gears becomes operational and is not a specific goal-driven process anymore. The difference of actions and operations reminds me strongly of the difference between explicit and tacit knowledge (Polanyi 1983). The important take away point from this is that it seems that activities are decomposed into actions, when the activity is not yet "automatic," while an activity that is already operationalized is not decomposed into lower-level actions, but can be seen as a primitive action.

Another key notion in activity theory is the notion of *mediation* by artifacts (Kuutti 1996). Artifacts include instruments, machines, etc, that mediate activity and are created or used by people to control their behavior. In this sense an activity constitutes the context itself. An activity creates a context through its enactment of actions and operations of the people engaged in the activity, and using artifacts to control their engagement. As such, we can see practice as the engagement in activities over a period of time.

3.1.3.3 Distributed cognition

Distributed cognition is a branch of cognitive science that studies the representation of knowledge both inside the heads of the people, as well as within the artifacts and systems they use. The cognitive system can be seen as an activity in activity theory. For example, Hutchins, in his study of the activity of "flying a plane," describes the cognitive system as the total setting of the cockpit (Hutchins 1995). He takes the *cockpit system* as the unit of analysis and observes the many representations that are inside the cockpit system, yet outside the head of the pilots. By taking this social-technical systems approach he can describe the "cognitive" properties of the system, meaning giving an account of the system's behavioral properties in terms of its internal representations, without saying anything about the processes that operated inside the heads of the individuals within the system.

Thus, distributed cognition moves the unit of analysis to the system as a whole, and analyzes the functioning of the system as a "functional unit," instead of as a cognitive system. In doing this, the emphasis

is on understanding the coordination among the individuals and the artifacts in the system. However, this understanding is created by focusing on the available information in the system, as represented in the artifacts and the heads of the individual. There is less of a focus on the activity and situated-actions as a whole, but more on how the lack of information creates a breakdown in the execution of plans and tasks by the individuals in the system.

One of the limitations of this approach is the necessity of drawing a boundary on the system to be analyzed at the start of an analysis. As opposed to letting the analysis of the setting be the driver in setting the boundary of the system. For example, Hutchins, in his study of the cockpit system, does not take into account the interaction and coordination between the pilots in the cockpit and the other crew and the passengers in the airplane (Hutchins 1995). Neither does he consider the interaction with the control tower and their view of the cockpit system. However, the interesting part of distributed cognitive analysis for getting an understanding of the work practice of pilots is the focus on the “memory” of the system as driving the activities of the pilots. This emphasizes the importance of a total systems view in the understanding of practical knowledge.

3.1.4 Work practice

Many researchers in the social sciences use the word *practice* as if it is a well-defined concept that everyone knows. However, it is difficult to describe what a practice is. People notice when something is not a practice, and can often describe why. It can be said that a group of people has developed a practice, but when asked to describe what it consists of, we find it difficult to describe in words. As such, practice is part of our tacit knowledge (Polanyi 1983).

An ad hoc definition of the word practice is:

Definition 1 (practice) *The (collaborative) performance of collective situated activities of a group of people who collaborate and communicate, while performing these activities synchronously or asynchronously, by making use of knowledge previously gained through experience in performing similar activities.*

In short, practice is doing in action (Suchman 1987). Scientists have described how a practice develops, like Wenger, who defines the creation of a practice as follows (Wenger 1997):

Being alive as human beings means that we are constantly engaged in the pursuit of enterprises of all kinds, from ensuring our physical survival to seeking the most lofty pleasures. As we define these enterprises and engage in their pursuit together, we interact with each other and with the world and we tune our relations with each other and with the world accordingly. In other words, we learn. Over time, this collective learning results in practices, which reflect both the pursuit of our enterprises and the attendant social relations. These practices are thus the property of a kind of community created over time by the sustained pursuit of a shared enterprise.

Everybody knows what Wenger means when he says, “this collective learning results in practices”, but what is it that results? Can it be described? Can it be modeled? To do this we need to be able to describe practice at an epistemological level.

3.1.5 How modeling practice is like Aaron’s drawing

Can there be a model of practice? Is a description of practice equal to practice itself? This is similar to the question; is a description of knowledge equal to knowledge itself? This is a debate in AI that has been going on for many years. Clancey makes an argument that allows us to get away from the arguments for or against this issue (Clancey 1997a). Clancey’s way of describing “the representation problem” allows us to ask the question differently, namely;

How can we create an internal representation of work practice, such that the observer interprets the external presentation of a simulation of a model of this work practice, as a reasonable description of the actual work practice?

Clancey describes Aaron, a robot built by Harold Cohen that creates original drawings. The question asked is; is Aaron an artist, or is Aaron a mere mechanical apparatus that can create drawings in a prescribed mechanical fashion? As Clancey states it:

[...] Cohen's dilemma is to understand the relation between internal descriptions, which he formulates and builds into the program, and outside behaviors, which observers will abstract and interpret in Aaron's drawings. (Clancey 1997a, p. 15)

In a private conversation with Clancey, Cohen revealed that his goal is not to create a robot artist, but to create a minimum representational configuration of drawings that will, when put on paper, be interpreted as an artistic image (Clancey 1997a, p. 16):

In this way, the product (what observers perceive) and the mechanism (what is inside the robot) are distinct.

Similarly, in this thesis the goal is not to create a mechanism for developing work practice, but to develop a representational language and simulation program that produces a model of work practice that is interpreted as such.

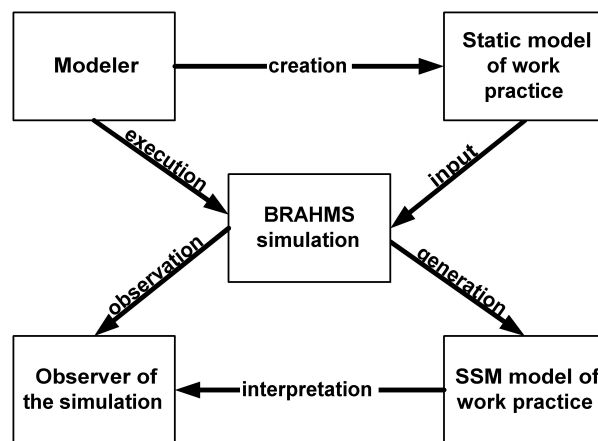


Figure 3-1. Relation of a model of work to a description of the work practice

The modeler in Figure 3-1 develops a model of the work using the representational power of the Brahms language. Model creation is an elaborate process of data collection and work description that leads to a static model of the situated activities of the individuals involved. Using the Brahms simulation program, the model is simulated and a dynamic behavioral model of the work (i.e. a model of the practice) is generation. The observer of the simulation model can observe the model during and after the simulation, interpreting the work practice model.

In the next chapters I define what should be represented in a model of work practice.

3.2 ON THE EPISTEMOLOGICAL LEVEL OF WORK PRACTICE

In the model of Technical Rationality, the notion of a practice is automatically associated with the application of scientific knowledge in “major” professions. Not only am I claiming that practical knowledge is an important category of knowledge, but the concept of work practice allows us to view practical knowledge

within the scope of all kinds of practitioners (not only within those of “major” professions). Here, I focus on creating a framework that allows us to investigate, collect data about, and model the work practices of any group of individuals from any type of profession. Even more so, I focus the attention on work situations where multiple individuals from different professional backgrounds are collaborating. In contrast with Schön’s epistemological model of reflection-in-action in a specific profession, I focus the attention not on the application problem-solving knowledge of an individual, but on the collaboration of activities between individuals.

Work practices is constituted by the way people act and interact in their daily tasks as part of their job, socially and psychologically situated within their environment. It is situated action described in terms of activities and their context. It is how people act and interact in order to accomplish what they have to do. In the next sections, I give definitions of the important elements: community of practice, activity, collaboration, communication, artifacts, and geographical environment.

3.2.1 Community of practice

People who are engaged in a work practice together belong to a *community* that has an identity (Wenger 1997). Together this group of people is engaged in choreographed activities, acting either together or on their own. For example, consider the interplay of activities of people working and dining in a restaurant. There are different roles that are played, the waiters, the chef, the dishwashers, the maître d’homme, et cetera. Even the dinner guests are part of the practice. They all engage in interplay, a kind of theatrical improvisation in real-time. An unwritten play, so to speak, unrehearsed, but still they never forget their lines. They seem to know what the play is about, reacting to each other, never stepping out of character. They all seem to know their parts. They react to and communicate with each other. They have all played their parts before they have ever met each other, because their actions are based on similar previous experiences working and eating in restaurants. This is what the activity of working in a restaurant, and going to eat in a restaurant is all about. It is a *conceptual choreography*. Everyone knows their roles, because they have done it so many times before. They are part of a *community of practice* that exists inside and outside the restaurant. This type of community of practice focuses on a group of people who produce something together.

Definition 2a (community of practice) *A community of practice is a group of individuals, each with different individual skills and knowledge, performing complementary activities while producing something together, that collectively can be seen as a unity within a practice.*

I define a second type of community of practice (see definition 2b). The distinction between the first definition and the second is the type of people that belong to a community. The first definition (2a) includes individuals playing different roles and performing different activities. The second definition of community of practice includes people with similar skills and knowledge, playing the same role and performing similar activities. This type of community of practice includes the professional communities, such as the Java programmers at company X, the architects at company Y, or the group of waiters at a restaurant, et cetera. However, it does not by definition have to be a professional community. For example, we could also talk about the practice of the group of people meeting each other regularly at the water cooler. Such communities are more *informal* or *social*, and do not have to include people from the same professional background. The point is that this definition of community of practice focuses on people that play similar roles and perform similar activities.

Definition 2b (community of practice) *A community of practice is a group of individuals playing similar roles, each with similar skills and knowledge that allow them to perform the same activities, that collectively can be seen as a unity within a practice.*

Both definitions are useful and hold true at the same time. The reason for making a distinction is for the purpose of identifying these types of communities of practice, and the ability to talk about their practice as a whole. For purpose of modeling, it is useful to make a distinction in the practice of a community in terms of different groups of people performing different activities, or in terms of a group of people performing similar activities. By describing a community of practice as a group to which individuals belong, we can represent people’s practice in terms of the sum of the communities (groups) they belong to.

3.2.2 Activities

I now turn the attention to how the behavior of people can be represented as activities. In a knowledge-based system approach, the descriptive modeler's perspective of people's behavior is focused on a narrow description of what people do in terms of tasks and goals. Knowledge modelers start by choosing to model one task and the predefined goals that are to be pursued. With such a design approach, human activity appears to be a relation between goals, data, and decisions. For example in the PEES project, in modeling the "front-door sharing rule" in the Dutch social security law, I choose the interview with the client as the activity of the social security officer (van Dijck et al. 1987). I even focused more narrowly on the client's data specific for making a decision on how much his or her social security check was to be cut, ignoring the actual interview and the setting in which this takes place (Sierhuis 1986). Nevertheless, the social security officer is in the activity of "interviewing the client," as well as at the same time in the broader activity of "working in the social security office." I ignored the context of clients coming and going, colleagues asking for information about cases, people looking for the right forms to be filled out, clients asking for help. In designing the expert system for the "front-door sharing rule," I left out the "work life" of the social security officer. I ignored meetings, discussions in the hallways, the search throughout the building for the correct stamp needed. When I analyzed the task of determining the amount of social security a client would receive, I ignored most of the *activity* of the people in a social security office.

3.2.2.1 Activities versus tasks and goals

Imagine yourself going through a day. There is one response function when you get yourself out of bed, one when you reach for your clothes, one when you face yourself in the mirror, another when you go to breakfast and talk to your spouse, another when you get in your car and drive to work. Each of those situations is radically different and each calls for a quite different function about how to respond to the environment. One involves beds, floors, and covers; another involves mirrors and faucets, another yet something entirely different [...] Describing behavior as multiple response functions implies some sort of decomposition *within* the organism [...] How then should we describe systems? How should we describe their response functions? (Newell 1990, p. 43-44, emphasis added)

These are questions Newell asks in order to describe the foundations of cognitive science. He is interested in describing the workings of the *individual* information processing system (IPS). In other words, the way the individual comes to behave a certain way, or as he says it, "the working of the response functions." This is the individual IPS view he developed with Herb Simon, focusing his theory on *individual problem solving* as the way to describe individual behavior (Newell and Simon 1972).

The theory of humans as an IPS defines problem solving in terms of pursuing pre-specified *goals* in order to accomplish pieces of work that need to be done (i.e. *tasks*). The specification of a goal is a way to make a stated problem actionable, i.e. solvable by means of well-defined decisions. Problem solving is the systematic search over the problem space describing how one can attain a goal. Such an approach is in contrast to a theory for describing *how people actually work* within the constraints of their environment, and how the environment determines their actions and the interactions with other people and artifacts in that environment. Describing the behavior in terms of what actually happens in the world does not lead to a description of the individual's problem-solving behavior. Rather, it leads to a description of the emergent *total system behavior* in terms of the individual interactions, responses to the other elements in the system (people and artifacts), as well as the emergent sequence of individual *activity* (i.e. the state of being active), something Newell calls "microepics."

As is evident in my attempt to make this subtle, but important distinction, the focus in modeling work practice is on the IPS being the *total system*, including the environment, its people, artifacts, places, and time (see chapter 3.1.3). The emphasis of behavior lays at a broader level, namely at a level of *interaction between* discrete entities in the system, each being an IPS in its own right, but influenced by the other elements (IPS's) in the system. Problem solving happens at the individual level, while conceptual construction of activity (i.e. practice) happens at the system level. By describing the individual activity and interactions of elements in the system we can understand the behavior of the total system, as a result of the problem-solving behavior at the individual level. In other words, goals and tasks are being executed within activities,

or better, activities at the meso-level are our social conception of goals and tasks at the micro problem-solving level.

In this view of system behavior, *activities* are socially constructed engagements situated in the real world, taking time, effort and application of knowledge. Activities have a well-defined beginning and end, but do not have goals in the sense of problem-solving models. Instead, the goals are conceptual constructs created and articulated within activities of individual IPS's. Viewing work as activities of individuals allows us to understand why a person is working on a particular task at a particular time, why certain tools are being used or not, and why others are participating or not. This contextual perspective helps us explain the quality of a task-oriented performance.




Figure 3-2. Dimensions of behavior

In this sense, as is shown in Figure 3-2, activities are orthogonal to tasks and goals. While engaged in an activity, people might articulate the task that they are working on, and the goal that they want to accomplish, but these are constructed within the activity. An example of an activity is pursuing a research career. A goal within this activity might be to get a research paper accepted for a conference. A task to reach that goal might be to gather all the relevant literature for the paper. The task and goal are created within the activity, but they are not determined by the activity (Clancey 1997b), meaning that they could similarly arise outside of that particular activity in another. Conceptually we can view activities as the "what we are doing at each moment in time". Goals can be viewed as the "why we are doing what we are doing," while tasks can be viewed as the "how we are doing what we are doing."

To understand activities we must first understand that human action is inherently social. The key is that "action" is meant in the broad sense of an "activity," and not in the narrow sense of altering the state of the world. Instead of viewing "social activity" as something that people do together, such as "socializing at a party" or "the social chat before the meeting," I take a social behaviorist's view. Describing human activities as social means that the tools and materials we use, and how we conceive of what we are doing, are culturally constructed. Although an individual may be alone, as when reading a book, there is always some larger social activity in which he or she is engaged. For instance, the individual is reading the book in his hotel, as relaxation, while on a business trip. Engaging in the activity of "being on a business trip," there is an even larger social activity that is being engaged in, namely "working for the company," and so on. The point is that we are always engaged in a social activity, which is to say that our activity, as human beings, is always shaped, constrained, and given meaning by our ongoing interactions within a business, family, and community. An activity is therefore not just something we do, but a manner of interacting. Viewing activities as a *form of engagement* emphasizes that the conception of activity constitutes a means of coordinating action, a means of deciding what task to do next, what goal to pursue, in other words, a manner of being engaged with other people and things in the environment. The idea of activity has been appropriately characterized in cognitive science as intentional, a mode of being. The social perspective adds the emphasis of time, rhythm, place, and a well-defined beginning and end.

As represented in Figure 3-3, we can be in more than one activity at the same time. While performing one particular activity, we are also engaged in a larger, broader activity. For example, while in the broader activity of working on my dissertation, I am in the middle of the activity of writing the section on activities when my sister-in-law comes in the room to say good-bye. At that moment I suspend the activity of writing

the section, get up and go downstairs to say good-bye, which is the activity that I then engage in for a couple of minutes.

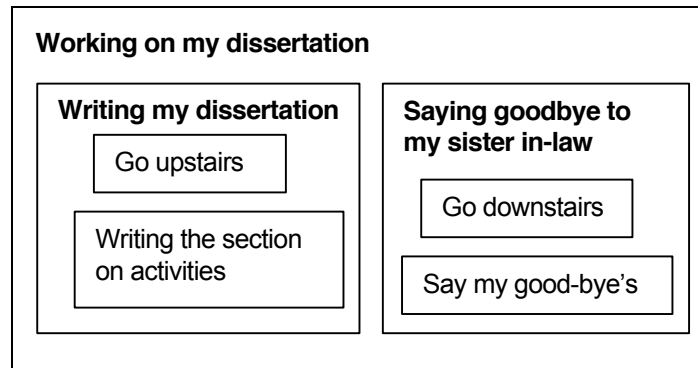


Figure 3-3. Activity subsumption

After my good-bye's I go back upstairs and continue my suspended activity of writing the section of my thesis. While saying my good-bye's I am still in the broader activity of working on my dissertation, otherwise there would be no reason for me to go back upstairs and continue writing. This is *situated action*, an activity that is not fully planned in detail, and can be interrupted and resumed (Suchman 1987); think about putting on your pants in the morning, and the phone rings. While there is not a control program that runs and controls our activities, a situation that suddenly comes up has to be dealt with, without articulated task knowledge. While switching context, the higher-level activity is still being engaged in. Therefore, it is such higher-level activities that constrain us from switching context from one lower-level activity to another lower-level activity and back.

The idea is that humans can control their own behavior—not 'from the inside', on the basis of biological urges, but 'from the outside', using and creating artifacts. (Engeström 1991, p.12)

People choose which activity they engage in, but cannot choose this for others. Therefore, when people suddenly enter our space to interact, we juggle the activities we engage in. We suspend the current activity, start a new one, stop a third one never to come back to it again, et cetera. We act in the situation and react to our environment. This is how the work practice of an organization is formed, and work happens or does not happen. If we are interrupted all the time during our work activities, we start acting a certain way, conscious or unconscious. We might hide, so that interruptions are minimized, or we might just do those activities that do not require a lot of time, or can be interrupted at any moment. In short, the situation and the environment determine our activities, which in turn form our work practice.

Definition 3 (activity) *An activity is a collection of actions performed by one individual, socially constructed, situated in the physical world, taking time, effort, and application of knowledge. An activity has a well-defined beginning and end, but can be interrupted.*

3.2.3 Collaboration

One of the fundamental elements of work practice is the *collaboration* between individuals. An individual rarely works in isolation. Even if we would focus on the practice of one of the major professions, like a medical doctor, an architect or an engineer, we have to realize that they are acting in a context that includes more than just themselves. For instance, the doctor serves patients, and is paid for his services by an HMO. In the office there are physician assistants, nurses, secretaries, et cetera. They are all part of the picture; they collaborate with each other and with each patient that walks through the door. Even when there are no patients there are collaborative activities that take place, such as doing laboratory tests, entering results of tests into the patient's records, calling the pharmacy about prescriptions, and doctors mentoring the physician assistant. In short, the people in the doctor's practice collaborate (Wenger 1997).

Collaboration is a conceptual phenomenon that happens during the collection of activities being performed by the collaborators. Most individuals speak of having “a collaboration” when they feel that the activities engaged in with others is helpful to whatever the objective of the collaboration is. Mead calls it a *social act*, in his point of view of social behaviorism.

A social act may be defined as one in which the occasion or stimulus which sets free an impulse is found in the character or conduct of a living form that belongs to the proper environment of the living form whose impulse it is. I wish, however, to restrict the social act to the class of acts which involve the co-operation of more than one individual, and whose object as defined by the act, in the sense of Bergson, is a social object. I mean by a social object one that answers to all the parts of the complex act, though these parts are found in the conduct of different individuals. The objective of the acts is then found in the life-process of the group, not in those of the separate individuals alone. (Mead 1934, p. 7, footnote 7)

Collaboration can happen when two or more people work together at the same or at different times, being either in the same place or at different places.

The social act is not explained by building it up out of stimulus plus response; it must be taken as a dynamic whole—as something going on—no part of which can be considered or understood by itself—a complex organic process implied by each individual stimulus and response in it. (Mead 1934, p. 7)

In addition, collaboration can happen without people being conscious about it.

The mechanism of the social act can be traced out without introducing into it the conception of consciousness as a separable element within that act; hence, the social act, in its more elementary stages or forms, is possible without, or apart from, some form of consciousness. (Mead 1934, p. 18)

Especially for these forms of collaboration a work practice model could be useful in showing them, making the phenomenon visible and thus explicit. For example, my work practice has changed since I have moved from New York to California. My colleague in New York and I now use e-mail to discuss our research, whereas before we were mostly collaborating face to face, during our daily commute. Our form of collaboration has changed from same-time/same place to different time/different place. It is interesting to observe that we changed our communication tools as well (see paragraph 3.2.4). All this would be difficult to show in a workflow model, but in a model of work practice we include the different geographical places, as well as the different times we are each in our separate activities of reading and replying to our e-mails. The model would also show our new tool for communication (i.e. using e-mail), as well as the information (the stuff we are writing) we are communicating through our e-mails.

Collaboration is a conceptual creation, a state of *mental awareness* by the individuals collaborating. This mental awareness does not necessarily have to exist at the same time, in the same place, and in the same way for every individual in the collaboration. Such awareness is created at the moment we are in our individual activities, making us feel we are collaborating. Collaboration integrates the activities of the individuals in the group, thus establishing a community of practice.

Definition 4a (collaboration) *A collaboration is a collection of activities of two or more individuals, all of them with the mental awareness (being conscious) of working together, either at the same time or at a different time, and either being in the same place or in a different place.*

However, this definition does not capture the fact that people can collaborate even when they are not aware that they are collaborating, i.e. the mental awareness does not exist for them (see the above quote of Mead). I call this *indirect collaboration*. For example, when telephone company sales representatives add

order information to the order databases, this information is used to provision the telephone circuit by the trunk assignor at a different time and in a different office. The two individuals are not aware that they are collaborating when they are in their respective activities. Nevertheless, they are *indirectly collaborating* by the communication of the order information through the order database. Indirect collaboration is sometimes an external observer's conception, however, most of the time the people who engage in such indirect collaboration know that this takes place. It is especially in the breakdown of such collaboration—as when the information in the database is incorrect—that they realize this indirect collaboration they are engaged in.

Definition 4b (indirect collaboration) *An indirect collaboration is a collection of activities of two or more individuals, whom together, without mental awareness (not conscious) of the collaboration, but satisfying their individual goals, using an indirect form of (i.e. asynchronous) communication, either at the same time or at a different time, and either being in the same place or in a different place.*

3.2.4 Communication

Having defined collaboration as a collection of activities, direct or indirect between people, I now turn to how people coordinate their collaboration. The short answer is, through *communication*. In order for two or more people to collaborate they need to communicate. In the Speech Act theory by Searle, the meaning and intent of speech acts are formalized (Searle 1969). Searle describes people's action in terms of sending and receiving speech acts triggering response actions. A *speech act* has at least four distinct types of acts that are all part of the act at the same time (Searle 1969, p. 24-25):

1. Uttering words is performing an *utterance act*.
2. Referring and predicating is performing a *propositional act*.
3. Stating, questioning, commanding, promising, et cetera. is performing an *illocutionary act*.
4. The consequence or effect on actions, thoughts, and beliefs of the hearers is the *perlocutionary act*.

Searle went as far as defining a taxonomy of types of speech acts in which he classified all types as embodying one of five illocutionary points: assertives, directives, commissives, expressives, and declarations (Searle 1975). Speech Act theory analyzes communication in terms of its illocutionary point, - force and propositional content. Using this type of communication analysis, we can model the sequence of communications in a collaboration activity between sender and receiver, as well as the intention and meaning of the speech act. However, in analyzing the way collaboration occurs in practice, we also need to analyze communication in terms of how it actually happens in the real world, thereby modeling collaboration as it really occurs. Speech Act theory abstracts communication in terms of patterns of commitment entered into by the speaker and the hearer. While this is important, in modeling communication as it happens in practice we also need to take into account if a communication activity between two people actually happens, or does not happen. We need to include the *communication tools* used in the speech act, because the type of tool has an impact on when and how the hearer receives the speech act.

Today, communication is more and more efficient and certain communication tools are used globally. Phones, voice mail, e-mail, and fax are communication tools that are more and more taken for granted in the way that we use them. However, it should not be taken for granted that we all have created our own practice around the use of these tools in certain situations. For example, when I work at home I am not checking my office voice mail as often as I should. Without justifying this, it is simply not part of my work practice. Therefore, if someone is trying to contact me, by calling me at my office phone and leaving a voice mail, I might not respond to it for a couple of days. It is not an efficient way of getting a hold of me. Sending e-mail is a better way, since I am constantly checking my e-mail at home. This emphasizes the point that collaboration is very much defined by our practice surrounding our communication tools, and that we, therefore, need to include the use of communication tools in modeling how people actually coordinate their collaboration in the real world. We need to include a model of the workings of communication tools, and how they are used in practice.

3.2.4.1 Content and Information transfer

Speech acts are abstractions of the content of a communication activity between speaker and hearer. For instance, directive speech acts attempt to get the hearer to do something. What is left out is how this collaboration actually takes place. The speaker is in a communication activity, communicating some question or command. The hearer, when receiving the communication, reacts to this communication—based on the illocutionary point—and will perform some activity that ends in a communication activity that communicates the hearer's response to the speaker—the perlocutionary act. In reality, this speech act is a collaboration between two people, and they are using a communication tool, such as a phone, e-mail, or even a face-to-face conversation. The time it takes for this collaboration to complete, and be successful, depends on when the hearer receives the initial communication, and is able to communicate his or her response back to the sender. If the phone rings and the hearer is not in the location of the phone, the communication will not succeed and the speech act will not be completed. If a voice mail is left, the hearer might check it later on and, depending on the message, will either do what is being commanded or will first need to call the speaker back to ask for clarification. This sequence of activities, constrained by the communication tool used, is part of the collaboration between the speaker and hearer, and needs to be taken into account in a model of work practice.

Definition 5 (communication) *A communication is the activity (speech act) of directional transferring of information (in the form of beliefs), held by one individual called the sender, to one or more individuals called the receiver(s), using a specific communication tool (face-to-face, telephone, e-mail, fax, document, etc). After the transfer activity is complete, and successful, the receiver(s) will hold the same information (belief) as the sender of the information, and can now react to it.*

3.2.4.2 Communication tools and their impact on work practice

There are different tools for communication dependent on the location and time spans of the collaborating individuals, having a major impact on the work practice of the group. In one of our investigations, we found that two different groups of workers would use different communication tools for accomplishing the same task. The first group, a group of technicians and a manager, communicated “the assignment of the day's jobs” in a morning coffee meeting. The technicians all come in to work around eight o'clock in the morning. The manager who comes in at seven, will have scheduled the jobs for the day, and will sit with the “force” to have a coffee meeting. During this social gathering, the manager would hand out the job assignments for the day. The second group, consisting of all the same level workers, with one having an acting role as a manager, does not engage in the assigning of the day's jobs during a coffee meeting. Rather, the acting manager assigns the jobs through the job scheduling system. As the workers come in to work, they check their work assignment through the computer. This example shows the difference in the communication activities—consequently the communication tools used—in the practice of assigning jobs for the day. The social interaction and work practice in these two groups is different, which is clearly impacted by the mode (tools) of communication.

It is worthwhile to emphasize that this example shows that the type of communication tools used is an important element in the communication mode, and is one of the defining factors in work practice. Thus, it is important to model the communication tools and its uses in activities, as they define the mode of communication and have an impact on the work practice.

3.2.4.3 Communication effectiveness and efficiency

A communication activity can be seen as simply an information transfer that constrains future actions for the receiver of the information. Either the information is received or not, in which case there is a communication breakdown. A communication activity can be qualified in terms of its *efficiency* and its *effectiveness*. In a communication breakdown its effectiveness is zero. Receiving information means that the information was transferred from the sender to the receiver with an effectiveness of one. Thus, *effectiveness* of communication is a measurement about whether the information is received or not.

Efficiency is a measurement of how many intermediate communication-activities are needed to receive the information. For instance, when the sender uses a telephone as communication tool and the receiver is not

there to answer the phone, the sender can leave a voice mail. When the receiver listens to the voice mail, and the message simply gives the receiver the intended information, the efficiency of the original communication activity was two. This means it took two communication activities for the transfer of the information from the sender to the receiver. If, on the other hand, the voice mail message states for the receiver to call back the sender, and the receiver calls back after which the information transfer takes place, the efficiency of the original communication activity is three, meaning it took a total of three communication activities to transfer the original information.

Using these measurements we can measure the effectiveness and efficiency of a speech act between people, or between people and artifacts (such as a computer system or robot).

3.2.4.4 Same-location communication

When collaborators are in the same location there are a number of *communication modes* they can choose from. If collaborating with just one individual, *face-to-face* communication is used. In face-to-face communication, two individuals are communicating synchronously and instantaneously. For communication with more than one individual, at the same time, a *broadcasting* communication mode is used. The distinction between these two modes is that in a face-to-face communication other individuals around are ignored. In a face-to-face communication people act and react only to the person they are communicating with. When broadcasting, people open the collaborative activity to everyone in the same location, as if speaking to everyone at the same time. The social coffee meeting described above is an example where a broadcasting mode allows the individuals in the group to not only get the information about their own jobs for the day, but also hear what jobs are assigned to the others in the group. Such a communication interaction facilitates learning, because suddenly the job assignment task becomes a social interaction of the group. The individuals in the group can exchange additional information; such as telling a colleague, who was just assigned a job at a location, about the problems at the location.

At the same time, individuals, who are in a location with a collaborating group and are not part of that collaboration, can ignore a broadcasting communication. This means that we, as individuals, can selectively react to communication. People are in control of their own actions; this is part of the meaning of collaboration.

Definition 7a (same-location communication) *Same-location communication is a communication form where the sender and receiver(s) are in the same geographical location. There are two modes of same-location communication, face-to-face communication, and broadcast communication. Face-to-face communication consists of one sender and one receiver. A broadcast communication consists of one sender and multiple receivers*

3.2.4.5 Communication over distance

Communication over distance happens when the communicators are not in the same geographical location. This form of communication can happen in different modes, *same-time communication* (synchronous communication over distance), or *different-time communication* (asynchronous communication over distance). Depending on these two modes, different types of communication technology can be used.

One of the oldest forms of different-time communication over distance is using a messenger who plays the role of a communication device. A more efficient form is mailing or faxing a written document. Alternatively, the use of workflow systems, e-mail or voice mail is becoming increasingly standard. Of course, the telephone is one of the most frequently used forms of same-time communication over distance. As technology is becoming more advanced, different types of communication devices will allow us to collaborate over larger and larger distances, more and more synchronously. As these technologies are being used in the daily work activities, they become a part of the practice.

Definition 7b (communication over distance) *Communication over distance is a communication form where the sender and the receiver(s) are in different geographical locations. In communication over distance, there is a communication device used to communicate. The sender sends the beliefs to the device. The receiver(s) receives the beliefs from the device. There are two modes of communication over distance, same-time communication over distance and different-time communication over distance. In*

same-time communication over distance (direct communication over distance), the sender and receiver communicate instantly or with some short transmit delay, using a communication device. In different-time communication over distance (indirect communication over distance), there is a time span between the sender's communication with the communication device, and the receiver's communication with the communication device

3.2.4.6 Taxonomy of communication types in work practice

From the above description and definitions of communication a taxonomy is presented. The taxonomy also includes possible communication tools that can be used for each type of communication:

Communication

The directional transfer of information from sender to receiver

Synchronous Communication

Same-time communication between sender and receiver

Same-Place Communication

Same-time communication where sender and receiver are in the same location

Face-to-face

Broadcast

Communication over distance

Same-time communication where the sender and receiver are in different locations

Phone-call

Voice-loop

Asynchronous Communication

Different-time communication with a delay between sending and receiving

Same-Place Communication

Different-time communication where the sender and receiver are in the same location

Using Artifacts

Documents

Using Electronic forums

E-mail

Database (or electronic document)

Communication over distance

Different-time communication where the sender and receiver are in different locations

Using Artifacts

Fax

Mailed documents

Using Electronic forums

Voice-mail

E-mail

Database (or electronic document)

3.2.5 Artifacts

People live and act within a physical world. People use and create artifacts in almost all activities that they engage in. When in the activity of hammering a nail, we use a hammer and a nail, and we end up with a nail in whatever artifact we have hammered it in. If we try to understand this activity in context of performing it in the real world, we cannot leave out the artifacts. The artifacts constrain the way we perform activities. It is part of our context, and we have no choice but to interact with the physical world in order to act. We need to include these artifacts into our model of work practice. Leaving them out would miss the opportunity to understand the reason for performing activities. In other words, the artifacts are as important in the work practice as the people are.

Definition 8 (artifact) *An artifact is a physical object in the world.*

George Mead's social-behaviorist notion of *instances of the universal*, as well as Heidegger's notion of *break down* and *readiness-at-hand*, explains the role of physical objects—artifacts—in an activity. Mead, as well as Heidegger, uses the hammer and the activity of hammering as the example in which the hammer is the object that turns into a tool—as an extension of the hand. Mead's idea is that the concept "hammer" is

the universal and the object used in the specific activity is *the instance of the universal*. Therefore, for Mead, the role of the hammer is socially bound to the activity, and is not a property of the object itself. If the person who is hammering uses a piece of wood to hammer in the nail, that piece of wood becomes the instance of the universal during its use in the activity, and thus plays the role of a hammer. In other words, the object is transformed into the *tool* used to hammer in the nail. Heidegger, in essence, says the same. Only he speaks to it through the understanding that objects and their properties are not inherent in the world, but arise only in an event of break down in which the object becomes *present-at-hand*. To the person hammering, the hammer as such does not exist. It is part of the *readiness-to-hand* that is taken for granted in the activity, without the user's identification as an object. It is only in the break down, for example when the person cannot find the hammer when he wants to hammer in the nail, that the object is present for the user. Whichever notion speaks to you, the issue that is important in modeling work practice is *how* the artifact is used and conceptually understood within the activity. Figure 3-4 shows this relationship.

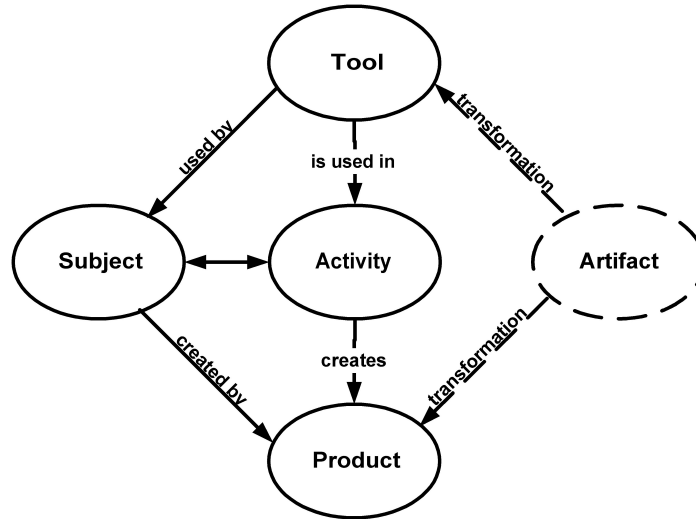


Figure 3-4. Mediated relationship of artifacts in activities

It is the use of the artifact in the activity—its role—that transforms the artifact into a *tool* or a *product* of the activity, used or created by the subject. Outside the activity the artifact is just an object in the world. To the observer the object is necessary for the activity to be performed.

Definition 9 (tool) *When an artifact is being used in an activity, it becomes a tool in the performance of the activity.*

Definition 10 (product) *When an artifact is created or changed in an activity, it becomes a product of the activity.*

3.2.6 Geographical environment

Work is performed within a three-dimensional geographical environment. The restaurant we have dinner at, the office that we work in, and the moon crater the astronauts explore, are all examples of places, spaces, and environments which constraint the way we do our work. The artifacts we use in our work, such as communication- and information tools, are also located in a three-dimensional space. We are constrained to our three-dimensional world, and it defines very much how we can perform our work. For example, when the phone rings, we cannot hear it if we are not in the same room as the telephone. We also cannot observe specific changes in a location when we are not there. For example, if someone turns off the light in a room, and you are not there, you will not observe this and therefore will not be aware of the fact that the light in this room is now off. To show the effect of the environment on the practice, we need to include a model of the geographical environment in a model of work practice.

3.2.6.1 Important aspects of modeling the environment

Modeling geographical spaces is an intricate subject in and of itself. The question we need to ask is; how much and how detailed do we need to model the geographical environment if we want to show its importance to the work practice? The answer is that it depends on the work practice and the geographical space we are trying to model. If we are interested in office work, we need to model the office space in terms of where artifacts are located, such as where the offices of the people are, their telephones, fax machines, computer terminals, meeting rooms, et cetera. When we are modeling astronauts on an extravehicular activity on the Moon, we want to model the traverses, such as which craters they go to, how long it takes to go from one point of interest to another, which rocks are they looking at, and even which soil samples are they taking back with them. We are also interested in how they are traversing, and how long it takes to go from point A to point B. Are they walking or using a moon rover to travel. Are they aimlessly wandering around or are they following a pre-selected route? All these aspects are specified and constraint by the environment and the geographical space in which the work takes place. To give a concrete example of how the geography plays an important part in the way work happens, think about the things that might go wrong during a moon traverse, and how the environment constrains how long you can stay outside on the traverse. How much consumable oxygen do we have to get back to the spacecraft? This is a question that was constantly in the back of the minds of the people at mission control. It defined whether the next activity was to be done or was to be skipped. Dealing with the environmental constraints shapes the work practice.

Definition 11 (geography) *Geography is the description of the physical environment in which the people and artifacts are located when performing their activities.*

3.3 MODEL-BASED APPROACH

In this section I investigate how to operationalize a model of work practice. I use the term *operationalization* to refer to the implementation of a model of work practice that can be executed, i.e. a *computational model of work practice*. In this section of the thesis, I have described a framework for modeling work practice at an epistemological level. Here I investigate how we can implement such a model of work practice. This is the *operationalization problem* (Schreiber 1992). Generally, the term operationalization is used to denote the process of designing and implementing a system. In the context of computational modeling, the operationalization problem includes the ability to execute the model.

In the last decade, model-based development approaches have become the prevailing paradigm in knowledge-based system (KBS) development, as well as in more traditional system development. In KBS development, model-based refers to a development approach in which problem-solving expertise is described (represented) at the *knowledge-level* (Newell 1982) (Clancey 1985). One of the more well known model-based KBS design methodology is the CommonKADS methodology (Schreiber et al. 2000) (Schreiber et al. 1993). The CommonKADS methodology defines a number of design models that allow us to describe problem-solving behavior at Newell's knowledge-level. Much research has been done about how to operationalize KADS models of expertise (Angele et al. 1991) (van Harmelen and Balder 1992) (Karch et al. 1991) (Linster and Musen 1992). In software engineering, model-based refers to a system design approach in which the system is described in terms of a number of well-defined design models (Yourdon 1989), using an object-oriented representation of the system that is being designed (Jacobson 1994). In this chapter I describe what is meant with a model-based approach for modeling work-practice.

We make observations from within our field of reality. A model is a description of that what we observe to exist in the real world. We create models all the time, mental or external, formal or informal. Mental models exist in our minds, and are our interpretation—description—of the world as we experience it. External models are models we create based on our mental models, and therefore, are manifestation of our mental models. In the context of this thesis, all external models are *system models* in the sense that they describe the world in components—*objects*—having properties, mirroring the properties of objects existing in the real world. When we create models that are not physically or geometrically identical with the world we are studying, we have to define system objects with properties that, for the purpose of our study, are similar to the real world objects. Secondly, the relations between the system objects have to be similar to the corresponding real world relations. In algebraic terms, the system objects and their relations have to be *isomorphic* with the real world objects and their relations in the real world.

3.3.1 Formal and informal system models

When we create non-physical, non-geometrical external models, we have a choice of creating these models with a formal representation or, as is often the case, with an informal representation. A *formal model* uses a description formalism that is predefined having a formal syntax and semantics. One of the benefits of a formal model is that the meaning of the model can be formally derived, and there can be no argument about this meaning. However, due to the formality of such models, creating and understanding formal models is often not a simple matter. On the other side of the spectrum, there are *informal models*. Informal models are models that do not have a well-defined meaning. Often the meaning of such models is in the eye of the beholder. Even though the meaning of informal models is not well defined, they can be useful in the understanding of a system. We all know the saying “a picture tells a thousand words.” This also holds for informal models. As such, I feel that the value of informal models is often similar to that of a picture of a scene. It gives context, an external description of reality that can be referred to and shared with others.

One of the benefits of creating external models is their use in analysis and design. External models can be used for explanation of relations and properties of a system that either already exists in the world or is to be developed; in which case the model is the only manifestation of the system.

3.3.2 Computational models

Despite some of the benefits, there are problems with informal models. Informal models cannot be used as a theoretical description of the real world. Therefore, we cannot use informal models to deduce new theorems—propositions about properties of the model. If we cannot do that, we cannot use the model to test hypothesis about properties of the world being modeled.

I distinguish two formal aspects of a system, namely a *structural* aspect of the system and a *behavioral* aspect of the system. Computational models are models that show the behavioral aspects of a system, by *simulating* the behavior of the system over time. This is in contrast with static models, which only show the structural aspects (i.e. the system elements and their relations at one moment in time). As the complexity of a system increases, understanding how the system changes over time—its behavior—becomes increasingly difficult. This is especially true for non-linear systems. A computational model allows us to *observe* the result of changes in the system as time moves forward.

A second problem with informal models is that they cannot be made computational, in the sense that they cannot be executed. Static models can only describe a system at a particular moment in time. They are a static representation of the interpretation of the modeler at the moment the world was interpreted. It depicts the model at a specific time slice. If a model is static it cannot be used to describe the changes over time of the world being modeled. In the case of modeling the work practice of a human activity system (Checkland and Scholes 1990) this is problematic. A static model could describe static properties of a system, but it fails to describe how dynamic properties change over time. Many elements of work practice contain *dynamic relations* between system objects, such as activities being performed by people, communications between people, changes in the environment, et cetera, et cetera. In other words, time is an important independent variable on which a lot of other variables depend. Therefore, if we want to model the work practice of a human activity system, we need to be able to create a dynamic model that can show how the system changes over time. In this case, we cannot use an informal description of work practice.

Figure 3-5 shows how our epistemology of work practice (described in this chapter), formalized in our Brahms modeling language and operationalized in the Brahms simulator (described in chapter 4), relates to a simulation of the work practice in a real world human activity system.

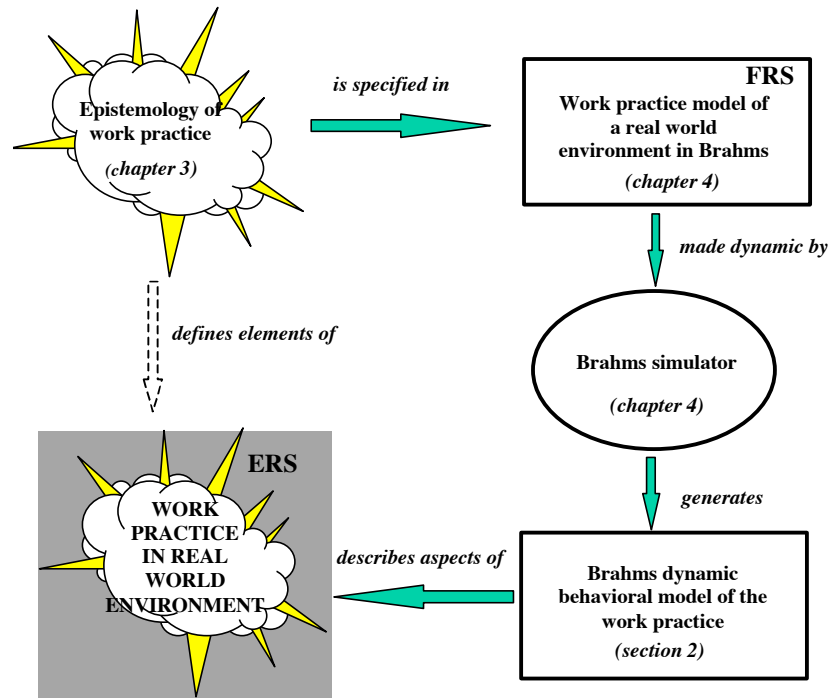


Figure 3-5. Describing real world work practice with computational modeling

3.3.2.1 The empirical relational system

The work practice in a real-world human activity system is an empirical relational system (ERS). It is empirical in the sense that it is the *source system* in which we can observe the objects and relations. The ERS refers to a group of people doing work in the real world, observed for the purpose of understanding the work practices of this group of people.

3.3.2.2 The epistemology of work practice

We observe the ERS by using the epistemological elements of work practice, described in chapter 3.2, as a sort of theoretical filter through which we view the empirical relations between the objects in the ERS. The elements of work practice we use in our filter are, again, community of practice, activities, collaboration, communication, artifacts, and geography.

3.3.2.3 The formal relational system

The elements of work practice, based on the epistemology, can be encoded into a computational Brahms model using the formal Brahms language (described in chapter 4). A Brahms model is a formal relational system with objects and relations isomorphic to real-world objects and relations in the ERS. The computational modeling language defines the formal relational system (FRS). In the FRS we describe the aspects of the work practice observed in the ERS. For each epistemological element observed in the ERS, there are formal Brahms language objects and relations that describe our observations.

3.3.2.4 The Brahms model simulator

From a computational Brahms model of the work practice a dynamic simulation model is generated, by executing the Brahms model using the Brahms simulation program (also described in chapter 4). This is the step in which the dynamic behavioral model of the work practice is generated.

3.3.2.5 The dynamic behavioral model

The behavioral model is a dynamic model in that it includes temporal activity-relations, and how they change over time.

The epistemological concepts of work practice define the theoretical basis of how to observe, capture and talk about work practice. The Brahms language operationalizes these epistemological concepts by defining a computational modeling language and a simulation program, allowing us to model and simulate a work practice from observations in the real world. Next, I describe how we develop a model of work practice.

3.3.3 Work practice models

This section describes the (sub)models of a work practice model, based on the epistemological work practice level described in section 3.2. I divide the work practice elements into related models that can be viewed independently. Dividing a model of work practice in this way helps the modeler with the decomposition of the domain, and makes the modeling effort easier.

3.3.3.1 Agent model

People are represented as *agents*. Just as people, agents do work. We can describe the work of people performing the same work, by describing the work of a *group*. Each member of the group is an *agent*, and is able to perform the work defined for the group. People can belong to multiple groups, and as such an agent can be a member of multiple groups. We represent the people in a community of practice as agents belonging to their respective groups. This way we can model any human activity system as communities of practice.

3.3.3.2 Activity model

The work that people do is described in terms of activities. *Activities* are defined at either the individual agent level, or the group level, in which case each member (agent) of the group can execute the activity. An activity represents the behavior of a person for a period of time. There are two types of activities, a *primitive activity* and a *composite activity*. A primitive activity is primitive, because it is not further decomposed, and takes some amount of time. The time element represents how long the agent is working *within* the activity. Thus, a primitive activity describes *what the agent is doing and how long the agent is doing that*. A composite activity is a higher-level activity. We can say, it is a more abstract representation of what the agent is doing. An activity can be decomposed in *sub-activities*, which can be primitive sub-activities, or again, composite sub-activities. Using primitive and composite activities we can describe what people are doing at any level of detail.

Work is the execution of activities under certain constraints. Agents' constraints for performing activities are matched against the beliefs they hold. We represent the constraints when agents can perform activities in an activity rule, called a *workframe*. A workframe defines the conditions under which the agent can execute the activity.

3.3.3.3 Communication model

We represent communication between people as an activity in which people engage when communicating with someone or something else. When communicating, people send or receive information. In our FRS, communication is represented as a type of primitive activity, called a *communication activity*. A communication activity is primitive, in that it takes a certain amount of time and is not decomposed into more primitive activities. In a communication activity we can specify what information the agent can communicate or receive, and with whom the agent is communicating when in the activity. Conditions in the workframes for communication activities specify under what circumstances an agent communicates.

3.3.3.4 Object model

People use and create artifacts in performing their activities. Artifacts are represented as *objects*. Types of artifacts, such as telephones, hammers, etc, are modeled as *classes*. New objects can be created as instances of a class. With these constructs we can model any type of artifact used within the work practice.

Some artifacts can perform activities, such as computer systems, telephones, microwaves, et cetera. Artifact behavior is represented similarly as the behavior in agents, meaning that the behavior in objects is also represented as activities and workframes.

3.3.3.5 Geography model

A human activity system is always located in some geographical space in which activities are performed. People and artifacts cannot be without location. Location also constrain when activities can be performed. For example, we cannot pick up the telephone if we're not located in the same geographical space as the telephone. We describe the location of where the agent and object's activities are performed in *geographical areas*.

Depending on the human activity system we can define types of areas with *area-definitions*, for example, buildings. A geographical area is an instance of an area-definition.

An agent and object performs its activities within only one geographical area. Moving from geographical area to area is represented as a *move activity*. A move activity is a primitive activity that takes time, and moves the agent from its current location to its new location. Using move activities we can formally describe the movements of people, during their activities.

3.3.4 Developing a model of work practice

Figure 3-6 describes an operational methodology for developing a formal computational model and a dynamic simulation model of a work practice, for an observable human activity system. A work practice is not simply the summation of the activities of all elements in the system, but it is the emergent behavior of the system as a whole, based on the interaction and collaboration between the elements in the system. Because a human activity system is about humans, we can observe *the way the humans are performing their activities*. In other words, we can observe the work practice of the system. The goal of the observation of the people in a human activity system is to create informal *static* models of the people, artifacts, the activities of those people and artifacts as they are being performed over time, as well as the geographical environment in which these activities take place.

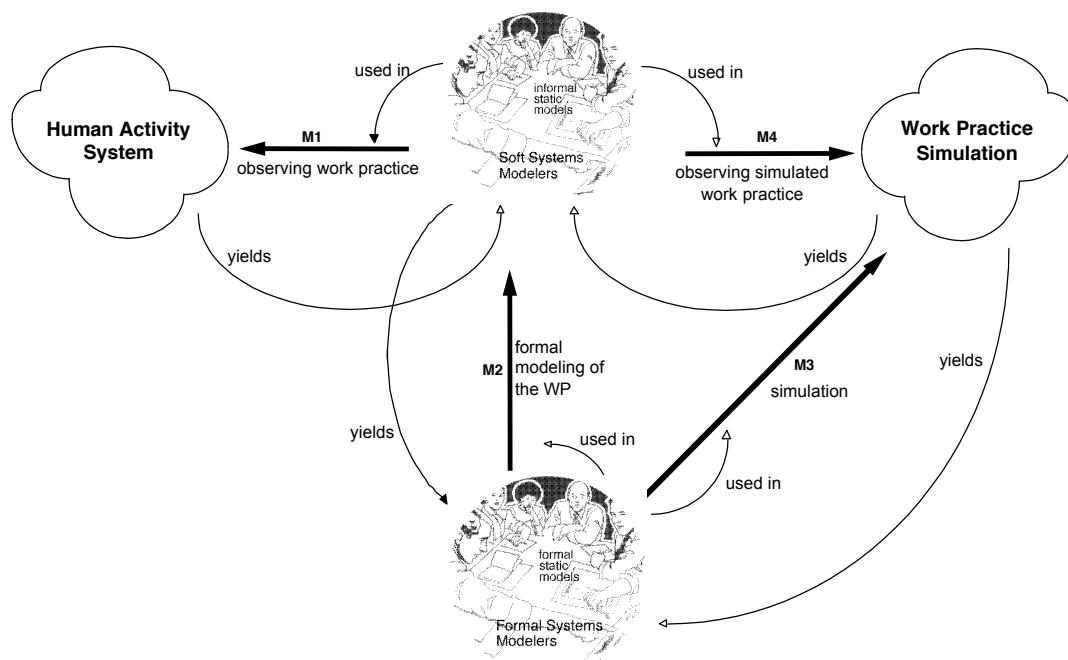


Figure 3-6. Modeling process

The empirical relational system is the human activity system being observed. The purpose of the methodology is to operationalize the modeling of the ERS, and create a Brahms model that can be executed by the Brahms simulator to create a simulation of the activities of agents and objects.

3.3.4.1 Method M1 – observing work practice

The purpose of method M1 is to observe the ERS and create an informal static description of an observation of the work practice of a human activity system. The goal of the observation is to create useful data to create static informal models, which will be used in M2 to develop formal models of work practice. There are different ways of observing a human activity system, and create data. I only mention two ways we can observe work practice in a human activity system, as examples. The first one is by analyzing video recordings of the actual work, and the second one is by using participant observation.

3.3.4.2 Method M2 – formal model of the work practice

The purpose of method M2 is to formalize the static informal models created during the application of M1, creating the FRS. In Brahms terms, this is where the Brahms model is developed. The formal system modelers need to be able to translate the informal models into formal models using a specialized kind of formal modeling knowledge. The formal modelers and the informal modelers do not necessarily have to be the same, and in fact, the skill set for these two types of modelers are very different. The informal modelers should be system analysts, knowledge engineers and anthropologists. The formal Brahms modelers should be people that understand the concept of agent-based modeling, and often have experience in developing rule-based systems.

3.3.4.3 Method M3 - simulation

The purpose of method M3 is to construct a simulation of the formal model, by running the simulator with the formal model as input and the work practice simulation as the output. The M3 method can be seen as the *model, compile, simulate, and debug cycle*.

3.3.4.4 Method M4 – observing the simulation

The purpose of method M4 is to observe and investigate the work practice simulation output, and compare it with the actual human activity system. It is during this cycle that the actual objective of the work practice simulation project is being accomplished. The result might be suggested changes to the formal model, in order to perform a what-if scenario. Thus, there is a modeling and simulation cycle between M1, M2, M3 and M4, which means that these methods have to be closely integrated if we want to make this cycle be as efficient as possible.

3.4 CONCLUSION

In this chapter, I discussed what I mean with “work practice.” In work practice modeling we focus on the collaborative activities of a community of individuals who collaborate together to accomplish a goal. I defined an epistemological framework for describing a work process at the *work-practice level*, using concepts such as collaboration, community of practice, communication, activity, and geography.

Having an informal *model of a work process at the work-practice level*¹⁸ could help us tremendously with our understanding of what is really happening within a work process. However, what has become clear from the framework is that practice is an emergent phenomenon that only shows its relationships and influences over time. Therefore, it is important not to leave out time. If we could simulate a model, we can observe how the work practice in an organization emerges. To allow for dynamics in a model, we need to make it computational. A model that is computational needs to be formal, so that it has a context-free grammar and a defined semantics. In the next chapter, I describe the formal Brahms language for modeling and simulating work practice.

¹⁸ From now on I will simply call this “modeling work practice.”